The Prophecy of Timely Rollback*

Martín Abadi

Google
Mountain View, California, USA

Abstract

Techniques for rollback recovery play a central role in ensuring fault-tolerance in many distributed systems [5]. This talk addresses the formal specification and analysis of those techniques. In particular, we will discuss the relevance of prophecy variables [4] (auxiliary program variables whose values are defined in terms of current program state and future behavior) to reasoning about systems with undo operations [1]. We will then focus on a model for data-parallel computation with a notion of virtual time [6, 2]. In this model, rollbacks allow the selective undo of work at particular virtual times [3]. A refinement theorem ensures the consistency of rollbacks.

This talk is largely based on joint work with Michael Isard.

1998 ACM Subject Classification F.3.1 Specifying and Verifying and Reasoning about Programs

Keywords and phrases Dataflow, refinement, rollback

Digital Object Identifier 10.4230/LIPIcs.CSL.2015.1

Category Invited Talk

References

1 Martín Abadi. The prophecy of undo. In Alexander Egyed and Ina Schaefer, editors, *Fundamental Approaches to Software Engineering – 18th International Conference, FASE 2015, Proceedings*, pages 347–361. Springer, 2015.

2 Martín Abadi and Michael Isard. Timely dataflow: A model. In Susanne Graf and Mahesh Viswanathan, editors, *Formal Techniques for Distributed Objects, Components, and Systems – 35th IFIP WG 6.1 International Conference, FORTE 2015, Proceedings*, pages 131–145. Springer, 2015.

3 Martín Abadi and Michael Isard. Timely rollback: Specification and verification. In Klaus Havelund, Gerard Holzmann, and Rajeev Joshi, editors, *NASA Formal Methods – 7th International Symposium, NFM 2015, Proceedings*, pages 19–34. Springer, 2015.

4 Martín Abadi and Leslie Lamport. The existence of refinement mappings. *Theoretical Computer Science*, 82(2):253–284, 1991.

5 E. N. Elnozahy, Lorenzo Alvisi, Yi-Min Wang, and David B. Johnson. A survey of rollback-recovery protocols in message-passing systems. *ACM Computing Surveys*, 34(3):375–408, 2002.

6 Derek Gordon Murray, Frank McSherry, Rebecca Isaacs, Michael Isard, Paul Barham, and Martín Abadi. Naiad: a timely dataflow system. In *ACM SIGOPS 24th Symposium on Operating Systems Principles*, pages 439–455, 2013.

[*] Most of this work was done at Microsoft Research.

24th EACSL Annual Conference on Computer Science Logic (CSL 2015).
Editor: Stephan Kreutzer; pp. 1–1
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany