

Preface

Overview and goals

This book has grown out of the authors' lecture notes and studies in the field of distributed processing. In our teaching of distributed processing, we have found it necessary to deal with the traditional views of concurrency (e.g., matters such as race conditions, semaphores and mutual exclusion), and also the practical realisation of networked systems and to consider the underlying operating systems.

There are many books that focus on the theory of concurrency, operating systems or programming, but few that take a coherent view of distributed processing as a practical subject. This book thus aims at presenting the reader with a clear overview of the *process* that is followed when building distributed systems, and describes the tools, techniques and theory that can be applied within this process. The book aims to consider the important points of the engineering process, including the *models* that can be used to assess and analyse parts of distributed systems, the implementation techniques —such as sockets— for these models, and the protocols and security concerns that must be considered as part of any realistic distributed system.

Organisation and features

This book can reasonably be divided into three parts: foundations (of both theory and practical implementation concerns), engineering issues (including security and language concerns), and examples and case studies, serving to

integrate material from the previous two parts.

The first part of the book starts with an overview of distributed processing in Chapter 1. The critical concepts of distributed processing and concurrency —like semaphores, deadlock and naming— are discussed in Chapter 2. In Chapter 3 we present several rigorous models of concurrency that help us understand the underlying theory and practice behind the critical concepts of Chapter 2. We also discuss how these models are used in building distributed systems. Chapter 4 discusses operating systems, and their role in building distributed systems, followed by Chapter 5 which deals with interprocess communication. In Chapter 6 we study protocols, which are an important part of building realistic distributed systems. Examples of protocols are presented, as well as the challenges associated with designing them.

The second part of the book moves to engineering issues. Chapter 7 considers general concerns of security for distributed systems. We take an *engineering* perspective of security; it is important to be aware that security is not simply a matter of encrypting communication: a whole-system view must be taken. Chapter 8 presents a selection of languages and focuses on how these languages can be used to build distributed systems. This serves to illustrate many of the ideas that appear in the previous chapters, both theoretical and practical.

The third part of the book focuses on examples and case studies, aiming to integrate all the previous chapters. In Chapter 9 we go through the engineering life-cycle to carry out a selection of case studies in building distributed systems. While we do not implement these systems, we aim to illustrate the *process* of their construction, focusing on identifying requirements, thinking about protocols, and design. Chapter 10 presents a worked example: building a networked game. We identify technical and business requirements, present a design, discuss protocols and security concerns, and construct an implementation. This chapter serves to bring all of the previous chapters together in a coherent structure. Finally, Chapter 11 concludes the book, and identifies some areas for future work, reading, and experiment.

Each chapter concludes with a coherent summary, and highlights a number of exercises. Some exercises focus on thinking about the issues and problems identified in the chapter; others concentrate on building small parts of distributed systems.

A comprehensive glossary is included. Example code is indicated by ‘EG’ in the margin; similarly, related content is marked ‘↗’.

Target audience

This book is aimed at students who have a reasonable amount of (sequential) programming experience, but who may not have taken a course on operating systems or concurrency. Some experience with C programming may be helpful, but is certainly not a prerequisite to reading and understanding the material in this book. Readers who do not yet have broad programming experience will need to consult additional texts (see the bibliography) on specific programming languages.

We have taught the material in this text to undergraduate students at the second and third year level. With some additions —particularly, a large-scale project— the book could prove suitable to final-year students taking a software engineering specialisation as well.

Notes to the instructor

This book can be used for a coherent, one-term course or module on distributed systems. It can also be used as part of a longer course on software engineering where distributed systems elements play a substantial role. For use in a stand-alone course, we recommend following the book roughly sequentially, though within chapters topics can be rearranged to a degree without losing the overall flow. Instructors may choose to selectively cover some of the formal models in Chapter 3, depending on student background and interests, although we strongly recommend the inclusion of statecharts, as these are important and are used in the worked example in Chapter 10. Students who have previously taken a course in operating systems may be able to skip the material on semaphores in Chapter 2, though it may be helpful to use as a refresher.

We have found the case studies in Chapter 9 useful as reading and presentation projects, to be presented by small groups of students in class.

The material in Chapter 10 could form the basis of larger projects in distributed systems. Indeed, we have run both research and development projects for undergraduate and taught Masters students based on the material covered in this chapter. Some suggestions for areas to investigate for larger and more long-term projects are given in Section 11.2.2.

Sketch solutions or hints to many of the exercises are included in Appendix A. Sample code, including the networked game considered in Chapter 10, is available from the book's web site, <http://www.scm.tees.ac.uk/p.j.brooke/dpb/>. Additional extensions to the multiplayer game presented

in Chapter 10 are likely to be added over time. Comments and suggestions are welcome, and can be directed to the authors.

Acknowledgments

We thank our colleagues at the University of Teesside and the University of York (and, previously, the University of Plymouth and York University, Canada) for their suggestions, advice and helpful guidance. We also thank the reviewers of this book for their useful and timely suggestions. Special thanks should go to Catherine Brett and Wayne Wheeler of Springer for their assistance during the editing and publication process, as well as to Ian Mackie and Helen Callaghan who started the whole process.

Several software tools were used in the development of this book; without these tools, the writing process would have been even more fraught than it was. We would particularly like to thank the authors and maintainers of \LaTeX and Subversion.

Phil would like to thank Christine, Rebecca and Alexander for their support during evenings and days spent coding and writing.

Rich would like to thank Angelika for putting up with him (and Phil) during the writing, rewriting and gaming processes.

Both Phil and Rich would like to thank their proof-readers for catching errors and improving their writing. All remaining errors are, of course, the responsibility of the authors.

Phil Brooke
Rich Paige

July 2007



<http://www.springer.com/978-1-84628-840-1>

Practical Distributed Processing

Brooke, P.J.; Paige, R.F.

2008, XIV, 262 p. 24 illus., Softcover

ISBN: 978-1-84628-840-1