Prof. Fernando Pedone
Faculty of Informatics
University of Lugano (USI)
Switzerland
fernando.pedone@usi.ch

Lugano, February 14, 2017

Reviewer's opinion
on Ph.D. dissertation authored by

Tadeusz Kobus

entitled:

Transactional Replication: Algorithms and Properties

1. Problem and its impact

This thesis addresses the general problem of replication in the context of transactions. Transactions are among the most fundamental abstractions in computing systems. Transactions simplify the design of otherwise complex applications by encapsulating application logic in smaller units with well-defined properties. Replication is a technique to ensure that computing systems can tolerate failures without service disruption, a property known as availability. Today's computing systems are expected to be always available, despite the failure of some of its components. Replication is the technique that allows this property to be implemented. Therefore, understanding how to implement transactions and replication efficiently is of utmost importance.

Although replication has been a topic of extensive research in the last forty years, designing replicated systems that ensure strong consistency (a property typically provided by transactions) is challenging and many tradeoffs must be addressed. Moreover, novel environments require revisiting existing techniques. For example, cloud computing allows deployments that comprise dozens of replicas, as opposed to the few replicas that most replication protocols targeted in the past. This thesis takes a fresh look at the two most important replication techniques used with transactions: active replication (also known as state machine replication) and multi-primary passive replication (also known as deferred update replication). Both techniques are at the core of many current systems, including coordination services, databases, file systems, among many more.

There is no doubt that the methods used in the thesis are scientifically sound and the contributions have important practical implications.

2. Contribution

The thesis makes three contributions. First, it extends correctness criteria for replicated schemes. In particular, it proposes two families of properties: $\diamond$-opacity and $\diamond$-linearizability. These extensions capture more precisely the differences in the ways state machine replication and deferred update replication can be implemented. Understanding precisely what replication techniques and optimizations ensure is the basis for comparing them at a fundamental level. This is clearly an important contribution.

The second contribution of the thesis is a comparison between state machine replication and deferred update replication. The thesis describes the techniques at a high level, characterizes them precisely in terms of their properties, reasons about the correctness of algorithms, and assesses different protocols experimentally. The thesis also proposes an extension of state machine replication to accommodate "cheap reads", that is, read operations whose execution does not involve all the replicas, and therefore lead to better performance than classical state machine replication. This part of the thesis identifies workloads and conditions that characterize a tradeoff involving the two replication schemes. State machine replication is abort-free, but requires update transactions to be executed at all replicas. With deferred update replication, each transaction is executed by one replica only and then the execution is certified by all replicas and committed if the transaction passes certification. The thesis shows for which workloads one technique is more efficient than the other.

The third contribution of the thesis is a hybrid replication technique that exploits the strengths of state machine replication and deferred update replication. The idea is to provide both execution schemes and use the one most appropriate for the transaction in execution. For example, transactions that content for data items should be executed using state machine replication since these transactions will not be aborted. The thesis proposes two different ways to switch between the two replication schemes, one scheme that requires the intervention of the system administrator and one fully automated mechanism that uses machine learning techniques. In my opinion this is the most important contribution of the thesis. Since the proposed approach is relatively simple, it could perhaps be integrated in a real database system and result in significant performance improvements.

Although there is no dispute that the thesis makes important contributions, one can also consider other quality indicators. In particular, the candidate has authored several excellent publications and patents. In relation to this thesis, the candidate has published in two prestigious journals, IEEE Transactions on Parallel and Distributed Systems and Elsevier Journal on Parallel and Distributed Computing, and several excellent conferences, including IEEE Symposium on Reliable Distributed Systems (SRDS) and IEEE International Conference on Distributed Computing Systems (ICDCS).

Finally, the candidate has also taken part in other research activities that led to important publications, including a brief announcement at the ACM Symposium on Principles of Distributed Computing (PODC) and papers at the IEEE Symposium on Reliable Distributed Systems (SRDS) and International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM).

3. Correctness

The thesis strikes a balance between theory and systems, something that is commendable and not often seen. All claims are supported by rigorous correctness proofs. At the same time, the candidate has detailed implementations of the proposed protocols and extensive empirical evaluations. The experimental setup is carefully described and the results are interpreted in detail. I must also add that the document is very well organized and the writing style quite clear.

I believe we can trust the claims in the dissertation and the arguments seem correct to me. I did not find any flaws in the reasoning. Among the aspects I value the most in this thesis are the balance between theory and practice, the rigour of the arguments and the depth of the analysis of empirical data. I would recommend that the candidate make the software used in the experiments available as open source so that other researchers and practitioners in the field can benefit from it.

4. Knowledge of the candidate

The second chapter of the thesis presents the state of the art in the field and covers replication schemes, correctness properties for transactional and non-transactional systems, distributed transactional memory, and other aspects related to the work. In addition to the material in this chapter, other chapters also provide additional details about related work. This indicates that the candidate has general knowledge in the discipline of **Computing**.

5. Other remarks[1]

To conclude, I really enjoyed reading this thesis. Replication in distributed systems and distributed data management systems is my main topic of research, and rarely I have seen works that combine the two. Besides, the thesis approaches these aspects in a very clear way.

6. Conclusion

Taking into account what I have presented above and the requirements imposed by Article 13 of *the Act of 14 March 2003 of the Polish Parliament on the Academic Degrees and the Academic Title* (with amendments),[2] my evaluation of the dissertation according to the three basic criteria is the following:

A. Does the dissertation present an original solution to a scientific problem? (the selected option is marked with **X**)

[X]	[ ]	[ ]	[ ]	[ ]
Definitely YES	*Rather yes*	*Hard to say*	*Rather no*	*Definitely NO*

B. After reading the dissertation, would you agree that the candidate has general theoretical knowledge and understanding of the discipline of **Computing**, and particularly the area of **Software Engineering**?

[X]	[ ]	[ ]	[ ]	[ ]
Definitely YES	*Rather yes*	*Hard to say*	*Rather no*	*Definitely NO*

C. Does the dissertation support the claim that the candidate is able to conduct scientific work?

[X]	[ ]	[ ]	[ ]	[ ]
Definitely YES	*Rather yes*	*Hard to say*	*Rather no*	*Definitely NO*

Moreover, taking into account all the considerations above, in particular the balance between theory and systems, and the quality of the publications, I **recommend to distinguish** the dissertation for its quality.[3]

Signature

[1] Optional
[2] http://www.nauka.gov.pl/g2/oryginal/2013_05/b26ba540a5785d48bee41aec63403b2c.pdf
[3] Obviously, this sentence is optional.