

Texts in Computer Science

Editors

David Gries

Fred B. Schneider

For further volumes:

www.springer.com/series/3191

Krzysztof R. Apt
Frank S. de Boer
Ernst-Rüdiger Olderog

Verification of Sequential and Concurrent Programs

Third, Extended Edition

 Springer

Krzysztof R. Apt
Centrum Wiskunde
& Informatica
Science Park 123
1098 XG Amsterdam
Netherlands
k.r.apt@cwi.nl

Frank S. de Boer
Centrum Wiskunde
& Informatica
Science Park 123
1098 XG Amsterdam
Netherlands
F.S.de.Boer@cwi.nl

Ernst-Rüdiger Olderog
Department für Informatik
Universität Oldenburg
26111 Oldenburg
Germany
olderog@informatik.uni-oldenburg.de

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 e-ISSN 1868-095X
ISBN 978-1-84882-744-8 e-ISBN 978-1-84882-745-5
DOI 10.1007/978-1-84882-745-5
Springer Dordrecht Heidelberg London New York

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

Library of Congress Control Number: 2009932111

© Springer-Verlag London Limited 2009, Reprinted with corrections 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.

Printed on acid-free paper

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

Endorsements

THE THIRD EDITION is an excellent new version of a valuable book. Enhanced with new material on recursion and object-oriented programs, this book now covers methods for verifying sequential, object-oriented, and concurrent programs using well-chosen sample programming languages that highlight fundamental issues and avoid incidental complications. With growing challenges today to produce correct software systems for the future, this book lets students wisely use a few months now to master concepts that will last them a lifetime.

John C. Mitchell, Stanford University

Verification of programs is the Holy Grail of Computer Science. This book makes its pursuit seem both pleasant and worthwhile. Its unique strength lies in the way the authors have deconstructed the apparently complex subject such that each piece carries exactly one idea. The beauty of the presentation extends from the overall structure of the book to the individual explanations, definitions and proofs.

Andreas Podelski, University of Freiburg

Program verification became an interesting research topic of computing science about forty years ago. Research literature on this topic has grown quickly in accordance with rapid development of various programming paradigms. Therefore it has been a challenge to university lecturers on program verification how to carefully select an easy but comprehensive approach, which can fit in with most programming paradigms and can be taught in a systematic way. The publication of this book is an answer to the challenge, and to my knowledge quite many university lecturers have been influenced by the earlier editions of this book if not chosen them as

textbook. Given that the third edition includes verification of object-oriented programs – the most fashionable programming paradigm, and presents it in a way coherent with the approach adopted by the earlier ones, we can expect a further impact of the new edition on university teachings.

Zhou Chaochen, Chinese Academy of Sciences, Beijing

Foreword

T HIS BOOK CONTAINS a most comprehensive text that presents syntax-directed and compositional methods for the formal verification of programs. The approach is not language-bounded in the sense that it covers a large variety of programming models and features that appear in most modern programming languages. It covers the classes of sequential and parallel, deterministic and non-deterministic, distributed and object-oriented programs. For each of the classes it presents the various criteria of correctness that are relevant for these classes, such as interference freedom, deadlock freedom, and appropriate notions of liveness for parallel programs. Also, special proof rules appropriate for each class of programs are presented. In spite of this diversity due to the rich program classes considered, there exist a uniform underlying theory of verification which is syntax-oriented and promotes compositional approaches to verification, leading to scalability of the methods.

The text strikes the proper balance between mathematical rigor and didactic introduction of increasingly complex rules in an incremental manner, adequately supported by state-of-the-art examples. As a result it can serve as a textbook for a variety of courses on different levels and varying durations. It can also serve as a reference book for researchers in the theory of verification, in particular since it contains much material that never before appeared in book form. This is specially true for the treatment of object-oriented programs which is entirely novel and is strikingly elegant. I strongly recommend this book to both teachers who wish to train students in the most advanced techniques of verification, and to researchers in this important area.

Amir Pnueli

New York University and the Weizmann Institute of Science, Rehovot

Preface

COMPUTER PROGRAMS ARE by now indispensable parts of systems that we use or rely on in our daily lives. Numerous examples include booking terminals in travel agencies, automatic teller machines, ever more sophisticated services based on telecommunication, signaling systems for cars and trains, luggage handling systems at airports or automatic pilots in airplanes.

