

# Introduction to Reliable and Secure Distributed Programming



Christian Cachin • Rachid Guerraoui  
Luís Rodrigues

Introduction to

# Reliable and Secure Distributed Programming

Second Edition

 Springer

Dr. Christian Cachin  
IBM Research Zürich  
Säumerstrasse 4  
8803 Rüschlikon  
Switzerland  
cca@zurich.ibm.com

Prof. Luís Rodrigues  
INESC-ID  
Instituto Superior Técnico  
Rua Alves Redol 9  
1000-029 Lisboa  
Portugal  
ler@ist.utl.pt

Prof. Dr. Rachid Guerraoui  
Ecole Polytechnique  
Fédérale Lausanne (EPFL)  
Fac. Informatique et Communications  
Lab. Programmation Distribuée (LPD)  
Station 14  
1015 Lausanne  
Bat. INR  
Switzerland  
Rachid.Guerraoui@epfl.ch

ISBN 978-3-642-15259-7 e-ISBN 978-3-642-15260-3  
DOI 10.1007/978-3-642-15260-3  
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2011921701

ACM Computing Classification (1998): C.2, F.2, G.2

© Springer-Verlag Berlin Heidelberg 2011, 2006

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

*Cover design:* KuenkelLopka GmbH

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

*To Irene, Philippe and André.  
To Maria and Sarah.  
To Hugo and Sara.*



# Preface

This book provides an introduction to distributed programming abstractions and presents the fundamental algorithms that implement them in several distributed environments. The reader is given insight into the important problems of distributed computing and the main algorithmic techniques used to solve these problems. Through examples the reader can learn how these methods can be applied to building distributed applications. The central theme of the book is the tolerance to uncertainty and adversarial influence in a distributed system, which may arise from network delays, faults, or even malicious attacks.

## Content

In modern computing, a program usually encompasses *multiple processes*. A process is simply an abstraction that may represent a physical computer or a virtual one, a processor within a computer, or a specific thread of execution in a concurrent system. The fundamental problem with devising such distributed programs is to have all processes *cooperate* on some *common* task. Of course, traditional centralized algorithmic issues still need to be dealt with for each process individually. Distributed environments, which may range from a single computer to a data center or even a global system available around the clock, pose additional challenges: how to achieve a robust form of cooperation despite process failures, disconnections of some of the processes, or even malicious attacks on some processes? Distributed algorithms should be dependable, offer reliability and security, and have predictable behavior even under negative influence from the environment.

If no cooperation were required, a distributed program would simply consist of a set of independent centralized programs, each running on a specific process, and little benefit could be obtained from the availability of several processes in a distributed environment. It was the need for cooperation that revealed many of the fascinating problems addressed by this book, problems that need to be solved to make distributed computing a reality. The book not only introduces the reader to these problem statements, it also presents ways to solve them in different contexts.

Not surprisingly, distributed programming can be significantly simplified if the difficulty of robust cooperation is encapsulated within specific *abstractions*. By encapsulating all the tricky algorithmic issues, such distributed programming abstractions bridge the gap between network communication layers, which are

usually frugal in terms of dependability guarantees, and distributed application layers, which usually demand highly dependable primitives.

The book presents various distributed programming abstractions and describes algorithms that implement them. In a sense, we give the distributed application programmer a library of abstract interface specifications, and give the distributed system builder a library of algorithms that implement the specifications.

A significant amount of the preparation time for this book was devoted to formulating a collection of exercises and developing their solutions. We strongly encourage the reader to work out the exercises. We believe that no reasonable understanding can be achieved in a passive way. This is especially true in the field of distributed computing, where the human mind too often follows some attractive but misleading intuition. The book also includes the solutions for all exercises, to emphasize our intention to make them an integral part of the content. Many exercises are rather easy and can be discussed within an undergraduate teaching classroom. Other exercises are more difficult and need more time. These can typically be studied individually.

