

Texts in Computer Science

Editors

David Gries

Fred B. Schneider

For other titles published in this series, go to
www.springer.com/series/3191

A.W. Roscoe

Understanding Concurrent Systems

 Springer

Prof. A.W. Roscoe
Oxford University Computing Laboratory
Wolfson Bldg., Parks Road
Oxford OX1 3QD
UK
Bill.Roscoe@comlab.ox.ac.uk

Series Editors

David Gries
Department of Computer Science
Upson Hall
Cornell University
Ithaca, NY 14853-7501, USA

Fred B. Schneider
Department of Computer Science
Upson Hall
Cornell University
Ithaca, NY 14853-7501, USA

ISSN 1868-0941
ISBN 978-1-84882-257-3
DOI 10.1007/978-1-84882-258-0
Springer London Dordrecht Heidelberg New York

e-ISSN 1868-095X
e-ISBN 978-1-84882-258-0

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Control Number: 2010936468

© Springer-Verlag London Limited 2010

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act 1988, this publication may only be reproduced, stored or transmitted, in any form or by any means, with the prior permission in writing of the publishers, or in the case of reprographic reproduction in accordance with the terms of licenses issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publishers.

The use of 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 laws and regulations and therefore free for general use.

The publisher makes no representation, express or implied, with regard to the accuracy of the information contained in this book and cannot accept any legal responsibility or liability for any errors or omissions that may be made.

Cover design: VTEX, Vilnius

Printed on acid-free paper

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

Preface

Since C.A.R. Hoare’s text *Communicating Sequential Processes* [60] was published in 1985, his CSP notation has been extensively used for teaching and applying concurrency theory. I published a book [123] on this topic in 1997 entitled “The Theory and Practice of Concurrency”. We will be referring to the latter many times in the present book, and will abbreviate it *TPC*. Hoare’s book and TPC are now freely available via the web. The present book draws heavily on material from TPC: parts of the text are updated versions of sections of TPC, mainly in Part I, and we will frequently refer to material in it. However, we omit most advanced material from TPC that does not require updating. Chapter 4 and each of Chaps. 8–20 are either mainly or entirely new.

When I started writing this book, I thought of it as a much revised second edition of TPC. That remains true for the first (introductory) part, but for the rest the completed book is perhaps better described as a “sequel”.

Overview and Goals

The aim of this book is to be a comprehensive introduction and reference volume for CSP, providing the material for a number of different courses. It should also be the first point of reference for anyone wanting to use CSP or find out about its theory. It introduces other views of concurrency, using CSP to model and explain these. This book is integrated with, and uses, CSP-based tools and especially FDR to a much greater extent than TPC, and in addition describes how to create new tools based on FDR.

FDR is freely downloadable for all but commercial use. Almost all chapters either explain or illustrate the use of FDR, and there are two chapters (8 and 16) specifically about it.

This book is divided into parts on similar lines to TPC: the first three are, respectively, an introductory course on CSP, a review of the theory of CSP, and some topics on the application of CSP and its tools. Whereas TPC has extensive appendices on

the mathematical background, the CSP_M language and FDR, the present book has a fourth proper part.

I have become convinced that CSP and FDR are ideal vehicles for understanding and reasoning about a wide variety of concurrent systems, not just those initially described in this language. I and others have written a number of compilers that take programs written in other notations and translate them to CSP for analysis on FDR—and most also take the output of FDR and translate it back to give feedback appropriate to the non-CSP input notation. The fourth part is devoted to this topic, while simultaneously providing an introduction to the other sorts of concurrent systems that we explain in CSP.

Organisation and Features

Part I: An Introduction to CSP

This part is designed to be used for an introductory course on CSP.

It owes much to TPC, though some of the more difficult topics in Part I of TPC have been delayed to later parts, or removed. In their place we have included some hopefully more interesting case studies.