For the customers of travel agencies and banks and for the passengers of trains and airplanes the proper functioning and safety of these systems is of paramount importance. Money orders should reflect the right bank accounts and airplanes should stay on the desired route. Therefore the underlying computer programs should work correctly; that is they should satisfy their requirements. A challenge for computer science is to develop methods that ensure program correctness.

Common to the applications mentioned above is that the computer programs have to coordinate a number of system components that can work concurrently, for example the terminals in the individual travel agencies accessing a central database or the sensors and signals used in a distributed railway signaling system. So to be able to verify such programs we need to have at our disposal methods that allow us to deal with correctness of concurrent programs, as well.

Structure of This Book

The aim of this book is to provide a systematic exposition of one of the most common approaches to program verification. This approach is usually called assertional, because it relies on the use of assertions that are attached to program control points. Starting from a simple class of sequential pro-

grams, known as **while** programs, we proceed in a systematic manner in two directions:

- to more complex classes of sequential programs including recursive procedures and objects, and
- to concurrent programs, both parallel and distributed.

We consider here sequential programs in the form of deterministic and nondeterministic programs, and concurrent programs in the form of parallel and distributed programs. Deterministic programs cover **while** programs, recursive programs, and a simple class of object-oriented programs. Nondeterministic programs are used to analyze concurrent programs and the concept of fairness by means of program transformations. Parallel programs consist of several sequential components that can access shared memory. By contrast, distributed programs consist of components with local memory that can communicate only by sending and receiving messages.

For each of these classes of programs their input/output behavior in the sense of so-called partial and total correctness is studied. For the verification of these correctness properties an axiomatic approach involving assertions is used. This approach was initiated by Hoare in 1969 for deterministic programs and extended by various researchers to other classes of programs. It is combined here with the use of program transformations.

For each class of programs a uniform presentation is provided. After defining the syntax we introduce a structured operational semantics as originally proposed by Hennessy and Plotkin in 1979 and further developed by Plotkin in 1981. Then proof systems for the verification of partial and total correctness are introduced, which are formally justified in the corresponding soundness theorems.

The use of these proof systems is demonstrated with the help of case studies. In particular, solutions to classical problems such as producer/consumer and mutual exclusion are formally verified. Each chapter concludes with a list of exercises and bibliographic remarks.

The exposition assumes elementary knowledge of programming languages and logic. Therefore this book belongs to the area of programming languages but at the same time it is firmly based on mathematical logic. All prerequisites are provided in the preparatory **Chapter 2** of Part I.

In Part II of the book we study deterministic programs. In **Chapter 3** Hoare's approach to program verification is explained for **while** programs. Next, we move to the more ambitious structuring concepts of recursive and object-oriented programs. First, parameterless recursive procedures are studied in **Chapter 4**, and then call-by-value parameters are added in **Chapter 5**. These two chapters are taken as preparations to study a class of object-oriented programs in **Chapter 6**. This chapter is based on the work of the second author initiated in 1990, but the presentation is entirely new.

In Part III of the book we study parallel programs with shared variables. Since these are much more difficult to deal with than sequential programs,

they are introduced in a stepwise manner in **Chapters 7, 8, and 9**. We base our presentation on the approach by Owicki and Gries originally proposed in 1976 and on an extension of it by the authors dealing with total correctness.

In Part IV we turn to nondeterministic and distributed programs. Nondeterministic sequential programs are studied in **Chapter 10**. The presentation is based on the work of Dijkstra from 1976 and Gries from 1981. The study of this class of programs also serves as a preparation for dealing with distributed programs in **Chapter 11**. The verification method presented there is based on a transformation of distributed programs into nondeterministic ones proposed by the first author in 1986. In **Chapter 12** the issue of fairness is studied in the framework of nondeterministic programs. The approach is based on the method of explicit schedulers developed by the first and third authors in 1983.

Teaching from This Book

This book is appropriate for either a one- or two-semester introductory course on program verification for upper division undergraduate studies or for graduate studies.

In the first lecture the zero search example in Chapter 1 should be discussed. This example demonstrates which subtle errors can arise during the design of parallel programs. Next we recommend moving on to Chapter 3 on **while** programs and before each of the sections on syntax, semantics and verification, to refer to the corresponding sections of the preparatory Chapter 2.

