

Hans Hüttel

---

# Transitions and Trees

An Introduction to Structural  
Operational Semantics



CAMBRIDGE

CAMBRIDGE

more information - [www.cambridge.org/9780521197465](http://www.cambridge.org/9780521197465)



## **Transitions and Trees**

### **An Introduction to Structural Operational Semantics**

Structural operational semantics is a simple, yet powerful mathematical theory for describing the behaviour of programs in an implementation-independent manner. This book provides a self-contained introduction to structural operational semantics, featuring semantic definitions using big-step and small-step semantics of many standard programming language constructs, including control structures, structured declarations and objects, parameter mechanisms and procedural abstraction, concurrency, non-determinism and the features of functional programming languages. Along the way, the text introduces and applies the relevant proof techniques, including forms of induction and notions of semantic equivalence (including bisimilarity).

Thoroughly class-tested, this book has evolved from lecture notes used by the author over a 10-year period at Aalborg University to teach undergraduate and graduate students. The result is a thorough introduction that makes the subject clear to students and computing professionals without sacrificing its rigour. No experience with any specific programming language is required.

HANS HÜTTEL is Associate Professor in the Department of Computer Science at Aalborg University, Denmark.



# Transitions and Trees

An Introduction to Structural Operational Semantics

HANS HÜTTEL

*Aalborg University, Denmark*



**CAMBRIDGE**  
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS  
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore,  
São Paulo, Delhi, Dubai, Tokyo

Cambridge University Press  
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

[www.cambridge.org](http://www.cambridge.org)

Information on this title: [www.cambridge.org/9780521197465](http://www.cambridge.org/9780521197465)

© H. Hüttel 2010

This publication is in copyright. Subject to statutory exception  
and to the provisions of relevant collective licensing agreements,  
no reproduction of any part may take place without the written  
permission of Cambridge University Press.

First published 2010

Printed in the United Kingdom at the University Press, Cambridge

*A catalogue record for this publication is available from the British Library*

ISBN 978-0-521-19746-5 Hardback

ISBN 978-0-521-14709-5 Paperback

Additional resources for this publication at [www.operationalsemantics.net](http://www.operationalsemantics.net)

---

Cambridge University Press has no responsibility for the persistence or  
accuracy of URLs for external or third-party internet websites referred to  
in this publication, and does not guarantee that any content on such  
websites is, or will remain, accurate or appropriate.

---

# Contents

<i>Preface</i>	<i>page</i>	ix
<i>About the illustrations</i>		xiii
<i>List of illustrations</i>		xiv
<i>List of tables</i>		xv
<b>PART I BACKGROUND</b>		<b>1</b>
<b>1 A question of semantics</b>		<b>3</b>
1.1 Semantics is the study of meaning		3
1.2 Examples from the history of programming languages		4
1.3 Different approaches to program semantics		6
1.4 Applications of program semantics		10
<b>2 Mathematical preliminaries</b>		<b>16</b>
2.1 Mathematical induction		16
2.2 Logical notation		17
2.3 Sets		19
2.4 Operations on sets		20
2.5 Relations		21
2.6 Functions		22
<b>PART II FIRST EXAMPLES</b>		<b>25</b>
<b>3 The basic principles</b>		<b>27</b>
3.1 Abstract syntax		27
3.2 Transition systems		30
3.3 Big-step vs. small-step semantics		31
3.4 Operational semantics of arithmetic expressions		31