Part II: Theory

The theory of CSP has made considerable advances since 1997 in a number of directions, notably the following:

- Models have been developed [126] that do not need the sometimes controversial *divergence strictness* previously needed in all models that handle infinite behaviours.
- Our understanding of the overall hierarchy of CSP's semantic models [129, 130] has advanced: new models have been discovered, and we also know that there is no alternative to the weakest few models.
- We now have compositional models [99] for *Discrete Timed CSP*, lying between the informal *tock*-CSP language from Chap. 14 of TPC and the (continuous) *Timed CSP* of [111, 146]. Importantly, advances [98, 99] have been made in linking all these together.
- Links have been made between operational semantics in a general setting and CSP models [132]: it has thus been *proved* that one can give a semantics *in CSP* to a wide variety of languages.
- We now understand [89, 128] what properties can be specified using the (FDR-like) refinement check $F(P) \sqsubseteq G(P)$ for CSP contexts F and G .
- Impressive theories, for example [116], have been built on mechanical theorem provers to check the properties of CSP models and semantics.

- It is necessary to change much of the structure of the algebraic semantics for CSP presented in TPC if they are to encompass some of the new models now available.
- The theory of model checking CSP (which is what FDR does) has developed substantially in a number of ways.

In this book we will discuss most of these topics, some in Part II and some in Part III.

It would be impossible to give an in-depth presentation of the theory of CSP as it now is, as TPC largely did (with the exception of Timed CSP) in 1997. In the present book we explain some basic concepts in detail and introduce the main ideas behind the advanced theory, referring the reader to the resources provided by TPC and many academic papers. Thus Part II provides the core of a textbook, building on itself by referencing these on-line resources. There are chapters on:

- Operational semantics
- The usual denotational models
- The hierarchy of models where everything is finitely observable
- The hierarchy of models that include infinite observations
- Algebraic semantics

Part III: Using CSP

This part covers a number of topics in the application of CSP, and has chapters on:

- Timed modelling and analysis using the “*tock-time*” model introduced in TCP.
- The discrete modelling and verification of Timed CSP.
- More about FDR: advanced topics in the use of FDR and advanced specification techniques.
- Parameterised verifications and the state explosion problem: we introduce some techniques for coping with the exponential growth in the number of states to explore as we look at instances of networks with more processes, and techniques for proving properties of large classes of network. These include data independence, induction, and buffer tolerance.

Part IV: Exploring Concurrency

Here we emphasise the ability of CSP to describe and enable reasoning about parallel systems modelled in other paradigms. There are chapters on:

- Compiling shared variable programs into CSP
- Shared variable concurrency
- Priority and mobile processes

The two chapters on shared variable concurrency introduce a new tool called SVA, created by myself and David Hopkins, as a front end for FDR. We will see that it is highly effective in revealing how shared variable programs behave.

As in TPC I have attempted, in writing Parts III and IV, to rely as little as possible on the theory presented in Part II. Therefore either of these, or various combinations of chapters and sections chosen from them, can be used as the basis of a course on concurrency to follow up Part I.

Target Audience

This book is aimed at everyone who wants to get an in-depth understanding of concurrent systems, and will be essential reading for anyone interested in Hoare's CSP.

Part I is designed for an audience of undergraduate and Masters'-level graduate computer science students. At Oxford it is used for a second-year undergraduate course, and for both full-time and part-time M.Sc. students.

Part II is designed for people who are familiar with Part I and have fairly theoretical interests. These could be students taking an advanced course based on this material, or researchers interested in the state of the art.

Part III is intended for people who already have some experience in using CSP and FDR in practice, and want to be able to use them better or who are specifically interested in timed systems.

Part IV is designed for people who already understand CSP. They might want to understand other models of concurrent systems in terms of CSP. They might want model shared-variable, mobile or prioritised systems in CSP. Or they might want to write a translator from another language into CSP.

Most of the present book relies on no theoretical background other than a basic knowledge of sets and sequences. Some of Part II relies on a knowledge of basic partial order and metric space theory such as can be obtained by studying Appendix A of TPC. Except in Part II, I have tried to avoid making the reader follow sophisticated mathematical arguments, though this proved unavoidable in parts of Chap. 17.

Whilst I was writing this book, many people asked me to provide many examples of how to program in CSP: *design patterns*. While this is not the book of case studies that some wanted, I have tried to include enough to keep them happy. The main case studies can be found in Chaps. 4, 8, 9, 14, 15, 17, 18, 19 and 20.

