Prof. Dr. Jörg Hoffmann
Saarland University
Saarbrücken, Germany
hoffmann@cs.uni-saarland.de

SAARLAND UNIVERSITY
COMPUTER SCIENCE

March 10, 2015

Review of PhD Thesis by Martin Suda

To Whom it May Concern.

I am writing this in review of the doctoral thesis by Mr. Martin Suda, supervised by Christoph Weidenbach, Petr Stepanek, and Roman Bartak.

Overall, in my view it is completely clear that the presented material is suitable. The amount of results included in the thesis is impressive. Mr. Suda starts off with extending a reasoning calculus from propositional logic to LTL (Chapter 2); developing a practically effective guided version of that calculus (Chapter 3); devising pre-processing techniques in that context (Chapter 4); deeply relating the method to property-directed reachability, a recently successful method that at first looks rather different, and devising an enhancement as well as running extensive experiments in hardware verification (Chapter 5); adapting property-directed reachability to planning, as yet another related area, and running extensive experiments there (Chapter 6). Each of these work steps is substantial. Several of these steps (especially Chapters 2, 3, and 5) involve highly non-trivial conceptual and theoretical developments. All except Chapter 2 involve sizable implementation works and experiments. Throughout, the work is excellently conducted, is of high technical quality, and is clearly and comprehensively described.

In terms of original ideas and empirical results contributed, the strongest contributions of the thesis strike me as a new highly competitive LTL reasoner (Chapters 2–4) which seems to be able to combine most of the good results of other solvers in one; a new and potentially fruitful perspective on property-directed reachability as well as an optimization that significantly pays off in multi-property testing (Chapter 5); and property-directed reachability as an almost completely new approach to planning that, in this very first incarnation, already is very competitive with the state of the art in

satisficing planning (Chapter 6). Each of these is quite remarkable on its own, and more so is their combination within a single PhD thesis.

A very strong characteristic of the thesis, which I applaud, is a drive to "dig deep", trying to view the methods developed from different angels, and understand their behavior as well as interestingly relate it to that of competing methods. The thesis attests to creativity, with the author consistently coming up with the right ideas for applying a given method in a new context. The thesis finally attests to stamina, the author "traveling" a very long but coherent journey from LTL reasoning over a relation to hardware verification methods to AI planning.

I have very little to criticize. The language is good but contains a few typos etc. I would have liked to see some explicit data about the reduction effects in Section 4.3.2. Clearly these are quite unimportant. My only somewhat important concern is with the empirical results. While often good, they do not really indicate a "breakthrough", it's more like, being on par with previous methods or combining some of their virtues. While that is a bit unsatisfying, it is (a) expected in competitive areas such as those addressed, and (b) does in no way at all harm the scientific soundness and importance of the work presented.

In summary, I am quite happy with the thesis. It is more than worthy of the PhD degree, and I evaluate it as "ausgezeichnet (summa cum laude)".

Best Regards,

Jörg Hoffmann