## Presentation

The book as such is self-contained. This has been made possible because the field of distributed algorithms has reached a certain level of maturity, where distracting details can be abstracted away for reasoning about distributed algorithms. Such details include the behavior of the communication network, its various kinds of failures, as well as implementations of cryptographic primitives; all of them are treated in-depth by other works. Elementary knowledge about algorithms, first-order logic, programming languages, networking, security, and operating systems might be helpful. But we believe that most of our abstractions and algorithms can be understood with minimal knowledge about these notions.

The book follows an incremental approach and was primarily written as a textbook for teaching at the undergraduate or basic graduate level. It introduces the fundamental elements of distributed computing in an intuitive manner and builds sophisticated distributed programming abstractions from elementary ones in a modular way. Whenever we devise algorithms to implement a given abstraction, we consider a simple distributed-system model first, and then we revisit the algorithms in more challenging models. In other words, we first devise algorithms by making strong simplifying assumptions on the distributed environment, and then we discuss how to weaken those assumptions.

We have tried to balance intuition and presentation simplicity on the one hand with rigor on the other hand. Sometimes rigor was affected, and this might not have been always on purpose. The focus here is rather on abstraction specifications and algorithms, not on computability and complexity. Indeed, there is no theorem in this book. Correctness arguments are given with the aim of better understanding the algorithms: they are not formal correctness proofs per se.



## Organization

The book has six chapters, grouped in two parts. The first part establishes the common ground:

- In Chapter 1, we *motivate* the need for distributed programming abstractions by discussing various applications that typically make use of such abstractions. The chapter also introduces the modular notation and the pseudo code used to describe the algorithms in the book.
- In Chapter 2, we present different kinds of *assumptions* about the underlying distributed environment. We introduce a family of distributed-system models for this purpose. Basically, a model describes the low-level abstractions on which more sophisticated ones are built. These include process and communication link abstractions. This chapter might be considered as a reference to other chapters.

The remaining four chapters make up the second part of the book. Each chapter is devoted to one problem, containing a broad class of related abstractions and various algorithms implementing them. We will go from the simpler abstractions to the more sophisticated ones:

- In Chapter 3, we introduce communication abstractions for distributed programming. They permit the *broadcasting* of a message to a group of processes and offer diverse reliability guarantees for delivering messages to the processes. For instance, we discuss how to make sure that a message delivered to one process is also delivered to all other processes, despite the crash of the original sender process.
- In Chapter 4, we discuss *shared memory* abstractions, which encapsulate simple forms of distributed storage objects, accessed by read and write operations. These could be files in a distributed storage system or registers in the memory of a multi-processor computer. We cover methods for reading and writing data values by clients, such that a value stored by a set of processes can later be retrieved, even if some of the processes crash, have erased the value, or report wrong data.
- In Chapter 5, we address the *consensus* abstraction through which a set of processes can decide on a common value, based on values that the processes initially propose. They must reach the same decision despite faulty processes, which may have crashed or may even actively try to prevent the others from reaching a common decision.
- In Chapter 6, we consider *variants of consensus*, which are obtained by extending or modifying the consensus abstraction according to the needs of important applications. This includes total-order broadcast, terminating reliable broadcast, (non-blocking) atomic commitment, group membership, and view-synchronous communication.

The distributed algorithms we study not only differ according to the actual abstraction they implement, but also according to the assumptions they make on the underlying distributed environment. We call the set of initial abstractions that an algorithm takes for granted a *distributed-system model*. Many aspects have a fundamental impact on how an algorithm is designed, such as the reliability of the links,

the degree of synchrony of the system, the severity of the failures, and whether a deterministic or a randomized solution is sought.

In several places throughout the book, the same basic distributed programming primitive is implemented in multiple distributed-system models. The intention behind this is two-fold: first, to create insight into the specific problems encountered in a particular system model, and second, to illustrate how the choice of a model affects the implementation of a primitive.

A detailed study of all chapters and the associated exercises constitutes a rich and thorough introduction to the field. Focusing on each chapter solely for the specifications of the abstractions and their underlying algorithms in their simplest form, i.e., for the simplest system model with crash failures only, would constitute a shorter, more elementary course. Such a course could provide a nice companion to a more practice-oriented course on distributed programming.