Notes to the Instructor

Chapters 1–6 (with Chap. 7 being optional) provide a comprehensive introductory course on CSP, dipping into Chap. 8 on FDR as required during the course. When deciding whether or not to include Chap. 7, the instructor should bear in mind that Chaps. 18 and 19 (on shared variables) depend heavily on sequential composition.

For an audience already familiar with CSP one could give a theory course based on Part II. Many different courses on the practical uses of CSP and FDR could be based on Chaps. 4, 8 and 14–19 of the present book and Chaps. 12–15 of TPC, and indeed there is probably enough material in Chaps. 18 and 19 on which to base a course on shared variable concurrency.

Teaching Resources

This book has a web-site www.comlab.ox.ac.uk/ucs where you can find links to complete texts of Hoare's book [60] and TPC and links from which FDR, the ProBE CSP animator, SVA and other CSP-based tools can be down-loaded.

You can also find machine-readable CSP_M versions of almost all the CSP programs in this book and in TPC, as well as overheads covering most of the material in the two books. Additionally there are practical exercises in the use of FDR that those learning this material can use, whether personally or in a course.

Further practicals and solutions to all the exercises in this book can be obtained from the author by academics using this book for teaching.

And Finally...

When I started to write this book I assumed that it would include a chapter or two on security. In fact there is very little here on this subject, and that is mainly a summary of material in TCP. This is not because there is no new material on the mixture of CSP and security, but rather the reverse. There is now a book on security protocols via CSP [142] and a lot of additional material that goes well beyond that. It has become too large a subject to be included in a general book about CSP, at least if we want to discuss the state of the art.

While this book is focused on CSP, it covers a very wide variety of concurrent systems including combinatorial, timed, priority-based, mobile, shared variable, statecharts, buffered and asynchronous systems. Furthermore, we see how to translate several other notations into CSP. I hope, therefore, that it justifies its title *Understanding Concurrent Systems*.

Oxford, UK
June 2010

Bill Roscoe

Acknowledgements

I had the good fortune to become Tony Hoare’s research student in 1978, which gave me the opportunity to work with him on the development of the ‘process algebra’ version of CSP and its semantics from the first.¹ I have constantly been impressed that the decisions he took in structuring the language have stood so well the twin tests of time and practical use in circumstances he could not have foreseen. The work in the present book all results, either directly or indirectly, from his vision. Those familiar with his book will recognise that much of my presentation, and many of my examples, have been influenced by it.

The core theory of CSP was developed in the late 1970s and 1980s. The two people most responsible, together with Tony and myself, for the development of the basic theoretical framework for CSP were Steve Brookes and Ernst-Rüdiger Olderog, and I am delighted to acknowledge their contributions. We were, naturally, much influenced by the work of those such as Robin Milner, Matthew Hennessy and Rocco de Nicola who were working at the same time on other process algebras.

The last few months have seen the deaths of Robin Milner and Amir Pnueli, both giants of computer science and concurrency theory in particular. A small aspect of Amir’s influence can be seen in Sect. 16.4 of this book. Both were wonderful and generous men who influenced generations of researchers. They will be sadly missed.

Over the years, both CSP and my understanding of it have benefited from the work of too many people for me to list their individual contributions. I would like to thank the following present and former students, colleagues, collaborators and correspondents for their help and inspiration: Samson Abramsky, *Phil Armstrong*, Geoff Barrett, Stephen Blamey, Tilo Buschmann, *Sadie Creese*, Naiem Dathi, Jim Davies, *John Fitzgerald*, Richard Forster, Paul Gardiner, *Michael Goldsmith*, Anthony Hall, He Jifeng, *Philippa Hopcroft (née Broadfoot)*, *David Hopkins*, *Huang Jian*, Jason Hulance, David Jackson, Lalita Jategaonkar Jagadeesan, Alan Jeffrey, Mark Josephs, Maneesh Khattri, Jonathan Lawrence, *Ranko Lazić*, *Eldar Kleiner*,

¹An account of the development of CSP, and in particular the transition to the process algebra version of [60] from the “imperative” version of [58], can be found in [71].

