Review of a Doctoral Thesis at FIT BUT

Doctoral thesis: Automata in Software Verification and Testing

Name of the doctoral student: Martin Hruška

Name and institution of the reviewer:

Ahmed Rezine, Docent

Department of Information and Computer Science. Linköping University, Sweden.

I. Doctoral Thesis

Appropriateness and Relevance

The thesis addresses two areas in formal verification and testing of computer software: shape analysis for sequential programs and test generation for digital twins of distributed systems. Both areas can be challenging and are highly relevant for the verification and testing of modern software.

The main emphasis in the thesis is on shape analysis. This is a particularly relevant and appropriate area to conduct state-of-the-art research in formal verification. Indeed, the area is relevant as it focuses on capturing, representing, and analysing the graph-like objects (or shapes) routinely generated and manipulated by computer programs. Such objects are unavoidable in data-structures found both in standard libraries and in fine-tuned constructions whether in general programs or in specialized software as typically adopted in the kernel of any operating system. Bugs in such software are known to be particularly difficult to reason about and to detect. Conducting research in this area is appropriate, and needed, given the involved theoretical and practical challenges. From a theoretical point of view, the obtained shapes can involve arbitrary many nodes resulting in intricate structures. Being able to succinctly represent all possibly generated shapes, let alone to analyse them and to assess whether certain errors may occur, is a formidable theoretical challenge that has motivated, and continues to motivate, state-of-the-art research in software verification. Moreover, the complexity of the software and the resulting shapes, together with the sophistication of the involved techniques do require clever, systematic, and non-trivial implementation efforts. Conducting research to make practical the proposed solutions is indeed appropriate and needed.

The thesis also makes contributions to the test generation area with a focus on the generation of tests for digital twins of distributed systems. The approach of adopting digital twins to monitor, reason about, and assess the working of physical systems is already established and is gaining momentum. It is therefore important to generate tests for such systems to assess whether they operate as expected. This area is relatively new and is appropriate to conduct research in it and to transfer results from other verification and testing areas while accounting for its particularities.

A summary of the Contributions of the Thesis

At a high level, the thesis aims to improve techniques and tools used to assess the quality of software systems. For this, it focuses on shape analysis and on test generation for (digital twins of) distributed systems. The two areas share a common goal to enable assessing software quality, and a common family of formal techniques: formalizing, improving, and developing automata-based techniques to

capture, analyze and generate complex structures. Apart from these common aspects, the contributions are fundamentally different and target quite different challenges.

In the reviewer's view, the main contributions of the shape analysis part of thesis are:

- Improving the applicability and the scalability of tree-automata based shape analysis.
- Proposing key notions (e.g., compatibility, automatic splitting, folding and unfolding of boxes) to enable central operations such as precise intersection, backward analysis of a counter-example, and automatic refinement of abstraction.
- An intuitive and clear depiction of the working and capabilities of fundamentally new tree-automata based techniques that do improve on state-of-the-art and allow for the automatic verification of problems that were not possible to verify automatically.
- A description of the involved efforts to develop a competing tool in 4 instances of the software-verification competition (SVCOMP) at TACAS, the most established competition for software verification tools.
- Proposing an original template and SMT based verification technique for the analysis of linked data-structures and implementing it in a modern verification framework.
- This line of work resulted in three peer-reviewed conference papers and a book chapter.

In addition, the authors used his expertise in reasoning about finite-state automata in order:

- Propose a new technique for learning test cases from concrete traces of (digital twins of) distributed systems.
- Use automata-based abstraction techniques to generalize the family of learned test cases.
- Implement a tool to generate test cases that can actually be used by digital twins to assess the quality of their software.

The thesis indeed proposes novel techniques and tools that strictly improve on state-of-the-art for assessing the quality of particularly challenging software.

In addition to deriving and implementing a new test generation approach based on finite-state-automata abstractions, the student did contribute to the development of new techniques to maintain compatibility when automatically and elegantly generating boxes to capture non-trivial abstractions of shapes that are well beyond the current state-of-the-art. This allowed the student to derive a precise intersection approach for the formalism used to capture graphs in terms of tree-automata and to perform automatic checking and refinements of obtained counterexamples. In addition, the student concretized these complex and intricate techniques into tools that participated under several years in well established and recognized verification competitions. These participations showed the power and the proposed techniques by allowing for the automatic analysis of benchmarks that are beyond the capabilities of other automatic techniques.