3.5	Proving properties	38
3.6	A semantics of Boolean expressions	39
3.7	The elements of an operational semantics	40
<b>4</b>	<b>Basic imperative statements</b>	<b>43</b>
4.1	Program states	43
4.2	A big-step semantics of statements	45
4.3	A small-step semantics of statements in <b>Bims</b>	53
4.4	Equivalence of the two semantics	55
4.5	Two important proof techniques	60
	<b>PART III LANGUAGE CONSTRUCTS</b>	<b>63</b>
<b>5</b>	<b>Control structures</b>	<b>65</b>
5.1	Some general assumptions	65
5.2	Loop constructs	66
5.3	Semantic equivalence	70
5.4	Abnormal termination	72
5.5	Nondeterminism	73
5.6	Concurrency	76
<b>6</b>	<b>Blocks and procedures (1)</b>	<b>79</b>
6.1	Abstract syntax of <b>Bip</b>	79
6.2	The environment–store model	80
6.3	Arithmetic and Boolean expressions	83
6.4	Declarations	84
6.5	Statements	85
6.6	Scope rules	86
<b>7</b>	<b>Parameters</b>	<b>94</b>
7.1	The language <b>Bump</b>	94
7.2	Call-by-reference	96
7.3	On recursive and non-recursive procedure calls	97
7.4	Call-by-value	99
7.5	Call-by-name	100
7.6	A comparison of parameter mechanisms	110
<b>8</b>	<b>Concurrent communicating processes</b>	<b>113</b>
8.1	Channel-based communication – <b>Cab</b>	113
8.2	Global and local behaviour	114
8.3	Synchronous communication in <b>Cab</b>	115
8.4	Other communication models	119

8.5	Bisimulation equivalence	122
8.6	Channels as data – the $\pi$ -calculus	123
<b>9</b>	<b>Structured declarations</b>	134
9.1	Records	134
9.2	The language <b>Bur</b>	135
9.3	The class-based language <b>Coat</b>	142
<b>10</b>	<b>Blocks and procedures (2)</b>	154
10.1	Run-time stacks	154
10.2	Declarations	155
10.3	Statements	155
<b>11</b>	<b>Concurrent object-oriented languages</b>	161
11.1	The language <b>Cola</b>	161
11.2	A small-step semantics of concurrent behaviour	163
11.3	Transition systems	163
<b>12</b>	<b>Functional programming languages</b>	171
12.1	What is a functional programming language?	171
12.2	Historical background	173
12.3	The $\lambda$ -calculus	174
12.4	<b>Flan</b> – a simple functional language	177
12.5	Further reading	181
	<b>PART IV RELATED TOPICS</b>	183
<b>13</b>	<b>Typed programming languages</b>	185
13.1	Type systems	185
13.2	Typed <b>Bump</b>	187
13.3	Typed <b>Flan</b>	198
13.4	Type polymorphism and type inference	209
<b>14</b>	<b>An introduction to denotational semantics</b>	211
14.1	Background	211
14.2	$\lambda$ -Notation	212
14.3	Basic ideas	214
14.4	Denotational semantics of statements	216
14.5	Further reading	220
<b>15</b>	<b>Recursive definitions</b>	222
15.1	A first example	222
15.2	A recursive definition specifies a fixed-point	224
15.3	The fixed-point theorem	225

15.4	How to apply the fixed-point theorem	231
15.5	Examples of cpos	232
15.6	Examples of continuous functions	236
15.7	Examples of computations of fixed-points	240
15.8	An equivalence result	241
15.9	Other applications	246
15.10	Further reading	248
<b><i>Appendix A</i></b>	<b>A big-step semantics of Bip</b>	249
<b><i>Appendix B</i></b>	<b>Implementing semantic definitions in SML</b>	257
	<i>References</i>	264
	<i>Index</i>	269

# Preface

## About this book

This is a book about structural operational semantics; more precisely it is a book that describes how this approach to semantics can be used to describe common programming language constructs and to reason about the behaviour of programs.

The text grew out of the lecture notes that I have used over a period of more than 10 years in the course *Syntax and semantics* which is taught to all students following the various degree programmes in computer science at Aalborg University. What began as a 10-page set of notes in Danish is now a textbook in English.

The book also includes chapters on related material, namely short introductions to type systems, denotational semantics and the mathematics necessary to understand recursive definitions.

## Related work

This work was inspired by lecture notes by Plotkin (1981) (also written in Denmark), where this approach to programming language semantics was first presented.

The topic of structural operational semantics also appears in later books, three of which I will mention here.

Reynolds' book (Reynolds, 1999) is an excellent text that covers some of the same topics as this book but uses denotational and axiomatic semantics as well as structural operational semantics.

The book by Winskel (1993) is another very good textbook that covers many of the same topics as Reynolds' book.

Finally, I should mention Nielson and Nielson (2007), which introduces

and relates denotational, axiomatic and structural operational semantics and then gives an introduction to how these can be used in connection with static program analysis.