Gavin Lowe, Helen McCarthy, Christie Marr (née Bolton), Jeremy Martin, Albert Meyer, Michael Mislove, *Nick Moffat*, Lee Momtahan, *Tom Newcomb*, *Long Nguyen*, David Nowak, *Joël Ouaknine*, *Hristina Palikareva*, Ata Parashkevov, David Park, *Sriram Rajamani*, *Joy Reed*, Mike Reed, *Jakob Rehof*, *Markus Roggenbach*, Bill Rounds, Peter Ryan, Jeff Sanders, Bryan Scattergood, Steve Schneider, Brian Scott, Karen Seidel, *Jane Sinclair*, *Antti Valmari*, *Rob van Glabbeek*, *David Walker*, *Wang Xu*, *Peter Welch*, Paul Whittaker, Jim Woodcock, James (Ben) Worrell, Zhenzhong Wu, Lars Wulf, Jay Yantchev, Irfan Zakiuddin and Zhou Chao Chen.

Many of them will recognise specific influences their work has had on my two books. Those *italicised* have had a direct influence on the new work reported in the present book.

Special thanks are due to the present and former staff of Formal Systems (some of whom are listed above) for their work in developing FDR and ProBE. The remarkable capabilities of FDR transformed my view of CSP. One of my main motivations for writing this book is to show how to use FDR effectively. Bryan Scattergood was chiefly responsible for both the design and the implementation of the ASCII version of CSP used on these and other tools. The passage of time since FDR 2 was released has only emphasised the amazing job he did in designing CSP_M, and the huge expressive power of the embedded functional language.

In the last few years the development and maintenance of FDR have become the responsibility of my team at Oxford University Computing Laboratory, led by Phil Armstrong. I am pleased to say that this has led to many exciting developments with the tool and should soon lead to a third major release: FDR 3.

The presentation of this book has been greatly assisted by all those who have commented on and pointed out errors in drafts. In particular I would like to thank Irfan Zakiuddin, who read through the entire final draft and gave many useful comments. Wayne Wheeler and Simon Rees from Springer have given me just the right mixture of help, encouragement and cajoling.

My work on CSP, including the writing of both this book and TPC, has benefited from funding from several bodies over the years, including EPSRC, DRA, QinetiQ, ESPRIT, industry and the US Office of Naval Research. I am particularly grateful to Ralph Wachter from the last of these, without whom most of the research on CSP tools would not have happened.

This book could never have been written without the support of my wife Coby. For the third time (the previous ones being my doctoral thesis and TPC) she has read through hundreds of pages of text on a topic entirely foreign to her, expertly pointing out errors in spelling and style. And, yet again, she put up with me writing it.

Contents

Part I A Foundation Course in CSP

1	Building a Simple Sequential Process	3
1.1	Basic Operators	3
1.1.1	Prefixing	3
1.1.2	Recursion	4
1.1.3	Guarded Alternative	5
1.2	Choice Operators	10
1.2.1	External Choice	10
1.2.2	Nondeterministic Choice	11
1.2.3	Conditional Choice	12
1.2.4	Multi-Part Events: Extending the Notation of Channels	14
1.3	A Few Important Processes	16
1.4	Traces Refinement	16
1.5	What is a Communication?	18
1.6	Tools	19
1.6.1	Finite-State Machines	21
2	Understanding CSP	23
2.1	Algebra	23
2.2	The Traces Model and Traces Refinement	29
2.2.1	Working out $traces(P)$	30
2.2.2	Traces and Laws	33
2.2.3	Unique Fixed Points	34
2.2.4	Specification and Refinement	36
2.2.5	Afters and Initials	40
2.3	Operational Semantics and Labelled Transition Systems	41
2.4	Tools	43
3	Parallel Operators	45
3.1	Synchronous Parallel	45
3.1.1	Turning Parallel Processes into Sequential Ones	46