After Chapter 3 there are three natural alternatives to continue. The first alternative is to proceed with more ambitious classes of sequential programs, i.e., recursive programs in Chapters 4 and 5 and then object-oriented programs in Chapter 6. The second alternative is to proceed immediately to parallel programs in Chapters 7, 8, and 9. The third alternative is to move immediately to nondeterministic programs in Chapter 10 and then to distributed programs in Chapter 11. We remark that one section of Chapter 10 can be studied only after the chapters on parallel programs.

Chapter 12 on fairness covers a more advanced topic and can be used during specialized seminars. Of course, it is also possible to follow the chapters in the sequential order as they are presented in the book.

This text may also be used as an introduction to operational semantics. We present below outlines of possible one-semester courses that can be taught using this book. The dependencies of the chapters are shown in Fig. 0.1.

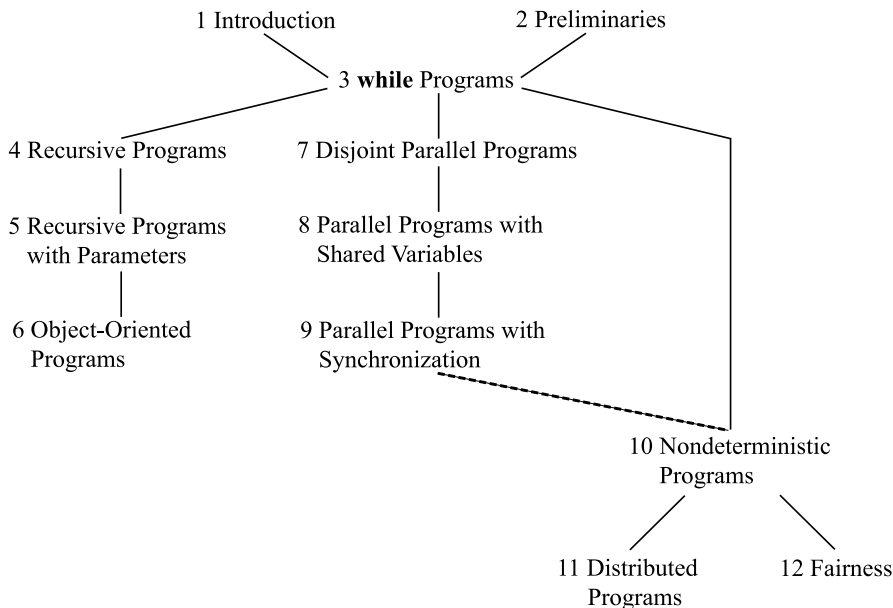


Fig. 0.1 Dependencies of chapters. In Chapter 10 only Section 10.6 depends on Chapter 9.

Changes in the Third Edition

The present, third edition of this book comes with a new co-author, Frank S. de Boer, and with an additional topic that for many years has been at the heart of his research: verification of object-oriented programs. Since this is a notoriously difficult topic, we approach it in a stepwise manner and in a setting where the notational complexity is kept at a minimum. This design decision has led us to add three new chapters to our book.

- In Chapter 4 we introduce a class of recursive programs that extends deterministic programs by parameterless procedures. Verifying such programs makes use of proofs from assumptions (about calls of recursive procedures) that are discharged later on.
- In Chapter 5 this class is extended to the recursive procedures with call-by-value parameters. Semantically, this necessitates the concept of a stack for storing the values of the actual parameters of recursively called procedures. We capture this concept by using a block statement and a corresponding semantic transition rule that models the desired stack behavior implicitly.
- In Chapter 6 object-oriented programs are studied in a minimal setting where we focus on the following main characteristics of objects: they pos-

sess (and encapsulate) their own local variables and interact via method calls, and objects can be dynamically created.

To integrate these new chapters into the original text, we made various changes in the preceding Chapters 2 and 3. For example, in Chapter 3 parallel assignments and failure statements are introduced, and a correctness proof of a program for partitioning an array is given as a preparation for the case study of the *Quicksort* algorithm in Chapter 5. Also, in Chapter 10 the transformation of parallel programs into nondeterministic programs is now defined in a formal way. Also the references have been updated.

