

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

Jan Kończak

entitled:

*Recovery algorithms in State Machine Replication with
volatile and non-volatile main memory*

1. Problem and its impact

Online systems and services are the core of the digital economy, an area that has experienced tremendous growth in recent years. To be of practical use, virtually all online services must withstand failures of the components that make up the environment in which these services are deployed. One central aspect is tolerating the failure of servers. This thesis discusses efficient ways to recover failed servers.

In more detail, the thesis considers recovery in the context of state machine replication, a replication technique that provides strong consistency, that is, the technique can be used with applications that are designed to be executed by a single server, without consideration for how failures should be handled. Arguably, this is the most useful approach to fault tolerance as it spares application designers from the complexity of tolerating failures in a distributed environment. Although the thesis focuses on Paxos, a particular way to implement state machine replication, the results presented in the thesis can be extended to other contexts.

Even though state machine replication and Paxos have been largely considered in the literature, recovery is a topic that has received relatively low attention, despite its importance. Fault-tolerant systems in general and state machine replication in particular introduce some level of redundancy (e.g., server replicas) to be able to continue to operate despite the crash of some servers. It is of utmost importance to consider how to reintegrate failed servers to the ensemble. Therefore, the thesis discusses a very important problem. The approach used in the thesis is sound and follows standard practices in systems research (i.e., detailed algorithms carefully evaluated experimentally). Finally, the results presented in the thesis clearly have the potential for practical impact.

2. Contribution

The thesis claims two main contributions to support crash recovery in state machine replication, one where a majority of replicas is assumed to be always operational, and another one that does not make this assumption (i.e., all replicas may crash simultaneously). These are the assumptions made by most proposals of interest. In this context, the thesis considers state recovery from operational replicas (i.e., volatile main memory) and state recovery from stable storage. The thesis proposes different algorithms in each case, some of which make use of modern technologies (i.e., non-volatile main memory). The

thesis also proposes as a contribution the experimental evaluation of these techniques and compares their performance to established approaches.

The contributions claimed in the thesis are valid. In addition to these, I could add the survey of related techniques provided in the thesis. Moreover, the candidate has several publications in top conferences and journals.

3. Correctness

The thesis follows a “systems approach”, that is, the typical approach used in the systems community in which a practical problem is identified (i.e., in the case of this thesis the recovery of failed replicas in Paxos-based state machine replication), one or more solutions are proposed (i.e., recovery protocols under different memory assumptions), and the proposed solutions are experimentally evaluated and compared. The thesis characterizes the proposed protocols under different circumstances. The experimental part of carried out carefully, which provides confidence in the results reported in the thesis. The protocols proposed are described in sufficient level of detail. On the negative side, a more formal argument for the correctness of the proposed protocols would be welcome. However, in my opinion, the lack of such formal argument does not diminishes the quality of the thesis.

4. Knowledge of the candidate

The thesis is structured in eight chapters, where the first and last contain the usual introduction and conclusion. While Chapter 4 explicitly mentions the state of the art in the area of the thesis, Chapters 2, 3 and 5 also contain tutorial information that confirm the candidate’s general knowledge in his discipline. Chapter 2 contains an overview of fundamental aspects of distributed systems, Chapter 3 delves into Paxos specifics, an area known for its inherent complexity, Chapter 5 overviews persistent memory. I considered all these chapters of very good quality and fairly complete.

5. Other remarks¹

Overall, I am satisfied with this thesis. It identifies a real issue in replicated systems that has received little attention, proposes solutions that make use of modern technologies, and evaluates the proposed contributions carefully. The text is well written, and the candidate has proved to be able to publish in top journals and conferences.

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)

Definitely YES

Rather yes

Hard to say

Rather no

Definitely NO

¹ Optional

² http://www.nauka.gov.pl/g2/oryginal/2013_05/b26ba540a5785d48bee41aec63403b2c.pdf

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

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?

Definitely YES

Rather yes

Hard to say

Rather no

Definitely NO

Moreover, taking into account the quality of the text and the significance of the contributions, I **recommend to distinguish** the dissertation for its quality³.


Signature

³ Obviously, this sentence is optional.