The book that you are now reading differs from the ones just mentioned in three important ways. First, the main topic is exclusively that of structural operational semantics. Second, both Reynolds and Winskel introduce domain theory early on; this book aims at developing the theory of structural operational semantics and making use of it with mathematical prerequisites of a more modest nature. Third, unlike the book by the Nielsons, the focus here is not that of program analysis. Instead, it is on how operational semantics can be used to describe common features of programming languages.

### What you need to know in advance

This text is *not* intended as an introduction to programming; if you are a reader expecting this to be the case, you will probably be disappointed! The ideal reader should already have some experience with programming in a high-level imperative programming language such as C, Java or Pascal.

Programming language semantics is a mathematical theory. Therefore, the reader should also have some mathematical maturity. In particular, you should be familiar with basic notions of discrete mathematics – sets, functions and graphs and the proof techniques of proof by induction and proof by contradiction. Chapter 2 gives a short overview of some of this material. There are several good textbooks that you can consult as a supplement; one that I would recommend in particular is Velleman’s book (Velleman, 2006).

### Ways through the book

The book falls into four parts. The first two parts must be covered in any course for which this book is the main text, since the contents of these first four chapters are necessary to understand the material in the rest of the book. After that, there are the following dependences:

- Chapter 7 and Chapter 9 are independent of each other but both extend the language introduced in Chapter 6,
- Chapter 8 assumes knowledge of the parallel operator introduced in Chapter 5,
- Chapter 10 on small-step semantics for procedures and blocks also assumes knowledge of Chapter 6,
- Chapter 11 assumes knowledge of Chapters 8 and 10 and finally

- Chapter 15 assumes knowledge of the contents of Chapter 14.

### Problems and thoughts

To learn a mathematical subject, one should of course read the text carefully but also learn to apply the content. For this reason, there are quite a few problems scattered throughout the text. As a rule of thumb, a problem will appear at the point in the text where it becomes relevant. I have chosen this approach since I would like you, the reader, to focus on the connection between the problem and the context in which it appears. You can read most of the text without solving the problems but I encourage you to solve as many of them as possible. In some places, I have put problems that are important for understanding the text and they are then marked as such.

I have also introduced mini-problems which I call **A moment's thought**. Here, the idea is to make you think carefully about what you have just read. Do *not* read the text without finding the answers to these mini-problems.

### Related resources

The book has its own website, <http://www.operationalsemantics.net>, which has more information, including hints to the mini-problems. The website also holds information about the Danish-language version of the book, *Pilen ved træets rod*, including how to obtain a copy of it.

### Acknowledgements

This book grew out of the years I have spent teaching, so, first, I would like to thank the students who have lived with the various incarnations of this text over the years and have made many useful comments that have helped improve and shape its content and form.

Second, I want to thank the people who have inspired me to reflect on the task of teaching mathematical subjects and teaching in general over the years: Jens Friis Jørgensen, Steffen Lauritzen, Finn Verner Jensen, Anette Kolmos, Helle Alrø and Ole Skovsmose.

Third, I would like thank all those who helped me make this book a reality. My thanks go to the people and organizations who have kindly allowed me to use the pictures in Chapter 1 and to David Tranah from Cambridge University Press for his encouragement.

A number of colleagues read parts of the manuscript and provided me with lots of important feedback. Special thanks are due to Denis Bertelsen,

Morten Dahl, Ulrich Fahrenberg, Morten Kühnrich, Michael Pedersen, Willard Rafnsson and last, but by no means least, Gordon Plotkin.

On an entirely personal level, there are others who also deserve thanks. Over the years, I have come to know many inspiring people through my extracurricular activities in human rights activism and music and, most recently, through the extended family of sisters and brothers that I now have. I am very grateful for knowing you all.

Finally, and most importantly, I want to thank my wife Maria and our daughter Nadia for being in my life. This book is dedicated to you.