Acknowledgments

The authors of this book have collaborated, often together with other colleagues, on the topic of program verification since 1979. During this time we have benefited very much from discussions with Pierre America, Jaco de Bakker, Luc Bougé, Ed Clarke, Werner Damm, Henning Dierks, Edsger W. Dijkstra, Nissim Francez, David Gries, Tony Hoare, Shmuel Katz, Leslie Lamport, Hans Langmaack, Jay Misra, Cees Pierik, Andreas Podelski, Amir Pnueli, Gordon Plotkin, Anders P. Ravn, Willem Paul de Roever, Fred Schneider, Jonathan Stavi and Jeffery Zucker. Many thanks to all of them.

For the third edition Maarten Versteegh helped us with the migration of files to adapt to the new style file. Alma Apt produced all the drawings in this edition.

The bibliography style used in this book has been designed by Sam Buss; Anne Troelstra deserves credit for drawing our attention to it.

Finally, we would like to thank the staff of Springer-Verlag, in particular Simon Rees and Wayne Wheeler, for the efficient and professional handling of all the stages of the production of this book. The T_EX support group of Springer, in particular Monsurate Rajiv, was most helpful.

Amsterdam, The Netherlands,
Oldenburg, Germany,

Krzysztof R. Apt and Frank S. de Boer
Ernst-Rüdiger Olderog

Outlines of One-Semester Courses

PREREQUISITES: Chapter 2.

Course on Program Semantics

Class of programs	Syntax	Semantics
while programs	3.1	3.2
Recursive programs	4.1	4.2
Recursive programs with parameters	5.1	5.2
Object-oriented programs	6.1	6.2
Disjoint parallel programs	7.1	7.2
Parallel programs with shared variables	8.1, 8.2	8.3
Parallel programs with synchronization	9.1	9.2
Nondeterministic programs	10.1	10.2
Distributed programs	11.1	11.2
Fairness	12.1	12.2

Course on Program Verification

Class of programs	Syntax	Semantics	Proof theory
while programs	3.1	3.2	3.3, 3.4, 3.10
Recursive programs	4.1	4.2	4.3, 4.4
Recursive programs with parameters	5.1	5.2	5.3
Object-oriented programs	6.1	6.2	6.3, 6.4, 6.5, 6.6
Disjoint parallel programs	7.1	7.2	7.3
Parallel programs with shared variables	8.1, 8.2	8.3	8.4, 8.5
Parallel programs with synchronization	9.1	9.2	9.3
Nondeterministic programs	10.1	10.2	10.4
Distributed programs	11.1	11.2	11.4

Course Towards Object-Oriented Program Verification

Class of programs	Syntax	Semantics	Proof theory	Case studies
while programs	3.1	3.2	3.3, 3.4	3.9
Recursive programs	4.1	4.2	4.3, 4.4	4.5
Recursive programs with parameters	5.1	5.2	5.3	5.4
Object-oriented programs	6.1	6.2	6.3, 6.4	6.8

Course on Concurrent Program Verification

Class of programs	Syntax	Semantics	Proof theory	Case studies
while programs	3.1	3.2	3.3, 3.4	3.9
Disjoint parallel programs	7.1	7.2	7.3	7.4
Parallel programs with shared variables	8.1, 8.2	8.3	8.4, 8.5	8.6
Parallel programs with synchronization	9.1	9.2	9.3	9.4, 9.5
Nondeterministic programs	10.1	10.2	10.4	10.5
Distributed programs	11.1	11.2	11.4	11.5

Course on Program Verification with Emphasis on Case Studies

Class of programs	Syntax	Proof theory	Case studies
while programs	3.1	3.3, 3.4	3.9
Recursive programs	4.1	4.3, 4.4	4.5
Recursive programs with parameters	5.1	5.4	5.4
Object-oriented programs	6.1	6.3–6.5	6.8
Disjoint parallel programs	7.1	7.3	7.4
Parallel programs with shared variables	8.1, 8.2	8.4, 8.5	8.6
Parallel programs with synchronization	9.1	9.3	9.4, 9.5
Nondeterministic programs	10.1, 10.3	10.4	10.5
Distributed programs	11.1	11.4	11.5

