

DEVELOPMENT OF PARALLEL PROGRAMS ON SHARED DATA-STRUCTURES REVISED VERSION

THIS REPORT IS A REVISED VERSION OF UMCS-91-1-1

Ketil Stølen
Oslo, March 2023

Preface

Thirty-two years ago, my PhD-thesis was published as Technical Report UMCS-91-1-1 at the computer science department at Manchester University. To save money, I was required to print the thesis on as few pages as possible. This meant small fonts (10 pt), no indentation, narrow page margins, etc., making the thesis unnecessarily hard to read. Over the years, I have also discovered several annoying misprints and minor mistakes. In this revised version these shortcomings have been removed.

Except for the corrections specified in detail on Pages iii-v, I have only changed the formatting of the original report.

Many thanks to Dag Frette Langmyhr, who helped me resolve some Latex-issues.

Corrections

The known mistakes in UMCS-91-1-1 and how they have been corrected in the current version.

- Page 32, Line +2: substitute ‘of’ for ‘of of’
- Page 32, Line +2: substitute ‘each finite’ for ‘each’
- Page 47, Line –10: remove ‘valid’
- Page 62, Section 11.2.1: substitute ‘ a_dcl ’ for ‘ $a-dcl$ ’
- Page 69, Line –10: substitute ‘ I ’ for ‘ $\overleftarrow{Inv} \Rightarrow Inv$ ’
- Page 75:
 - Line –9: substitute ‘ I ’ for ‘ $Dyn \wedge (\overleftarrow{Inv} \Rightarrow Inv)$ ’
 - Line –4: substitute ‘0’ for ‘ j ’
- Page 77, Line +7: substitute ‘ $\neg Empty(1)$ ’ for ‘ $\neg Empty(l)$ ’
- Page 81, Line +11: substitute ‘ $\overleftarrow{TrmS} \wedge \overleftarrow{Flag}$ ’ for ‘ $\overleftarrow{TrmS}$ ’
- Page 82:
 - Line +12: substitute ‘ $(L \cup \{Min\} = \overleftarrow{L} \vee L = \overleftarrow{L})$ ’ for ‘ $L \cup \{Min\} = \overleftarrow{L}$ ’
 - Line +16: substitute ‘ $(S \cup \{Max\} = \overleftarrow{S} \vee S = \overleftarrow{S})$ ’ for ‘ $S \cup \{Max\} = \overleftarrow{S}$ ’
- Page 83:
 - Line +9: same as on Page 82, Line +16
 - Line +14: same as on Page 82, Line +12
- Page 84:
 - Line –14: same as on Page 82, Line +12
 - Line –10: same as on Page 82, Line +16

- Page 86:
 - Line +12: same as on Page 82 Line +16
 - Line +17: same as on Page 82 Line +12
- Page 100, Line +8: substitute ‘,’ for ‘.’
- Page 117: substitute the two occurrences of ‘ $x_n := r_{x_n}$ ’ by ‘ $x_w := r_{x_w}$ ’
- Page 117: substitute the occurrence of ‘ $1 \leq k \leq n$ ’ by ‘ $1 \leq k \leq w$ ’
- Page 117: substitute the occurrence of ‘ $1 \leq j < k \leq n$ ’ in the footnote by ‘ $1 \leq j < k \leq w$ ’
- Page 118, Line +14: substitute ‘then $Rel[R, V, U]$ denotes the assertion’ by ‘then $Rel[R, V, U]$, $Rel_1[R, V, U]$ and $Rel_2[R, V, U]$ denote respectively the assertions:’
- Page 118, Line +15: substitute the full stop by comma and insert the following

$$\begin{aligned}
 & (R \wedge (h = 3 \curvearrowright \overline{h} \vee h \equiv 2 \curvearrowright \overline{h}) \wedge \\
 & \quad \bigwedge_{v \in V} h_v = [\underline{v}] \curvearrowright \overline{h_v} \wedge \bigwedge_{u \in U} h_u = [\underline{h_u(1)}] \curvearrowright \overline{h_u})^*, \\
 & (R \wedge (h = 3 \curvearrowright \overline{h} \vee h \equiv 1 \curvearrowright \overline{h}) \wedge \\
 & \quad \bigwedge_{v \in V} h_v = [\underline{v}] \curvearrowright \overline{h_v} \wedge \bigwedge_{u \in U} h_u = [\underline{h_u(1)}] \curvearrowright \overline{h_u})^*.
 \end{aligned}$$

- Page 118, Line +17: substitute ‘function Rel records any update due to the overall’ by ‘ Rel functions record any update due to the relevant’
- Page 118, Line +18: substitute ‘ Rel ’ by ‘the Rel functions’
- Page 118, Line –12: substitute the left-most occurrence of ‘ $::$ ’ by ‘sat’
- Page 119:
 - Line –15: substitute ‘ $Rel_1[\text{true}, \vartheta, U]$ ’ for ‘ $Rel[\text{true}, \vartheta, U]$ ’