## Changes Made for the Second Edition

This edition is a thoroughly revised version of the first edition. Most parts of the book have been updated. But the biggest change was to expand the scope of the book to a new dimension, addressing the key concept of *security against malicious actions*. Abstractions and algorithms in a model of distributed computing that allows adversarial attacks have become known as *Byzantine fault-tolerance*.

The first edition of the book was titled “Introduction to Reliable Distributed Programming.” By adding one word (“secure”) to the title – and adding one co-author – the evolution of the book reflects the developments in the field of distributed systems and in the real world. Since the first edition was published in 2006, it has become clear that most practical distributed systems are threatened by intrusions and that insiders cannot be ruled out as the source of malicious attacks. Building dependable distributed systems nowadays requires an interdisciplinary effort, with inputs from distributed algorithms, security, and other domains.

On the technical level, the syntax for modules and the names of some events have changed, in order to add more structure for presenting the algorithms. A module may now exist in multiple instances at the same time within an algorithm, and every instance is named by a unique identifier for this purpose. We believe that this has simplified the presentation of several important algorithms.

The first edition of this book contained a companion set of running examples implemented in the Java programming language, using the *Appia* protocol composition framework. The implementation addresses systems subject to crash failures and is available from the book’s online website.

## Online Resources

More information about the book, including the implementation of many protocols from the first edition, tutorial presentation material, classroom slides, and errata, is available online on the book’s website at:

<http://distributedprogramming.net>

## References

We have been exploring the world of distributed programming abstractions for almost two decades now. The material of this book has been influenced by many researchers in the field of distributed computing. A special mention is due to Leslie Lamport and Nancy Lynch for having posed fascinating problems in distributed computing, and to the *Cornell school* of reliable distributed computing, including Özalp Babaoglu, Ken Birman, Keith Marzullo, Robbert van Renesse, Rick Schlichting, Fred Schneider, and Sam Toueg.

Many other researchers have directly or indirectly inspired the material of this book. We did our best to reference their work throughout the text. All chapters end with notes that give context information and historical references; our intention behind them is to provide hints for further reading, to trace the history of the presented concepts, as well as to give credit to the people who invented and worked out the concepts. At the end of the book, we reference books on other aspects of distributed computing for further reading.

## Acknowledgments

We would like to express our deepest gratitude to our undergraduate and graduate students from the École Polytechnique Fédérale de Lausanne (EPFL) and the University of Lisboa (UL), for serving as reviewers of preliminary drafts of this book. Indeed, they had no choice and needed to prepare for their exams anyway! But they were indulgent toward the bugs and typos that could be found in earlier versions of the book as well as associated slides, and they provided us with useful feedback.

Partha Dutta, Corine Hari, Michal Kapalka, Petr Kouznetsov, Ron Levy, Maxime Monod, Bastian Pochon, and Jesper Spring, graduate students from the School of Computer and Communication Sciences of EPFL, Filipe Araújo and Hugo Miranda, graduate students from the Distributed Algorithms and Network Protocol (DIALNP) group at the Departamento de Informática da Faculdade de Ciências da Universidade de Lisboa (UL), Leila Khalil and Robert Basmadjian, graduate students from the Lebanese University in Beirut, as well as Ali Ghodsi, graduate student from the Swedish Institute of Computer Science (SICS) in Stockholm, suggested many improvements to the algorithms presented in the book.

Several implementations for the “hands-on” part of the book were developed by, or with the help of, Alexandre Pinto, a key member of the *Appia* team, complemented with inputs from several DIALNP team members and students, including Nuno Carvalho, Maria João Monteiro, and Luís Sardinha.

Finally, we would like to thank all our colleagues who were kind enough to comment on earlier drafts of this book. These include Felix Gaertner, Benoit Garbinato, and Maarten van Steen.

## **Acknowledgments for the Second Edition**

Work on the second edition of this book started while Christian Cachin was on sabbatical leave from IBM Research at EPFL in 2009. We are grateful for the support of EPFL and IBM Research.

