Boosting STM Replication via Speculation*

Paolo Romano[1], Roberto Palmieri[2], Francesco Quaglia[2], and Luis Rodrigues[1]

[1] INESC-ID, Lisbon, Portugal
[2] DIS, Sapienza University Rome

Abstract

Over the last years, several papers have addressed the issue of how to enhance both performance and dependability of Software Transactional Memories (STMs) using replication techniques, e.g. [1, 4, 2, 13]. Since STM and Database management Systems (DBMS) share the notion of atomic transaction, quite unsurprisingly, the vast literature on database replication, e.g. [9, 8, 5], has represented a natural source of inspiration also for the design of replicated STMs.

Unfortunately, in typical STM settings, the transaction execution time is up to two or three orders of magnitudes than in database environments [10]. Unlike conventional DBMS transactions, in fact, STM transactions only access in memory data items, hence not incurring in the latencies associated with disk accesses. Also, in STMs the overhead of SQL parsing and plan optimization are absent, which further contributes to shortening the transaction lifetime.

As a direct consequence, the relative overhead introduced by classical database replication schemes, when employed in STM systems, is correspondingly amplified, leading to severely performance degradations.

In order to address these efficiency issues, we have recently proposed a new approach to replication of STMs, which we named speculative replication [11]. The speculative replication algorithms that we have proposed so far [12, 6, 3, 7] explore a number different trade-offs, but all of them are based on two key common principles:

1. minimizing replication overhead by overlapping the replica synchronization and transaction execution phases, and
2. maximizing robustness in presence of mispeculations by exploring multiple transaction serialization orders.

In this talk, we will present a unified view of the speculative STM replication protocols that we have published so far, illustrating how the existing solutions fit in the whole design space of speculative replication protocols, and identifying which regions of this design space are currently unexplored in order to stimulate further research in this challenging area.

* This work has been partially supported by the COST Action IC1001 (Euro-TM), and by FCT (INESC-ID multiannual funding) through the PIDDAC Program Funds.

References

1. M. K. Aguilera, A. Merchant, M. Shah, A. Veitch, and C. Karamanolis. Sinfonia: a new paradigm for building scalable distributed systems. In *SOSP '07: Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, pages 159–174, New York, NY, USA, 2007. ACM.
2. N. Carvalho, P. Romano, and L. Rodrigues. Asynchronous lease-based replication of software transactional memory. In *Proc. ACM/IFIP/USENIX Middleware*. Springer-Verlag, 2010.
3. N. Carvalho, P. Romano, and L. Rodrigues. Scert: Speculative certification in replicated software transactional memories. In *Proc. of ACM SYSTOR*, 2011.
4. M. Couceiro, P. Romano, N. Carvalho, and L. Rodrigues. D^2STM: Dependable Distributed Software Transactional Memory. In *Proc. of PRDC*. IEEE Computer Society Press, 2009.
5. B. Kemme, F. Pedone, G. Alonso, and A. Schiper. Processing transactions over optimistic atomic broadcast protocols. In *Proc. of the 19th IEEE International Conference on Distributed Computing Systems*, page 424, Washington, DC, USA, 1999. IEEE Computer Society.
6. R. Palmieri, F. Quaglia, and P. Romano. Aggro: Boosting stm replication via aggressively optimistic transaction processing. In *Proc. of NCA*, pages 20–27, Los Alamitos, CA, USA, 2010. IEEE Computer Society.
7. R. Palmieri, F. Quaglia, and P. Romano. Osare: Opportunistic speculation in actively replicated transactional systems. In *Proc. of SRDS*, 2011.
8. M. Patino-Martínez, R. Jiménez-Peris, B. Kemme, and G. Alonso. Scalable replication in database clusters. In *Proc. of the 14th International Conference on Distributed Computing*, pages 315–329, London, UK, 2000. Springer-Verlag.
9. F. Pedone, R. Guerraoui, and A. Schiper. The database state machine approach. *Distributed and Parallel Databases*, 14(1):71–98, 2003.
10. P. Romano, N. Carvalho, and L. Rodrigues. Towards distributed software transactional memory systems. In *Proc. of the Workshop on Large-Scale Distributed Systems and Middleware*, Sept. 2008.
11. P. Romano, R. Palmieri, F. Quaglia, N. Carvalho, and L. Rodrigues. Brief announcement: On speculative replication of transactional systems. In *Proc. of SPAA*, 2010.
12. P. Romano, R. Palmieri, F. Quaglia, N. Carvalho, and L. Rodrigues. An optimal speculative transactional replication protocol. In *Proc. of ISPA*, volume 0, pages 449–457, 2010.
13. M. M. Saad and B. Ravindran. Supporting stm in distributed systems: Mechanisms and a java framework. In *Proc. of TRANSACT*, 2011.