Contents

<i>Endorsements</i>	v
<i>Foreword</i>	vii
<i>Preface</i>	ix
Outlines of One-semester Courses	xiv
 Part I In the Beginning	
1 <i>Introduction</i>	3
1.1 An Example of a Concurrent Program	4
Solution 1	4
Solution 2	5
Solution 3	6
Solution 4	8
Solution 5	9
Solution 6	10
1.2 Program Correctness	11
1.3 Structure of this Book	13
1.4 Automating Program Verification	16
1.5 Assertional Methods in Practice	17
 2 <i>Preliminaries</i>	19
2.1 Mathematical Notation	21
Sets	21
Tuples	22
Relations	23
Functions	23
Sequences	24
Strings	25
Proofs	26
Induction	27
Grammars	29
2.2 Typed Expressions	29
Types	29

	Variables	30
	Constants	30
	Expressions	31
	Subscripted Variables	32
2.3	Semantics of Expressions	32
	Fixed Structure	33
	States	34
	Definition of the Semantics	35
	Updates of States	36
2.4	Formal Proof Systems	38
2.5	Assertions	39
2.6	Semantics of Assertions	41
2.7	Substitution	42
	Simultaneous Substitution	45
2.8	Substitution Lemma	47
2.9	Exercises	50
2.10	Bibliographic Remarks	51

Part II Deterministic Programs

3	while <i>Programs</i>	55
3.1	Syntax	57
3.2	Semantics	58
	Properties of Semantics	62
3.3	Verification	63
	Partial Correctness	65
	Total Correctness	70
	Decomposition	73
	Soundness	73
3.4	Proof Outlines	79
	Partial Correctness	79
	Total Correctness	83
3.5	Completeness	85
3.6	Parallel Assignment	91
3.7	Failure Statement	94
3.8	Auxiliary Axioms and Rules	97
3.9	Case Study: Partitioning an Array	99
3.10	Systematic Development of Correct Programs	113
	Summation Problem	115
3.11	Case Study: Minimum-Sum Section Problem	116
3.12	Exercises	121
3.13	Bibliographic Remarks	124

4	<i>Recursive Programs</i>	127
4.1	Syntax	129
4.2	Semantics	129
	Properties of the Semantics	131
4.3	Verification	132
	Partial Correctness	132
	Total Correctness	134
	Decomposition	137
	Discussion	138
	Soundness	139
4.4	Case Study: Binary Search	144
	Partial Correctness	145
	Total Correctness	147
4.5	Exercises	149
4.6	Bibliographic Remarks	150
5	<i>Recursive Programs with Parameters</i>	151
5.1	Syntax	152
5.2	Semantics	154
5.3	Verification	157
	Partial Correctness: Non-recursive Procedures	158
	Partial Correctness: Recursive Procedures	162
	Modularity	165
	Total Correctness	165
	Soundness	167
5.4	Case Study: <i>Quicksort</i>	172
	Formal Problem Specification	173
	Properties of <i>Partition</i>	173
	Auxiliary Proof: Permutation Property	174
	Auxiliary Proof: Sorting Property	175
	Total Correctness	180
5.5	Exercises	182
5.6	Bibliographic Remarks	182
6	<i>Object-Oriented Programs</i>	185
6.1	Syntax	187
	Local Expressions	187
	Statements and Programs	188
6.2	Semantics	192
	Semantics of Local Expressions	192
	Updates of States	194
	Semantics of Statements and Programs	195
6.3	Assertions	197
	Substitution	199
6.4	Verification	200

	Partial Correctness	201
	Total Correctness.....	204
6.5	Adding Parameters	206
	Semantics	207
	Partial Correctness	208
	Total Correctness.....	210
6.6	Transformation of Object-Oriented Programs	211
	Soundness	214
6.7	Object Creation.....	217
	Semantics	218
	Assertions	219
	Verification	223
	Soundness	225
6.8	Case Study: Zero Search in Linked List.....	226
	Partial Correctness	226
	Total Correctness.....	229
6.9	Case Study: Insertion into a Linked List	232
6.10	Exercises	238
6.11	Bibliographic Remarks	240