We thank again the students at EPFL and the University of Lisboa, who worked with the book, for improving the first edition. We extend our gratitude to the students at the Instituto Superior Técnico (IST) of the Universidade Técnica de Lisboa, at ETH Zürich, and at EPFL, who were exposed to preliminary drafts of the additional material included in the second edition, for their helpful feedback.

We are grateful to many attentive readers of the first edition and to those who commented on earlier drafts of the second edition, for pointing out problems and suggesting improvements. In particular, we thank Zinaida Benenson, Alysso Bessani, Diego Biurrun, Filipe Cristóvão, Dan Dobre, Felix Freiling, Ali Ghodsi, Seif Haridi, Matúš Harvan, Rüdiger Kapitza, Nikola Knežević, Andreas Knobel, Mihai Letia, Thomas Locher, Hein Meling, Hugo Miranda, Luís Pina, Martin Schaub, and Marko Vukolić.

*Christian Cachin  
Rachid Guerraoui  
Luís Rodrigues*

# Contents

<b>1</b>	<b>Introduction</b> .....	1
1.1	Motivation .....	1
1.2	Distributed Programming Abstractions .....	3
1.2.1	Inherent Distribution .....	4
1.2.2	Distribution as an Artifact .....	6
1.3	The End-to-End Argument .....	7
1.4	Software Components .....	8
1.4.1	Composition Model .....	8
1.4.2	Programming Interface .....	11
1.4.3	Modules .....	13
1.5	Classes of Algorithms .....	16
1.6	Chapter Notes .....	17
<b>2</b>	<b>Basic Abstractions</b> .....	19
2.1	Distributed Computation .....	20
2.1.1	Processes and Messages .....	20
2.1.2	Automata and Steps .....	20
2.1.3	Safety and Liveness .....	22
2.2	Abstracting Processes .....	24
2.2.1	Process Failures .....	24
2.2.2	Crashes .....	24
2.2.3	Omissions .....	26
2.2.4	Crashes with Recoveries .....	26
2.2.5	Eavesdropping Faults .....	28
2.2.6	Arbitrary Faults .....	29
2.3	Cryptographic Abstractions .....	30
2.3.1	Hash Functions .....	30
2.3.2	Message-Authentication Codes (MACs) .....	30
2.3.3	Digital Signatures .....	31
2.4	Abstracting Communication .....	32
2.4.1	Link Failures .....	33
2.4.2	Fair-Loss Links .....	34
2.4.3	Stubborn Links .....	35
2.4.4	Perfect Links .....	37
2.4.5	Logged Perfect Links .....	38

- 2.4.6 Authenticated Perfect Links ..... 40
- 2.4.7 On the Link Abstractions ..... 43
- 2.5 Timing Assumptions ..... 44
  - 2.5.1 Asynchronous System ..... 44
  - 2.5.2 Synchronous System ..... 45
  - 2.5.3 Partial Synchrony ..... 47
- 2.6 Abstracting Time ..... 48
  - 2.6.1 Failure Detection ..... 48
  - 2.6.2 Perfect Failure Detection ..... 49
  - 2.6.3 Leader Election ..... 51
  - 2.6.4 Eventually Perfect Failure Detection ..... 53
  - 2.6.5 Eventual Leader Election ..... 56
  - 2.6.6 Byzantine Leader Election ..... 60
- 2.7 Distributed-System Models ..... 63
  - 2.7.1 Combining Abstractions ..... 63
  - 2.7.2 Setup ..... 64
  - 2.7.3 Quorums ..... 65
  - 2.7.4 Measuring Performance ..... 65
- 2.8 Exercises ..... 67
- 2.9 Solutions ..... 68
- 2.10 Chapter Notes ..... 71
- 3 Reliable Broadcast ..... 73**
  - 3.1 Motivation ..... 73
    - 3.1.1 Client–Server Computing ..... 73
    - 3.1.2 Multiparticipant Systems ..... 74
  - 3.2 Best-Effort Broadcast ..... 75
    - 3.2.1 Specification ..... 75
    - 3.2.2 Fail-Silent Algorithm: Basic Broadcast ..... 76
  - 3.3 Regular Reliable Broadcast ..... 77
    - 3.3.1 Specification ..... 77
    - 3.3.2 Fail-Stop Algorithm: Lazy Reliable Broadcast ..... 78
    - 3.3.3 Fail-Silent Algorithm: Eager Reliable Broadcast ..... 79
  - 3.4 Uniform Reliable Broadcast ..... 81
    - 3.4.1 Specification ..... 81
    - 3.4.2 Fail-Stop Algorithm:
      - All-Ack Uniform Reliable Broadcast ..... 82
    - 3.4.3 Fail-Silent Algorithm:
      - Majority-Ack Uniform Reliable Broadcast ..... 84
  - 3.5 Stubborn Broadcast ..... 85
    - 3.5.1 Specification ..... 85
    - 3.5.2 Fail-Recovery Algorithm: Basic Stubborn Broadcast ..... 86
  - 3.6 Logged Best-Effort Broadcast ..... 87
    - 3.6.1 Overview ..... 87
    - 3.6.2 Specification ..... 88
    - 3.6.3 Fail-Recovery Algorithm: Logged Basic Broadcast ..... 89

