UNIVERSITY of VIRGINIA | SCHOOL of ENGINEERING & APPLIED SCIENCE
Department of Computer Science

Matthew B. Dwyer
Professor, John C. Knight Faculty Fellow
Department of Computer Science
PO Box 400740University of Virginia
Charlottesville, VA 22904
(434) 982-210

August 25, 2019

Professor Jan Trlifaj
Vice-Dean
Faculty of Mathematics and Physics
Charles University
Prague
Czech Republic,

Dear Professor Trlifaj,

In this letter I evaluate the scholarly contributions of Dr. Pavel Parízek and assess the quality of those contributions and provide my opinion on whether his work meets the standards for promotion to the rank of Associate Professor.

I have known of Dr. Parízek's work since reading a 2012 OOPSLA paper that he authored on the topic of predicate abstraction. I was the program committee chair for that scientific meeting and recall that the conversation among experts, who were unanimously positive about the work, inspired me to read the paper myself. I was impressed and this has led me, over the years, to keep track of Dr. Parízek's work in the intervening years.. I do not have a personal relationship with him nor have I ever collaborated with him.

In reviewing Dr. Parízek's habilitation thesis, I found papers that I had already read and some that I was not familiar with. The two major areas that he identifies in his thesis are areas in which I have published several papers – including papers that Dr. Parízek's work builds on. Consequently, I believe that I am well-qualified to assess the quality of the work.

The first theme in Dr. Parízek's thesis extends, in multiple directions work that I published with colleagues in 2004 that sought to improve the accuracy of the *independence relation* that underlies partial order reduction (POR) optimizations in model checking of concurrent programs. Our idea was to shift the calculation of this relation to runtime so that precise information about the state could be exploited. The technique we developed is the method implemented in the JavaPathFinder (JPF) model checker upon which much of Dr. Parízek's work is based. Dr. Parízek goes further in several of the papers pursuing this theme, but computing ahead of time, using static analyses, information that can be consulted at runtime to improve POR.

In the ASE 2011 paper, co-authored with Dr. Lhoták from the University of Waterloo, they tackle the challenge that POR of heap manipulating programs must be able to reason about future field accesses to safely determine the independence relation. Our technique was sound, but conservative, and Dr. Parízek's insight was that soundness can be perserved while relaxing the conservative approach by computing a sound approximation of the fields that will be accessed by a thread at a given program point. This is analagous to a *live variable* analysis, but it is much more complex because it requires precise points-to information to track the sets of objects pointed to by a reference and it is naturally field-sensitive. The paper develops this idea and describes it unusually clearly. Moreover, they implement a prototype, integrate it with JPF, and conduct an evaluation on artifacts used in model checking studies of the time. Their findings conclusively demonstrate the value of the approach.

The VMCAI 2016 paper, solely authord by Dr. Parízek, tackles the challenge of extending the above strategy to array accesses. They key challenge here is the need to soundly, and accurately, compute an estimate of the values that can arise in an index expression at any subsequent point in the program's execution. The paper makes the clever insight that a focused symbolic analysis can work in this setting. Typically these are too expensive to incorporate into a static analysis that targets the entire program, but the nature of this application provides sufficient focus to restrict the cost to a reasonable time. This work follows the pattern of the ASE 2011 paper in presenting a clear explication, describing and prototype tool, and conducting an evaluation that demonstrates the benefits of the approach.

I'll comment on two aspects of the above papers. First, the ASE and VMCAI conferences conduct extremely rigorous review processes and are highly-selective – they are both considered top conferences in the field. Second, it is unusual to see a sole-authored paper these days in such a prestigious venue and I view it as a marker that Dr. Parízek is truly an independent researcher capable of world-class research. Third, these papers led to deeper followup work by him, as reflected by the papers comprising the later chapters in the thesis, and others, as reflected by citations to his work. Finally, in the two papers above, and in others by Dr. Parízek, I noticed an unusual degree of thoroughness in the design of the experimental studies. In fact, these studies constitute what is now termed an *ablation study* – where a technique built up out of multiple individual components is evaluated under the powerset of combinations of those components to demonstrate their unique contribution. These types of studies are now expected in the field, but Dr. Parízek was more than 7 years ahead of the research community in this regard and I view this as a sign of his maturity and leadership as a scholar.

I found Dr. Parízek's work on fault detection to be quite interesting as well – what he notes as the second theme in his thesis. In 2007 I published a paper at ICSE that used randomization to improve the fault-detection effectiveness of JPF, but here again Dr. Parízek took the work further. We looked just at randomizing the order in which enabled transitions were explored, but in Dr. Parízek's SPIN 2011 paper, and in his followup work, he took a broader view. In addition to randomizing the order of exploration his technique introduces an extra alternative at each state – backtracking. While SAT solvers used this heuristic before, Dr. Parízek's is the first to apply it to state-space search and fault detection. When I read this paper I had one of those "I wish I had thought of that" experiences. It is an idea that I immediately knew would improve on our approach and after reading Dr. Parízek's careful development of the ideas and evaluation of a prototype, it does.

In the two thematic areas that Dr. Parízek's thesis develop he has clearly advanced the state-of-the-art. The techniques have been implemented and evaluated thoroughly. Moreover, the implementations have been shared with the international community to promote further advances and others have built on them. I was impressed when I first read Dr. Parízek's OOPSLA 2012 paper, which is not included in the thesis, and after reading through the thesis I am significantly more impressed with the depth and quality of his work.

In conclusion, after reviewing the research described in Dr. Parízek's habilitation thesis, it is my assessment that he meets the standards for promotion to Associate Professor.

If you have any questions about my analysis of Pavel's record I would be happy to answer them.

[redacted]

Matthew B. Dwyer

Biography: I am a Professor in the Department of Computer Science at the University of Virginia where I am the John C. Knight Faculty Fellow. I have published more than 120 refereed scholarly papers since 1995 (when I graduated from the University of Massachusetts). I have served as principal investigator on more than 30 competitively funded research projects. I have had the good fortune of working with some very good students and have graduated 9 Doctoral students (4 who hold faculty positions and two who have received significant awards as independent researchers, e.g., NSF CAREER). I have authored papers that have received 4 best paper awards and 3 "test of time" awards. I have chaired the top research conferences in the field of Software Engineering (FSE in 2004, ICSE in 2008, OOPSLA in 2012) and served as Editor-in-Chief of the top journal in the field (IEEE TSE). I was named a Distinguished Scientist by ACM in 2007, named a Fulbright Research Scholar in 2011, and named an IEEE Fellow in 2013.