## About the illustrations

- The picture of Alfred Tarski on p. 4 is by George M. Bergman and used courtesy of Wikimedia Commons under the GNU Free Documentation License.
- The picture of Dana Scott on p. 7 is used courtesy of Dana Scott.
- The picture of Christopher Strachey on p. 7 is used courtesy of Martin Campbell-Kelly.
- The picture of Gordon Plotkin on p. 8 is used courtesy of Gordon Plotkin.
- The picture of Robin Milner on p. 9 is used courtesy of Robin Milner.
- The picture of Tony Hoare on p. 9 is used courtesy of Tony Hoare.
- The picture of Joseph Goguen on p. 10 is used courtesy of Healfdene Goguen.
- The pictures of Ariane 5 on p. 12 are used courtesy of the ESA/CNES and are ©AFP/Patrick Hertzog, 1996.
- The picture of the Mars Climate Orbiter on p. 14 is used courtesy of NASA.

# Illustrations

1.1	Alfred Tarski	4
1.2	An ALGOL 60 procedure. What is its intended behaviour?	5
1.3	Dana Scott (left) and Christopher Strachey (right)	7
1.4	Gordon Plotkin	8
1.5	Robin Milner	9
1.6	Tony Hoare	9
1.7	Joseph Goguen	10
1.8	The start and end of the maiden voyage of the Ariane 5	12
1.9	The Mars Climate Orbiter	14
3.1	A very small transition system	31
3.2	Derivation tree for a big-step transition for $(\underline{2}+\underline{3})*(\underline{4}+\underline{9}) \rightarrow 65$	35
3.3	Comparison between the derivation trees for the individual steps of a small-step transition sequence and that of a big-step transition	41
6.1	Example of a variable environment and a store	82
6.2	An example <b>Bip</b> statement whose behaviour is dependent on the choice of scope rules	88
7.1	A <b>Bump</b> statement with recursive calls	98
7.2	Exploiting call-by-name for computing the sum $\sum_{i=1}^{10} i^2$	104
7.3	A <b>Bump</b> statement showing where name clashes could occur as a result of an incorrectly defined substitution	108
9.1	A program with nested record declarations	135
9.2	A <b>Coat</b> program example	144
13.1	A type system is an overapproximation of safety	198
15.1	Part of the Hasse diagram for $(\mathbb{N}, \leq)$	226
15.2	A Hasse diagram for $(\mathcal{P}(\{1, 2, 3\}), \subseteq)$	233

# Tables

3.1	Abstract syntax of <b>Bims</b>	29
3.2	Big-step transition rules for <b>Aexp</b>	33
3.3	Small-step transition rules for <b>Aexp</b>	37
3.4	Big-step transition rules for <b>Bexp</b>	40
4.1	Big-step transition rules for <b>Aexp</b>	45
4.2	Big-step transition rules for <b>Bexp</b>	46
4.3	Big-step transition rules for <b>Stm</b>	47
4.4	Small-step transition rules for <b>Stm</b>	53
5.1	Big-step transition rules for repeat-loops	67
5.2	Big-step transition rules for for-loops	72
5.3	Big-step transition rules for the <b>or</b> -statement	74
5.4	Small-step transition rules for the <b>or</b> -statement	75
5.5	Small-step semantics of the <b>par</b> -statement	77
5.6	An attempt at big-step transition rules for the <b>par</b> -statement	77
6.1	Big-step operational semantics of <b>Aexp</b> using the environment-store model	83
6.2	Big-step semantics of variable declarations	84
6.3	Big-step transition rules for <b>Bip</b> statements (except procedure calls)	87
6.4	Transition rules for procedure declarations (assuming fully dynamic scope rules)	89
6.5	Transition rule for procedure calls (assuming fully dynamic scope rules)	89
6.6	Transition rules for procedure declarations assuming mixed scope rules (dynamic for variables, static for procedures)	90
6.7	Transition rules for procedure calls assuming mixed scope rules (dynamic for variables, static for procedures)	91
6.8	Transition rules for procedure declarations assuming fully static scope rules	92

