FACULTY
OF MATHEMATICS
AND PHYSICS
Charles University

To whom it may concern

Prague, March 1, 2023

Re: Martin Blicha, Doctoral thesis – Supervisor's reference

The submitted doctoral thesis concentrates on automated software verification. It addresses the problems of both correctness and efficiency of the software verification process, which is generally a hard – undecidable – problem, where also practical obstacles arise, especially when real-life programs are considered. The complexity of the verification usually results in huge usage of computational resources – both time and memory. Suboptimal algorithms or their implementations may then easily disable verification of non-trivial pieces of software. Here, it is of utter importance not only to carefully design and implement particular parts of the verification tools but also to pay extra attention to the composition and orchestration of the verifier's blocks.

The text of the thesis is devoted to four topics whose research results were, to a great extent, published in proceedings of renowned international conferences. The first chapter – Introduction – describes the thesis context, articulates the challenges, and summarises the contribution of each chapter in the rest of the thesis. The Preliminaries chapter contains definitions and explanations of basic concepts and principles used throughout the thesis. I consider both first chapters very nice and satisfactory.

The third chapter describes the author's contribution to the area of (Craig) interpolation in linear arithmetic. The motivation for this work is to address the problem of the diverging verification process due to the imprecise computed interpolants in some cases. The author proposes decomposed interpolants that address this issue. Decomposed interpolants are more precise and, thus, more likely to help the whole verification process converge. A deterministic linear algorithm for computing the coefficients of decomposed interpolants is presented, and its correctness is formally proven. Also, the contribution in terms of solved benchmarks is presented, showing the practical usefulness of the approach.

The fourth chapter is devoted to transition power abstractions (TPA). This novel concept, exploitable in model checking of safety properties of transition systems, is included in an algorithm also proposed in this part. TPA attempt to address the issue of so-called deep counterexamples, i.e. counterexamples reachable only after a high number of steps. Such errors in the systems are hard to find, yet harder to find and analyse their causes. TPA addresses the problem by introducing the concept of over-approximating abstrac-

Department of Distributed and Dependable Systems
Malostranské nám. 2/25, 118 00 Praha 1
phone: 951554245, e-mail: info@d3s.mff.cuni.cz

tions of sequences of exponentially growing transition series. Correctness and termination of the proposed algorithm are again formally proven.

In the fifth chapter, the Golem Horn solver is described. It is a solver for the satisfiability of constrained Horn clauses (CHC) developed (not only) to support demonstration of the author's research contributions. The tool has been designed to enable tight cooperation between model-checking front-ends and the used SMT solver. This way, the SMT-solving part of the computation can take advantage of additional pieces of information beyond the formulae themselves. The experiments in this chapter reveal the potential of Golem and evaluate its implementation in connection with several model-checking algorithms.

The sixth chapter describes a parallelisation approach to model checking involving k-induction and IC3-style search for inductive invariants. It naturally complements the results presented in the previous chapters by introducing a parallelisation layer. This way, the thesis's contribution spans several layers, from foundations through algorithms to parallelisation, resulting in a very complex approach.

The author proved his capabilities to perform high-quality research with international impact – the results were published at several top conferences in the field (some of them ranked A in CORE) and in one impacted international journal (Springer). The dissertation is written comprehensibly, and despite the technical details and the broad spectrum of levels included, it is very nice to read. The presented results are clearly above an average dissertation and more than sufficient to prove the author's ability to perform independent, high-quality research.

Regarding the dissertation itself, I have no comments. I would like the author to address the following questions during the defence:

1. In the context of TPA, the base for the transition sequence length is chosen as two. In the section about possible future work, you mention that other options may also be useful. Can you elaborate a little on the effects of choosing different bases? Especially in terms of abstraction precision and error detection speed.

2. In the case of Golem, is there any particular research/implementation direction (listed in the future-work section) you plan to take first?

Given my assessment of the dissertation of Martin Blicha above, I strongly recommend the thesis for defence and to grant the Doctor degree to him.

Jan Kofroň
supervisor