- 3.7 Logged Uniform Reliable Broadcast . . . . . 90
  - 3.7.1 Specification . . . . . 90
  - 3.7.2 Fail-Recovery Algorithm:
    - Logged Majority-Ack Uniform Reliable Broadcast . . . . . 90
- 3.8 Probabilistic Broadcast . . . . . 92
  - 3.8.1 The Scalability of Reliable Broadcast . . . . . 92
  - 3.8.2 Epidemic Dissemination . . . . . 93
  - 3.8.3 Specification . . . . . 94
  - 3.8.4 Randomized Algorithm: Eager Probabilistic Broadcast . . . . . 94
  - 3.8.5 Randomized Algorithm: Lazy Probabilistic Broadcast . . . . . 97
- 3.9 FIFO and Causal Broadcast . . . . . 100
  - 3.9.1 Overview . . . . . 101
  - 3.9.2 FIFO-Order Specification . . . . . 101
  - 3.9.3 Fail-Silent Algorithm: Broadcast with Sequence Number . . . 101
  - 3.9.4 Causal-Order Specification . . . . . 103
  - 3.9.5 Fail-Silent Algorithm: No-Waiting Causal Broadcast . . . . . 104
  - 3.9.6 Fail-Stop Algorithm: Garbage-Collection of Causal Past . . . 106
  - 3.9.7 Fail-Silent Algorithm: Waiting Causal Broadcast . . . . . 108
- 3.10 Byzantine Consistent Broadcast . . . . . 110
  - 3.10.1 Motivation . . . . . 110
  - 3.10.2 Specification . . . . . 111
  - 3.10.3 Fail-Arbitrary Algorithm:
    - Authenticated Echo Broadcast . . . . . 112
  - 3.10.4 Fail-Arbitrary Algorithm: Signed Echo Broadcast . . . . . 114
- 3.11 Byzantine Reliable Broadcast . . . . . 116
  - 3.11.1 Specification . . . . . 117
  - 3.11.2 Fail-Arbitrary Algorithm:
    - Authenticated Double-Echo Broadcast . . . . . 117
- 3.12 Byzantine Broadcast Channels . . . . . 120
  - 3.12.1 Specifications . . . . . 120
  - 3.12.2 Fail-Arbitrary Algorithm: Byzantine Consistent Channel . . . 122
  - 3.12.3 Fail-Arbitrary Algorithm: Byzantine Reliable Channel . . . . . 123
- 3.13 Exercises . . . . . 124
- 3.14 Solutions . . . . . 126
- 3.15 Chapter Notes . . . . . 134
- 4 Shared Memory . . . . . 137**
  - 4.1 Introduction . . . . . 138
    - 4.1.1 Shared Storage in a Distributed System . . . . . 138
    - 4.1.2 Register Overview . . . . . 138
    - 4.1.3 Completeness and Precedence . . . . . 141
  - 4.2  $(1, N)$  Regular Register . . . . . 142
    - 4.2.1 Specification . . . . . 142
    - 4.2.2 Fail-Stop Algorithm:
      - Read-One Write-All Regular Register . . . . . 144