6.9	Transition rules for procedure calls assuming fully static scope rules	93
7.1	Rules for declaring procedures with a single parameter assuming static scope rules	96
7.2	Transition rule for calling a call-by-reference procedure	97
7.3	Revised transition rule for procedure calls that allow recursive calls	99
7.4	Transition rules for procedure calls using call-by-value	101
7.5	Transition rules for declaration of call-by-name procedures assuming fully static scope rules	105
7.6	Transition rules for procedure calls using call-by-name	105
7.7	Substitution in statements	111
8.1	Transition rules defining local transitions	116
8.2	Rules defining the capability semantics	117
8.3	Synchronous communication: transition rules for the global level	118
8.4	Asynchronous communication: communication capabilities	120
8.5	Asynchronous communication: the global level	121
8.6	The structural congruence rules	128
8.7	Rules of the reduction semantics of the $\pi$ -calculus	130
8.8	Rules of the labelled semantics of the $\pi$ -calculus	132
9.1	Transition rules for generalized variables	137
9.2	Transition rules for arithmetic expressions	138
9.3	Transition rules for variable declarations	139
9.4	Transition rules for procedure declarations	139
9.5	Transition rules for generalized procedure names	140
9.6	Transition rules for record declarations	140
9.7	Transition rules for statements in <b>Bur</b>	141
9.8	The semantics of class declarations	146
9.9	The semantics of variable declarations	147
9.10	The semantics of method declarations	147
9.11	The semantics of object declarations	148
9.12	The semantics of object sequences	149
9.13	The semantics of object expressions	150
9.14	Evaluating extended variables (in the semantics of arithmetic expressions)	150
9.15	Important transition rules for statements	151
9.16	Transition rule for programs	152
10.1	Transition rules for statements other than procedure calls	157
10.2	Transition rules for procedure calls assuming static scope rules	159
10.3	Transition rule for procedure calls assuming dynamic scope rules	160
11.1	Transition rules for the local transition system	165
11.2	Transition rules of the labelled transition systems	166
11.3	Transition rules for the global level (1) – initialization	167

11.4	Transition rules for the global level (2) – the connection between global and local behaviour	168
11.5	Transition rules for the global level (3) – rendezvous	169
12.1	Big-step semantics for <b>Flan</b>	179
12.2	Some of the small-step semantics of <b>Flan</b>	180
13.1	Big-step operational semantics of <b>Exp</b> (arithmetic part)	188
13.2	Big-step transition rules for Bump statements (except procedure calls)	189
13.3	Type rules for <b>Bump</b> expressions	191
13.4	Type rules for variable and procedure declarations in <b>Bump</b>	191
13.5	Type rules for <b>Bump</b> statements	192
13.6	The error predicate for variable declarations	194
13.7	The error predicate for statements	195
13.8	Type rules for <b>Flan</b>	200
13.9	The small-step semantics of <b>Flan</b>	201
13.10	The type rule for <b>letrec</b>	209
A.1	Big-step operational semantics of <b>Aexp</b>	252
A.2	Big-step transition rules for $\rightarrow_b$	253
A.3	Big-step semantics of variable declarations	254
A.4	Transition rules for procedure declarations assuming fully static scope rules	254
A.5	Big-step transition rules for <b>Bip</b> statements	256



**PART I**  
BACKGROUND



# 1

## A question of semantics

The goal of this chapter is to give the reader a glimpse of the applications and problem areas that have motivated and to this day continue to inspire research in the important area of computer science known as programming language semantics.

### 1.1 Semantics is the study of meaning

Programming language semantics is the study of *mathematical models of and methods for describing and reasoning about the behaviour of programs*.

The word *semantics* has Greek roots<sup>1</sup> and was first used in linguistics. Here, one distinguishes among *syntax*, the study of the structure of languages, *semantics*, the study of meaning, and *pragmatics*, the study of the use of language.

In computer science we make a similar distinction between syntax and semantics. The languages that we are interested in are *programming languages* in a very general sense. The ‘meaning’ of a program is its behaviour, and for this reason programming language semantics is the part of programming language theory devoted to the study of *program behaviour*.

Programming language semantics is concerned only with purely internal aspects of program behaviour, namely what happens within a running program. Program semantics does not claim to be able to address other aspects of program behaviour – e.g. whether or not a program is user-friendly or useful.

In this book, when we speak of semantics, we think of *formal semantics*,

<sup>1</sup> The Greek word (transliterated) is *semantikós*, meaning ‘significant’. The English word ‘semantics’ is a singular form, as are ‘physics’, ‘mathematics’ and other words that have similar Greek roots.