Novelty and Significance:

The work encompasses the following novelties:

- The part on test generation for (digital twins of) distributed systems innovates by leveraging on an abstraction technique (collapse states sharing the same bounded language) to extrapolate and learn classes of possible trace sets. Although this can be fine-tuned, it introduces interesting techniques that can greatly simplify generation of tests for such systems.
- In addition, the work makes several innovations in the shape analysis area:

- Very few families of verification techniques can consider shapes that are as complex as those targeted by the thesis. These include (separation) logic and graph grammars. The thesis is a sophisticated improvement on a powerful family of techniques. Namely those building on Abstract Regular Model Checking. More precisely, the thesis builds on previous works on using nested tree-automata to abstract and to analyze shapes generated by pointer manipulating programs. The state-of-the-art in the field was to use sets of tree-automata to capture different shapes including (doubly) linked lists, trees or combinations of these. Finite height and language-based abstractions were recently introduced. The proposed work improves on this by introducing the notion of compatibility of forest automata. This allows the thesis to propose precise intersection and entailment operators, which combined to improved nesting of tree automata (i.e., unfolding and folding of boxes) and abstractions allows the work to propose a framework with the unique capability to perform counter-example refinement abstraction for sophisticated shapes that no other current technique can handle automatically.
- Moreover, the work proposes a promising approach to represent arbitrary graphs with tree automata. The thesis conjunctures that this novel representation allows for efficient manipulation of more general shapes than those currently handled.
- The work also proposes a new technique for analyzing pointer manipulating programs using k-induction and templates capturing may-point-to between the nodes of the generated shapes. The technique handles simpler shapes than those handled by the abstract tree forests, but it can easily account for data.
- Finally, and importantly, a strong emphasis has been placed on implementing the techniques and on competing with the tools during well-established competitions. This, and making the tools open source, certainly impact research on shape analysis and can strongly support further investigations in automata based graph representations and analysis.

Evaluation of the Formal Aspects of the Thesis:

The doctoral thesis is clearly organized into two parts: one about shape analysis and the other about test generation. There are several grammatical/spelling mistakes that should be easy to fix. There is more focus on the first part as it is the part where most work and contributions occurred. Satisfying efforts were spent in enumerating and comparing to the current state-of-the-art, especially for the shape analysis part. The second part could have taken up/comparing against techniques for learning regular languages and their applications to testing. The concepts taken up in the thesis are intuitively introduced despite their sophistication. To summarize, the thesis is well written, but some grammatical/spelling need to be fixed.

Quality of Publications

The work resulted in five technical papers at the peer-reviewed conferences VMCAI17, FMCAD18, NETYS21, EUROCAST22 and CADE23. The conferences CADE, VMCAI and FMCAD are very well-established venues among formal verification practitioners. The core of the thesis has been published at an appropriate level given the key contributions. This will give visibility to the work and to the unique investigations in adopting tree automata to capture and reason about complex shapes, but also to create a link between automata-based abstraction techniques and generating tests for distributed systems. Importantly, an impressive effort has been spent in creating and improving tools for shape analysis. The four documented participations to SVCOMP at TACAS, the most established and

respected competition for software verification tools, does give visibility to the work and showcases its potential to compete with the best approaches in the field.

II. Candidate's Overall Achievements

Overall R&D Activities Evaluation:

The thesis, the impressive amount of work involved in producing verification tools at the forefront of the current research by achieving strong contributions to an original and unique approach to carry shape analysis indicate the candidate holds strong creative and scientific abilities. The command of the landscape of current research in shape analysis and the depth of the understanding of the relation of the proposed contributions to this landscape witness the more than sufficient level of scientific erudition of the candidate.

Assessment of Other Candidate Characteristics (optional):

The reviewer would like to underscore the effort and expertise required to favorably and continuously compete, with original and non-trivial techniques based on nested tree representations of graphs, at such a well-established venue for software verification.

III. Conclusion

Given the sophistication and depth of the work, the required level of expertise, the quality and number of publications, and the quality and span of the thesis, it is the reviewer's assessment that the doctoral thesis and the student's achievements until now meet the requirements for the award of a PhD degree.

Place (Linköping/Sweden) 20.11.2023

Linköping 20/11/2023

Signature of the reviewer:

Ahmed Rezine