IRISA

IRISA • Campus universitaire de Beaulieu • 35042 Rennes Cedex France • Tél. : +33 2 99 84 71 00 • Télécopie : +33 2 99 84 71 71 • Internet : www.irisa.fr

Report on the Dissertation Submitted by Konrad SIEK to Obtain the the PhD Degree in Informatics from Poznań University

The work of Konrad SIEK to obtain a PhD in Informatics was done at the *Faculty of Computer Science* of the *Poznań University of Technology*. This work was realized under the supervision of Professor Pawel T. Wojciechowski.

The document is titled "*Distributed Pessimistic Transactional Memory: Algorithms and Properties*". This document considers the transactional memory (TM) approach to address the synchronization issues encountered when writing distributed applications on top of multicore processors. The aim of a TM system is to simplify the writing of parrallel and distributed programs, by requiring the programmer to specify only which parts of the algorithms associated with each processor have to be executed atomically. These parts of code are called *transactions*. In such a context, Konrad SIEK promotes the *pessimistic approach* in which a transaction never abort on conflict, but may be delayed so that conflicts never materialize.

Content of the document

- After a first chapter introducing the problem and offering a global view of the work, three chapters focus on defining the context of the study. These chapters constitute an in-depth presentation and precise analysis of the TM context.
  - Chapter Two defines the fundamental concepts and objects found in transactional memory (TM) systems, namely, processes, transactions, legality, execution history, safety, locks, early release, and progressiveness. These definitions relie on a formalism, which makes the definitions very precise and unambiguous.
  - Chapter Three discusses TM safety properties, some coming from database such as serializability, strictness, and rigorousness, others originating from the TM area such opacity, virtual world consistency, live opacity, and elastic opacity. The chapter also analyses if these properties allow transactions to use the "early release" technique. Simple examples are used to ease the understanding, and a table at the end of the chapter gives a nice summary of this analysis.
  - Chapter Fourth presents distributed and non-distributed TM algorithms using the pessi-

mistic transactional approach, the optimistic, and the early release technique. Basically, the author reviews transactional two-phase locking, versioning algorithms, Shavit's transactional locking II, and the non-distributed Matveev-Shavit's Pessimistic algorithm. Each algorithm is carefully presented anand analyzed with meaningful examples. A table at the end of the chapter summarizes its technical content.

- The second part of the dissertation is made up of five chapters, which constitutes the technical core of the document and the new technical results due to the author.
  - Chapter Five presents a TM property called *Last-use*. This is a new TM safety property that provides strong consistency guarantees without compromising early release. This property relies on the notion of a *closing* write operation. In addition to formal definitions, examples are described which motivate and give the intuition that underlies this new TM property.
  - Chapter Six presents a family of new pessimistic TM algorithms. These algorithms are based on a versioning technique, and the Last-use property. Each algorithm is clearly explained and and its properties analyzed.
    Then, Chapter Seven presents safety proofs of some of of the previous algorithms. These proofs are not "sketch of proof" but formal proofs. Their reading requires some efforts, but this is rewarding as their correctness relies on precise definitions and rigorous reasonning.
  - Chapter Eight and Nine are on the implementation side. Chapter Eight is on the implementation side from an "execution point of view". It presents two distributed implementations of versionning algorithms. These implementations are "complete" in the sense that their API is precisely defined. They are evaluated with appropriate benchmarks, and a lot of (convincing) evaluation figures are included in the document.
    Chapter Nine is on the implementation side from a "pre-execution (compilation) point of view". It presents a precompiler which relies on a static analysis algorithm, which derives a priori information which can then be used by the TM algorithms.
- Lastly, a conclusion (Chapter ten) provides a short abstract of the work and presents directions for future works.

My point of view

- **Problem and impact**. Transactional memory is a very important topic when one wants to address synchronization issues in the context of multticore architecture. Lot of results are in the litterature, but lot of work remains to be done on both its theoretical side and its practical side. Hence Konrad ZIEK'work is really both a relevant and timely work.
- **Contribution and Correctness**. The work of Konrad ZIEK is on a difficult and fundamental topic of synchronization, which is one of the most important topic of Informatics. In a very interesting way, its contributions belong to two sides. One is the theoretical side: introduction of a new relevant TM safety property and design of TM algorithms; the other is on the practical side: implementation of a TM system and in-depth analysis of it.
  On the correctness side, Konrad ZIEK made a big effort to have precise definitions, rigorous

proofs, and motivated implementation analysis. This is great.

- **Knowledge of the candidate.** As demonstrated by chapters One-Fourth, and chapter Nine, Konrad ZIEK had a deep knowledge not only of Synchronization but of many domains of Informatics (including very basics domains such as language theory, compilation, and operating systems).
- **Conclusion.** Answers to questions as required by Article 13 of the Act of March 14, 2003, of the Polish Parliament on the Academic Degrees and the Academic Title.
  - Question: Does this dissertation present an original solution to a scientific problem?
    Definitely yes.
  - Question: After reading the dissertation, would you agree that the candidate has general theoretical knowledge and understanding of the discipline of Computing and particularly of Parallel and Distributed Systems?
    Definitely yes.
  - Question: Does the dissertation support the claim that the candidate is able to conduct scientific work?
    Definitely yes.

To conclude

According to the *high scientific quality* of the document (which is very well presented) produced by Konrad SIEK, I give a very positive recommendation for the defense of his work to obtain the PhD Degree in Informatics, and I recommend to distinguish this dissertation for its quality.

August 23, 2016

Professor Michel Raynal

Member of Academia Europaea
Senior Member, Institut Universitaire de France

IRISA, Université de Rennes (France)

Adjunct Professor, Polytechnic University, Hong Kong

http://www.irisa.fr/prive/raynal/
(see also DBLP "raynal")