Figure 1.1 Alfred Tarski

understood as an approach to semantics that relies on precise mathematical definitions.

Formal semantics arose in the early twentieth century in the context of mathematical logic. An early goal of mathematical logic was to provide a precise mathematical description of the language of mathematics, including the notion of truth. An important contributor in this area was the logician Alfred Tarski (Figure 1.1) (Tarski, 1935).

Many of the first insights and a lot of the fundamental terminology used in programming language semantics can be traced back to the work of Tarski. For instance, the important notion of *compositionality* – that the meaning of a composite language term should be defined using the meanings of its immediate constituents – is due to him. So is the insight that we need to use another language, a *metalanguage*, to define the semantics of our target language.

## 1.2 Examples from the history of programming languages

The area of programming language semantics came into existence in the late 1960s. It was born of the many problems that programming language designers and implementors encountered when trying to describe various constructs in both new and existing programming language.

The general conclusion that emerged was that an informal semantics, however precise it may seem, is not sufficient when it comes to defining the behaviour of programs.

### 1.2.1 ALGOL 60

The programming language ALGOL 60 was first documented in a paper from 1960 (Backus and Naur, 1960), now often referred to simply as ‘the ALGOL 60 report’. ALGOL 60 was in many ways a landmark in the evolution of programming languages.

Firstly, the language was the result of very careful work by a committee of prominent researchers, including John Backus, who was the creator of FORTRAN, John McCarthy, the creator of Lisp, and Peter Naur, who became the first Danish professor of computer science. Later in their careers, Backus, McCarthy and Naur all received the ACM Turing Award for their work on programming languages.

Secondly, ALGOL 60 inspired a great many subsequent languages, among them Pascal and Modula.

Thirdly, ALGOL 60 was the first programming language whose syntax was defined formally. The notation used was a variant of context-free grammars, later known as Backus–Naur Normal Form (BNF).

However, as far as the semantics of ALGOL 60 is concerned, Backus and his colleagues had to rely on very detailed descriptions in English, since there were as yet no general mathematical theories of program behaviour. It turned out to be the case that even a group of outstanding researchers (who for the most part were mathematicians) could not avoid being imprecise, when they did not have access to a formalized mathematical theory of program behaviour. In 1963 the ALGOL 60 committee therefore released a revised version of the ALGOL 60 report (Backus and Naur, 1963) in which they tried to resolve the ambiguities and correct the mistakes that had been found since the publication of the original ALGOL 60 report.

However, this was by no means the end of the story. In 1967, Donald E. Knuth published a paper (Knuth, 1967) in which he pointed out a number of problems that still existed in ALGOL 60.

One such problem had to do with global variables in procedures. Figure 1.2 illustrates the nature of the problem. The procedure `awkward` is a procedure returning an integer value.

```
integer procedure awkward
begin comment x is a global variable
  x := x+1
  awkward := 3
end awkward
```

Figure 1.2 An ALGOL 60 procedure. What is its intended behaviour?

The procedure `awkward` manipulates the value of a global variable and therefore has a side effect. However, the ALGOL 60 report does not explain whether or not side effects are allowed in procedures. Because of this, there is also no explanation of how arithmetic expressions should be evaluated if they contain side effects.

Let us consider a global variable `x` whose value is 5 and assume that we now want to find the value of the expression `x+awkward`. Should we evaluate `x` before or after we evaluate `awkward`? If we evaluate `x` first, the value of the expression will be 8; should we evaluate `x` after having called `awkward`, we get the value 9!

One consequence of Knuth's paper was that the ALGOL 60 committee went back to the drawing board to remove the ambiguities. The main reason why it took so long to discover these problems was that the language designers had no precise, mathematical criterion for checking whether or not all aspects of the language had been defined.

### *1.2.2 Pascal*

Pascal, a descendant of the Algol family, was created by Niklaus Wirth and first documented in a book written with Kathleen Jensen (Jensen and Wirth, 1975). Ever since then, Pascal has been a common introductory language in computer science degree programmes around the world.