3.2	Alphabetised Parallel	49
3.3	Interleaving	57
3.4	Generalised Parallel	59
3.5	Parallel Composition as Conjunction	61
3.6	Tools	64
3.7	Postscript: On Alphabets	65
4	CSP Case Studies	67
4.1	Sudoku in CSP	67
4.2	Deadlock-Free Routing	73
4.2.1	Head for the Trees!	75
4.2.2	Uncloggable Rings	80
4.2.3	The Mad Postman	84
4.3	Communications Protocols	86
5	Hiding and Renaming	93
5.1	Hiding	93
5.1.1	The Consequences of Hiding	97
5.1.2	Hiding <i>versus</i> Constructiveness	100
5.2	Renaming and Alphabet Transformations	102
5.2.1	Injective Functions	103
5.2.2	Non-injective Functions	104
5.2.3	Relational Renaming	105
5.3	Linking Operators	109
5.4	Tools	112
6	Beyond Traces	115
6.1	A Brief Introduction to Failures and Divergences	115
6.2	Failures and Divergences in Specifications	120
6.3	Ungranted Requests and the Limits of Failures	123
6.4	Avoiding Divergence	124
6.5	Abstraction by Hiding	125
6.6	Tools	129
7	Further Operators	131
7.1	Termination and Sequential Composition	131
7.1.1	What is Termination?	131
7.1.2	Distributed Termination	136
7.1.3	Effects on the Failures-Divergences Model	137
7.2	Interrupting Processes	138
7.3	Tools	141
8	Using FDR	143
8.1	What is FDR?	143
8.1.1	Running and Debugging Checks	146
8.1.2	FDR's Settings	148

- 8.1.3 Defining Non-Process Objects 150
- 8.1.4 The Limits of FDR 155
- 8.2 Checking Parallel Processes 157
- 8.3 The Structure of a Refinement Check 161
- 8.4 Failures and Divergences 163
- 8.5 Watchdogs 166
- 8.6 Breadth versus Depth, Symmetry and Divergence 166
- 8.7 Determinism Checking 170
- 8.8 Compression 173
 - 8.8.1 Using Compression 175
- 8.9 Notes and Reflections 184

Part II Theory

- 9 Operational Semantics 191**
 - 9.1 Transition Systems and State Machines 192
 - 9.2 Firing Rules for CSP 200
 - 9.2.1 SOS Style Operational Rules 201
 - 9.2.2 Combinator Style Operational Rules 204
 - 9.3 From Combinators to Supercombinators! 212
 - 9.4 Translating Combinators to CSP 214
 - 9.5 Relationships with Abstract Models 221
 - 9.5.1 Extracting Failures and Divergences 221
 - 9.5.2 Infinite Traces and Infinite Branching 222
 - 9.6 Tools 227
 - 9.7 Notes 227
- 10 Denotational Semantics and Behavioural Models 229**
 - 10.1 Introduction 229
 - 10.1.1 Fixed-Point Theory 231
 - 10.2 Analysing Traces Semantics 232
 - 10.3 The Stable Failures Model 236
 - 10.3.1 Applications 239
 - 10.3.2 Channel-Based Failures 241
 - 10.4 The Failures-Divergences Model 241
 - 10.5 Determinism, Confluence and Proof 247
 - 10.6 Full Abstraction and Congruence 251
 - 10.7 Notes 252
- 11 Finite Observation Models 255**
 - 11.1 What is a Behavioural Model? 255
 - 11.1.1 The Finest Model of Them All 256
 - 11.1.2 Clouding the Glass 258
 - 11.2 A Tour through the Hierarchy 259
 - 11.2.1 The Ready Sets, or Acceptances Model \mathcal{A} 260
 - 11.2.2 The Stable Refusal Testing Model \mathcal{RT} 262

11.2.3	The Stable Revivals Model	264
11.2.4	Other Models	266
11.3	The Big Picture	268
11.4	Tools	270
12	Infinite-Behaviour Models	271
12.1	Divergence-Strict Models for Finite Nondeterminism	272
12.1.1	Determinism amongst the Richer Models	274
12.1.2	Applications	275
12.2	Strict Divergence for General CSP	276
12.2.1	Healthiness Conditions	276
12.2.2	Fixed Point Theories for Unbounded Nondeterminism	278
12.2.3	Applying Infinite-Behaviour Models	280
12.3	The Hierarchy of Divergence-Strict Models	282
12.4	Seeing Beyond Divergence	283
12.4.1	Applications of \mathcal{M}^\sharp	285
12.5	The Fixed Point Theory of \mathcal{M}^\sharp	286
12.6	The Hierarchy	289
12.7	Tools	290
12.8	Notes	290
13	The Algebra of CSP	293
13.1	Introduction	293
13.2	AOS Form	295
13.3	Algebraic Operational Semantics	299
13.4	Normal Forms	305
13.5	A Tour through the Hierarchy	309
13.6	General Programs	315
13.7	Notes	317
Part III Using CSP in Practice		
14	Timed Systems 1: <i>tock</i>-CSP	321
14.1	Introduction	321
14.2	A Brief History of Time in CSP	321
14.3	<i>tock</i> -CSP	323
14.3.1	Expressing Timing Constraints	324
14.4	Case Study: Bully Algorithm	327
14.4.1	Part 1: Specification	329
14.4.2	Part 2: Implementation	330
14.4.3	Analysis	335
14.4.4	Conclusions	337
14.5	Maximal Progress and Priority	339
14.5.1	Case Study: Timed Routing	340
14.6	Specifying and Metering <i>tock</i> -CSP	341
14.7	Tools	343

