

Preface

This book is the result of the seminar “*A 30-Year Perspective on Replication*,” which took place at Monte Verità, Ascona, Switzerland, in November 2007. As suggested by the title, the goal of the seminar was not to speculate about the future of replication, but rather to understand the present, by analyzing past successes and past failures, and to make an assessment of about 30 years of research on replication. Replication is a topic addressed by several communities: the distributed computing community, the distributed system community, and the database community. Each of these communities has looked at replication from different points of view and with different goals, e.g., performance vs. fault tolerance. Recently, these different goals have started to converge, and there has been work showing that efficiency and strong consistency can sometimes be reconciled.

During the seminar the observation was made that we had reached a point of understanding of the different issues of replication, and this knowledge should be materialized in a book covering the different aspects of replication. This book results from this observation. Its goal is to present a comprehensive view of the achievements of 30 years of research on replication. The book was written by most of the people who have contributed to developing the state-of-the-art replication techniques. It brings a comprehensive view of existing solutions, from a theoretical as well as from a practical point of view. It covers replication of processes/objects and of databases; replication for fault tolerance and replication for performance; benign faults and malicious (Byzantine) faults. By covering these different issues in an integrated way, we believe the book fills a gap, and as such it should find a place in the graduate teaching of distributed computing, distributed systems, and databases.

The book is organized in thirteen chapters. Chapter 1 introduces consistency models for replicated data, both in the context of process/object and database replication. Chapter 2 discusses replication techniques commonly used in process replication, focusing on primary back-up and related techniques; it considers both the fail-stop and the crash failure models. Chapter 3 considers modular approaches to process replication; it starts with state-machine replication based on atomic broadcast and shows how this can be built on top of consensus. Although the literature on consensus is vast, there are many misunderstandings, often involving dif-

ferent communities. Chapter 4 discusses these misunderstandings. Chapter 5 covers replication for performance; it contains different strategies and examples, and discusses trade-offs. Chapters 6 and 7 provide a historical account of the Virtual Synchrony Replication Model and Viewstamped Replication, two early replication systems and how they have evolved over the years. Chapters 8 and 9 are dedicated to state-machine replication with Byzantine faults; the first considers distributed trust systems, and the second introduces protocols for state-machine replication. Chapter 10 surveys Byzantine quorum systems, suitable for use when parts of the system cannot be trusted. Chapters 11 through 13 consider database replication. Chapter 11 bridges the gap between process/object replication and database replication, while Chap. 12 surveys database replication techniques; it discusses different replication approaches, consistency criteria for replicated databases and existing systems. Chapter 13 illustrates database replication with a case study: the details of an architecture for practical database replication.

Each one of the chapters in the book is self-contained, and can be read individually. Readers interested in certain specific aspects of replication, however, may prefer to focus on some of the chapters. Chapters 1 and 11 through 13 provide a detailed description of replication in the context of databases. Theoretical aspects of replication under benign failures are discussed in Chapters 1, 3 and 4. Chapters 5, 12 and 13 cover many issues involving practical replication issues. Chapters 8 through 10 address replication under malign failures (i.e., Byzantine failures). Readers mostly interested in historical aspects of replication should read Chaps. 6 and 7.

The Monte Verità seminar organizers are thankful to all the participants for accepting to take part in this unique seminar, and to all authors for taking their time to produce this book. We would also like to thank a number of institutions for the financial support to the seminar: the Monte Verità Foundation, the Hasler Foundation, Microsoft, Eidgenössische Technische Hochschule Zürich (ETHZ), École Polytechnique Fédérale de Lausanne (EPFL), Università della Svizzera italiana (USI), and the École polytechnique in Palaiseau.

October 2009

Bernadette Charron-Bost
Fernando Pedone
André Schiper



<http://www.springer.com/978-3-642-11293-5>

Replication

Theory and Practice

Charron-Bost, B.; Pedone, F.; Schiper, A. (Eds.)

2010, XV, 290 p., Softcover

ISBN: 978-3-642-11293-5