Even though great care was taken in the exposition of the language features, Pascal also suffers from the problems associated with an informal semantics. In particular, there are problems with explaining scoping rules – in fact, the scoping rules are barely explained in the book. There is mention of global variables; however, nowhere in the text is it explained what a global variable is, let alone what its scope should be. Nor are there any rules that specify that a variable must be declared before it is used!

All existing implementations of Pascal assume this (except for pointer variables), but the declaration-before-use convention is not part of the original definition of the language.

## **1.3 Different approaches to program semantics**

The development of a mathematical theory of program semantics has been motivated by examples such as the ones given above. There are several ways of providing such a mathematical theory, and they turn out to be related.




Figure 1.3 Dana Scott (left) and Christopher Strachey (right)

*Denotational semantics* was the first mathematical account of program behaviour; it arose in the late 1960s (Strachey, 1966, 1967; Scott and Strachey, 1971) and was pioneered by Dana Scott and Christopher Strachey (Figure 1.3), who at the time were both working at Oxford University.

In denotational semantics, the behaviour of a program is described by defining a function that assigns meaning to every construct in the language. The meaning of a language construct is called its *denotation*. Typically, for an imperative program, the denotation will be a *state transformation*, which is a function that describes how the final values of the variables in a program are found from their initial values.

*Structural operational semantics* – the main topic of this book – came into existence around 1980 and is due to Gordon Plotkin (Figure 1.4), who gave the first account of his ideas in a set of lecture notes written during his sabbatical at Århus University in 1980 (Plotkin, 1981). An important early contribution is that of Robin Milner (Figure 1.5), who used Plotkin’s approach to give a labelled semantics to the process calculus CCS (Calculus of Communication Systems) (Milner, 1980). Plotkin (2004) gives a detailed account of the origins and early history of the area.

In structural operational semantics one specifies the behaviour of a program by defining a transition system whose transition relation describes the evaluation steps of a program. One of the underlying motivations for this approach was that it is possible to give a simple account of concurrent programs; previous attempts to give a semantic description of even simple



Figure 1.4 Gordon Plotkin

parallel programming languages had used denotational semantics and had turned out to be quite complicated.

A central insight of this approach, and one to which we shall return repeatedly throughout this book, is that one can describe the evaluation steps of a syntactic entity (such as a program) in a structural fashion, that is, by means of an inductive definition based on the abstract syntax.

*Axiomatic semantics* is due to Tony Hoare (Hoare, 1969; Apt, 1981) (Figure 1.6) and, like denotational semantics, it is a product of the late 1960s. Here one describes a language construct by means of mathematical logic. More precisely, one defines a set of rules that describe the assertions that must hold before and after the language construct has been executed.

*Algebraic semantics* is related to denotational semantics and describes the behaviour of a program using universal algebra (Guessarian, 1981; Goguen and Malcolm, 1996). The members of the research collective behind the OBJ specification language, with Joseph Goguen (Figure 1.7) as a prominent contributor, have been important figures in the development of this approach.

These four approaches to programming language semantics are not rivals. Rather, they complement each other. Some approaches are more suitable than others in certain situations. For instance, it is much easier to describe




Figure 1.5 Robin Milner



Figure 1.6 Tony Hoare

parallel and nondeterministic program behaviour using structural operational semantics than by means of denotational semantics.

There are many precise mathematical results relating the four approaches. In this book we give an example of such a result in Chapter 15, where we



Figure 1.7 Joseph Goguen

show that the structural operational semantics and the denotational semantics of the **Bims** language are equivalent in a very precise sense.

#### 1.4 Applications of program semantics

The area of program semantics has turned out to be extremely useful in situations where it is important to give a precise description of the behaviour of a program. Here are some prominent examples.

##### 1.4.1 *Standards for implementation*

The formal semantics of a programming language is not meant as an alternative to the informal descriptions of programming constructs found in introductory programming textbooks. A formal semantics serves a very different purpose, namely to act as a yardstick that any implementation must conform to.

The examples mentioned in Section 1.2 all helped make computer scientists aware of the fact that *only a precise semantic definition can provide an exhaustive and implementation-independent account of all aspects of a programming language*. Such an account is particularly necessary if one is a ‘superuser’ of the language whose task is to implement an interpreter or a compiler or, in general, to create a language-dependent programming