- 4.2.3 Fail-Silent Algorithm:
  - Majority Voting Regular Register . . . . . 146
- 4.3  $(1, N)$  Atomic Register . . . . . 149
  - 4.3.1 Specification . . . . . 149
  - 4.3.2 Transformation:
    - From  $(1, N)$  Regular to  $(1, N)$  Atomic Registers . . . . . 151
  - 4.3.3 Fail-Stop Algorithm:
    - Read-Impose Write-All  $(1, N)$  Atomic Register . . . . . 156
  - 4.3.4 Fail-Silent Algorithm:
    - Read-Impose Write-Majority  $(1, N)$  Atomic Register . . . . . 157
- 4.4  $(N, N)$  Atomic Register . . . . . 159
  - 4.4.1 Multiple Writers . . . . . 159
  - 4.4.2 Specification . . . . . 160
  - 4.4.3 Transformation:
    - From  $(1, N)$  Atomic to  $(N, N)$  Atomic Registers . . . . . 161
  - 4.4.4 Fail-Stop Algorithm:
    - Read-Impose Write-Consult-All  $(N, N)$  Atomic Reg. . . . . 165
  - 4.4.5 Fail-Silent Algorithm:
    - Read-Impose Write-Consult-Majority  $(N, N)$  Atomic Reg. . . . . 167
- 4.5  $(1, N)$  Logged Regular Register . . . . . 170
  - 4.5.1 Precedence in the Fail-Recovery Model . . . . . 170
  - 4.5.2 Specification . . . . . 170
  - 4.5.3 Fail-Recovery Algorithm: Logged Majority Voting . . . . . 172
- 4.6  $(1, N)$  Byzantine Safe Register . . . . . 175
  - 4.6.1 Specification . . . . . 176
  - 4.6.2 Fail-Arbitrary Algorithm: Byzantine Masking Quorum . . . . . 177
- 4.7  $(1, N)$  Byzantine Regular Register . . . . . 179
  - 4.7.1 Specification . . . . . 179
  - 4.7.2 Fail-Arbitrary Algorithm:
    - Authenticated-Data Byzantine Quorum . . . . . 180
  - 4.7.3 Fail-Arbitrary Algorithm:
    - Double-Write Byzantine Quorum . . . . . 182
- 4.8  $(1, N)$  Byzantine Atomic Register . . . . . 188
  - 4.8.1 Specification . . . . . 189
  - 4.8.2 Fail-Arbitrary Algorithm:
    - Byzantine Quorum with Listeners . . . . . 189
- 4.9 Exercises . . . . . 194
- 4.10 Solutions . . . . . 195
- 4.11 Chapter Notes . . . . . 200
- 5 Consensus . . . . . 203**
  - 5.1 Regular Consensus . . . . . 204
    - 5.1.1 Specification . . . . . 204
    - 5.1.2 Fail-Stop Algorithm: Flooding Consensus . . . . . 205
    - 5.1.3 Fail-Stop Algorithm: Hierarchical Consensus . . . . . 208