Part III Parallel Programs

7	<i>Disjoint Parallel Programs</i>	245
7.1	Syntax	247
7.2	Semantics	248
	Determinism	249
	Sequentialization	252
7.3	Verification	253
	Parallel Composition.....	254
	Auxiliary Variables	256
	Soundness	259
7.4	Case Study: Find Positive Element	261
7.5	Exercises	264
7.6	Bibliographic Remarks	266
8	<i>Parallel Programs with Shared Variables</i>	267
8.1	Access to Shared Variables	269
8.2	Syntax	270
8.3	Semantics	271
	Atomicity	272
8.4	Verification: Partial Correctness	274
	Component Programs	274
	No Compositionality of Input/Output Behavior	275
	Parallel Composition: Interference Freedom	276
	Auxiliary Variables Needed	279

	Soundness	282
8.5	Verification: Total Correctness	284
	Component Programs	284
	Parallel Composition: Interference Freedom	286
	Soundness	288
	Discussion	289
8.6	Case Study: Find Positive Element More Quickly	291
8.7	Allowing More Points of Interference	294
8.8	Case Study: Parallel Zero Search	299
	Step 1. Simplifying the program	299
	Step 2. Proving partial correctness	300
8.9	Exercises	303
8.10	Bibliographic Remarks	305
9	<i>Parallel Programs with Synchronization</i>	307
9.1	Syntax	309
9.2	Semantics	310
9.3	Verification	311
	Partial Correctness	311
	Weak Total Correctness	313
	Total Correctness	314
	Soundness	316
9.4	Case Study: Producer/Consumer Problem	319
9.5	Case Study: The Mutual Exclusion Problem	324
	Problem Formulation	324
	Verification	326
	A Busy Wait Solution	327
	A Solution Using Semaphores	331
9.6	Allowing More Points of Interference	334
9.7	Case Study: Synchronized Zero Search	335
	Step 1. Simplifying the Program	336
	Step 2. Decomposing Total Correctness	337
	Step 3. Proving Termination	337
	Step 4. Proving Partial Correctness	342
9.8	Exercises	344
9.9	Bibliographic Remarks	345

Part IV Nondeterministic and Distributed Programs

10	<i>Nondeterministic Programs</i>	349
10.1	Syntax	351
10.2	Semantics	352
	Properties of Semantics	353
10.3	Why Are Nondeterministic Programs Useful?	354
	Symmetry	355

	Nondeterminism	355
	Failures	356
	Modeling Concurrency	356
10.4	Verification	357
	Partial Correctness	357
	Total Correctness	357
	Soundness	359
10.5	Case Study: The Welfare Crook Problem	360
10.6	Transformation of Parallel Programs	363
10.7	Exercises	368
10.8	Bibliographic Remarks	370
11	<i>Distributed Programs</i>	373
11.1	Syntax	375
	Sequential Processes	375
	Distributed Programs	376
11.2	Semantics	380
11.3	Transformation into Nondeterministic Programs	382
	Semantic Relationship Between S and $T(S)$	382
	Proof of the Sequentialization Theorem	385
11.4	Verification	390
	Partial Correctness	390
	Weak Total Correctness	391
	Total Correctness	391
	Proof Systems	392
	Soundness	393
11.5	Case Study: A Transmission Problem	396
	Step 1. Decomposing Total Correctness	397
	Step 2. Proving Partial Correctness	397
	Step 3. Proving Absence of Failures and of Divergence	399
	Step 4. Proving Deadlock Freedom	400
11.6	Exercises	402
11.7	Bibliographic Remarks	405
12	<i>Fairness</i>	407
12.1	The Concept of Fairness	409
	Selections and Runs	410
	Fair Nondeterminism Semantics	412
12.2	Transformational Semantics	413
12.3	Well-Founded Structures	413
12.4	Random Assignment	414
	Semantics	415
	Verification	415
12.5	Schedulers	419
	The Scheduler FAIR	421

	The Scheduler RORO	424
	The Scheduler QUEUE	426
12.6	Transformation	427
12.7	Verification	430
	Total Correctness.....	430
	Soundness	438
12.8	Case Study: Zero Search	442
12.9	Case Study: Asynchronous Fixed Point Computation	446
12.10	Exercises	452
12.11	Bibliographic Remarks	455
A	<i>Semantics</i>	457
B	<i>Axioms and Proof Rules</i>	459
C	<i>Proof Systems</i>	471
D	<i>Proof Outlines</i>	475
	<i>References</i>	477
	<i>Index</i>	491
	<i>Author Index</i>	497
	<i>Symbol Index</i>	501