- 15 Timed Systems 2: Discrete Timed CSP** 345
 - 15.1 Modelling Timed CSP 345
 - 15.1.1 Semantics 349
 - 15.2 Examples 351
 - 15.3 Digitisation 353
 - 15.4 Notes for Chapters 14 and 15 355

- 16 More About FDR** 357
 - 16.1 Normalisation 357
 - 16.2 More About Compression 363
 - 16.2.1 Strong Bisimulation 364
 - 16.2.2 DRW-Bisimulation 364
 - 16.2.3 τ -Loop Elimination 365
 - 16.2.4 Diamond Elimination 367
 - 16.2.5 Laziness and chase 370
 - 16.3 Handling Large Checks 373
 - 16.3.1 Memory Locality 373
 - 16.3.2 Parallel Implementation 376
 - 16.4 Generalising the Specification Model 379

- 17 State Explosion and Parameterised Verification** 385
 - 17.1 Induction 386
 - 17.1.1 The Limitations of Induction 390
 - 17.2 Data Independence 390
 - 17.2.1 Thresholds 392
 - 17.2.2 Beyond Thresholds 397
 - 17.3 Data-Independent Induction 398
 - 17.4 Buffer Tolerance 407
 - 17.4.1 Definitions, Basics and Tree Networks 408
 - 17.4.2 Functional and Confluent Processes 411
 - 17.5 Approximation Based Methods 413
 - 17.6 Notes 415

Part IV Exploring Concurrency

- 18 Shared-Variable Programs** 419
 - 18.1 Writing a Compiler in CSP_M 421
 - 18.1.1 Data Types 422
 - 18.1.2 Variable Names 423
 - 18.1.3 Compilation Strategy 424
 - 18.1.4 Compiling a Thread 425
 - 18.1.5 Evaluating an Expression 427
 - 18.2 Applying Compression 430
 - 18.3 The Front End of SVA 434
 - 18.3.1 SVL 434
 - 18.3.2 The GUI 436

- 18.3.3 Possible Extensions 436
- 18.4 Specifications in SVA 438
- 18.5 Case Study: Lamport’s Bakery Algorithm 439
- 18.6 Case Study: The Dining Philosophers in SVL 441
- 18.7 Notes 446
- 19 Understanding Shared-Variable Concurrency 447**
 - 19.1 Dirty Variables 447
 - 19.2 Case Study: Simpson’s 4-Slot Algorithm 451
 - 19.3 Observing and Refining SVL Programs 457
 - 19.3.1 The Bakery Algorithm Revisited 464
 - 19.4 Atomic Equivalent Programs 465
 - 19.5 Overseers: Modelling Complex Data-Types 469
 - 19.6 Abstraction in SVA 471
 - 19.6.1 Two Threads out of Many 471
 - 19.6.2 Finitely Representing an Infinite Linear Order 474
 - 19.7 Notes 480
- 20 Priority and Mobility 481**
 - 20.1 Case Study: Knight’s Tour 481
 - 20.2 Priority 486
 - 20.2.1 Statecharts 490
 - 20.2.2 Synchronous Two-Phase Automata 494
 - 20.3 Mobility 497
 - 20.3.1 An Introduction to Mobility 497
 - 20.3.2 Towards a “Mobile CSP” 498
 - 20.3.3 Pass the Port! 499
 - 20.3.4 Closed Worlds 500
 - 20.3.5 Opening Out 505
 - 20.4 Notes 506
- Notation 509**
- References 513**
- Index 519**