- 5.2 Uniform Consensus . . . . . 211
  - 5.2.1 Specification . . . . . 211
  - 5.2.2 Fail-Stop Algorithm: Flooding Uniform Consensus . . . . . 212
  - 5.2.3 Fail-Stop Algorithm: Hierarchical Uniform Consensus . . . . . 213
- 5.3 Uniform Consensus in the Fail-Noisy Model . . . . . 216
  - 5.3.1 Overview . . . . . 216
  - 5.3.2 Epoch-Change . . . . . 217
  - 5.3.3 Epoch Consensus . . . . . 220
  - 5.3.4 Fail-Noisy Algorithm: Leader-Driven Consensus . . . . . 225
- 5.4 Logged Consensus . . . . . 228
  - 5.4.1 Specification . . . . . 228
  - 5.4.2 Logged Epoch-Change . . . . . 229
  - 5.4.3 Logged Epoch Consensus . . . . . 230
  - 5.4.4 Fail-Recovery Algorithm:  
 Logged Leader-Driven Consensus . . . . . 234
- 5.5 Randomized Consensus . . . . . 235
  - 5.5.1 Specification . . . . . 236
  - 5.5.2 Common Coin . . . . . 237
  - 5.5.3 Randomized Fail-Silent Algorithm:  
 Randomized Binary Consensus . . . . . 238
  - 5.5.4 Randomized Fail-Silent Algorithm:  
 Randomized Consensus with Large Domain . . . . . 242
- 5.6 Byzantine Consensus . . . . . 244
  - 5.6.1 Specifications . . . . . 244
  - 5.6.2 Byzantine Epoch-Change . . . . . 246
  - 5.6.3 Byzantine Epoch Consensus . . . . . 248
  - 5.6.4 Fail-Noisy-Arbitrary Algorithm:  
 Byzantine Leader-Driven Consensus . . . . . 259
- 5.7 Byzantine Randomized Consensus . . . . . 261
  - 5.7.1 Specification . . . . . 261
  - 5.7.2 Randomized Fail-Arbitrary Algorithm:  
 Byzantine Randomized Binary Consensus . . . . . 261
- 5.8 Exercises . . . . . 266
- 5.9 Solutions . . . . . 268
- 5.10 Chapter Notes . . . . . 277
  
- 6 Consensus Variants . . . . . 281**
  - 6.1 Total-Order Broadcast . . . . . 281
    - 6.1.1 Overview . . . . . 281
    - 6.1.2 Specifications . . . . . 283
    - 6.1.3 Fail-Silent Algorithm:  
 Consensus-Based Total-Order Broadcast . . . . . 284
  - 6.2 Byzantine Total-Order Broadcast . . . . . 287
    - 6.2.1 Overview . . . . . 287
    - 6.2.2 Specification . . . . . 288

- 6.2.3 Fail-Noisy-Arbitrary Algorithm:
  - Rotating Sender Byzantine Broadcast ..... 288
- 6.3 Terminating Reliable Broadcast ..... 292
  - 6.3.1 Overview ..... 292
  - 6.3.2 Specification ..... 293
  - 6.3.3 Fail-Stop Algorithm: Consensus-Based
    - Uniform Terminating Reliable Broadcast ..... 293
- 6.4 Fast Consensus ..... 296
  - 6.4.1 Overview ..... 296
  - 6.4.2 Specification ..... 297
  - 6.4.3 Fail-Silent Algorithm:
    - From Uniform Consensus to Uniform Fast Consensus ..... 297
- 6.5 Fast Byzantine Consensus ..... 300
  - 6.5.1 Overview ..... 300
  - 6.5.2 Specification ..... 300
  - 6.5.3 Fail-Arbitrary Algorithm:
    - From Byzantine Consensus to Fast Byzantine Consensus ... 300
- 6.6 Nonblocking Atomic Commit ..... 303
  - 6.6.1 Overview ..... 303
  - 6.6.2 Specification ..... 304
  - 6.6.3 Fail-Stop Algorithm:
    - Consensus-Based Nonblocking Atomic Commit ..... 304
- 6.7 Group Membership ..... 307
  - 6.7.1 Overview ..... 307
  - 6.7.2 Specification ..... 308
  - 6.7.3 Fail-Stop Algorithm: Consensus-Based
    - Group Membership ..... 309
- 6.8 View-Synchronous Communication ..... 311
  - 6.8.1 Overview ..... 311
  - 6.8.2 Specification ..... 312
  - 6.8.3 Fail-Stop Algorithm:
    - TRB-Based View-Synchronous Communication ..... 314
  - 6.8.4 Fail-Stop Algorithm: Consensus-Based
    - Uniform View-Synchronous Communication ..... 319
- 6.9 Exercises ..... 323
- 6.10 Solutions ..... 324
- 6.11 Chapter Notes ..... 337
- 7 Concluding Remarks ..... 341**
  - 7.1 Implementation in *Appia* ..... 341
  - 7.2 Further Implementations ..... 342
  - 7.3 Further Reading ..... 344

Contents xix

**References** ..... 347

**List of Modules** ..... 355

**List of Algorithms** ..... 357

**Index** ..... 361