- Line –14: substitute ‘ $Rel_2[\mathbf{true}, \vartheta, U]$ ’ for ‘ $Rel[\mathbf{true}, \vartheta, U]$ ’
- Page 121:
 - Line +4: substitute ‘ $Rel_2[\mathbf{true}, \vartheta', U]$ ’ for ‘ $Rel[\mathbf{true}, \vartheta', U]$ ’
- Page 125:
 - Line –8: substitute ‘ $Rel_1[\mathbf{true}, \vartheta, U]$ ’ for ‘ $Rel[\mathbf{true}, \vartheta, U]$ ’
 - Line –7: substitute ‘ $Rel_2[\mathbf{true}, \vartheta, U]$ ’ for ‘ $Rel[\mathbf{true}, \vartheta, U]$ ’
- Page 130:
 - Line –8: substitute ‘ $Rel_1[\mathbf{true}, \vartheta \cup \alpha, U]$ ’ for ‘ $Rel[\mathbf{true}, \vartheta \cup \alpha, U]$ ’
 - Line –7: substitute ‘ $Rel_2[\mathbf{true}, \vartheta \cup \alpha, U]$ ’ for ‘ $Rel[\mathbf{true}, \vartheta \cup \alpha, U]$ ’
- Page 149, [Bac59]: substitute ‘conference’ for ‘conferance’
- Page 151, [Gol90]: substitute ‘Conference’ for ‘Conferance’

DEVELOPMENT OF PARALLEL PROGRAMS ON SHARED DATA-STRUCTURES

THIS TECHNICAL REPORT IS A SLIGHTLY MODIFIED VERSION OF
[Stø90]

By
Ketil Stølen
December 1990

Contents

Abstract	12
Declaration	13
Education	14
Acknowledgements	15
1 Introduction and Summary	16
1.1 Motivation	16
1.1.1 Program Development	16
1.1.2 Concurrency	17
1.1.3 Compositionality	17
1.1.4 Rely- and Guar-Conditions	18
1.1.5 Temporal Logic	19
1.1.6 Compositional State Specifications	20
1.1.7 Without Temporal Logic	20
1.1.8 Other Systems	21
1.2 Our Approach	21
1.2.1 Properties	21
1.2.2 Lack of Generality	22
1.2.3 Operational Semantics	23
1.2.4 Auxiliary Variables	23
1.2.5 Specified Programs	24
1.2.6 Specifications	24
1.2.7 Convergent Environments	24
1.2.8 Divergence, Convergence and Deadlock	24
1.2.9 Satisfaction — Ignoring Auxiliary Variables	25

1.2.10	Satisfaction — General Case	25
1.2.11	Total Correctness	26
1.2.12	Logic of Specified Programs	26
1.2.13	Parallel-Rule	26
1.2.14	Soundness	27
1.2.15	Relative Completeness	27
1.2.16	Compositional Completeness	27
1.2.17	Adaptation Completeness	28
1.2.18	Achievements	28
1.2.19	Organisation of Thesis	29
2	First-Order Language	30
2.1	Syntax	30
2.1.1	Restrictions	30
2.1.2	Variables	30
2.2	Semantics	31
2.2.1	Structures	31
2.2.2	Valuations and States	32
2.2.3	Validity	33
2.2.4	Unary and Binary Assertions	33
2.2.5	Reflexivity	34
2.2.6	Transitivity	34
2.2.7	Respects	34
2.2.8	Substitution	34
3	Syntax of Programming Language	35
3.1	Motivation	35
3.2	Syntax	35
3.2.1	BNF	35
3.2.2	Additional Restrictions	36
3.2.3	Further Comments	39
3.2.4	Notation	39
4	Operational Semantics	40
4.1	Motivation	40
4.2	Computations	40
4.2.1	Configurations	40
4.2.2	Transitions	41

4.2.3	No Fairness	42
4.2.4	Naive Attempt	43
4.2.5	Progress	43
4.2.6	Hidden Variables	44
4.2.7	Summary	45
4.2.8	Why Local Variables Must be Initialised	46
4.2.9	Useful Notation	46
4.2.10	Meanings	47
4.2.11	Uninterrupted Transitions	47
4.3	Composition	47
4.3.1	Motivation	47
4.3.2	Summary	49
4.4	Decomposition	49
4.4.1	Motivation	49
4.4.2	Summary	50
4.5	Composition Proposition	50
4.6	Decomposition Proposition	51
5	Specifications	54
5.1	Motivation	54
5.2	Syntax	54
5.2.1	Tuple	54
5.2.2	Assertions	54
5.2.3	Variable Sets	55
5.3	Semantics	55
5.3.1	Assumptions	55
5.3.2	Commitments	56
5.4	Summary	57
5.5	An Example	58
6	Auxiliary Variables	60
6.1	Motivation	60
6.2	As a Specification Tool	61
6.2.1	In Isolation	61
6.2.2	Allowing Interference	61
6.2.3	Introducing Auxiliary Variables	62
6.3	As a Verification Tool	63
6.3.1	Encoding	63

6.3.2	Too Weak Specifications	64
6.3.3	Introducing Auxiliary Variables	65
6.3.4	Level of Detail	66
7	Specified Programs	67
7.1	Motivation	67
7.2	Syntax	67
7.2.1	Two Components	67
7.2.2	Summary	68
7.3	Satisfaction — Ignoring Auxiliary Variables	68
7.3.1	External	68
7.3.2	Internal	69
7.3.3	Summary	70
7.4	Satisfaction — General Case	70
7.4.1	Existence Check	70
7.4.2	Removals	71
7.4.3	Restrictions	71
7.4.4	Summary	72
7.4.5	Example	73
7.4.6	General Case	74
7.4.7	Auxiliary Form	74
7.4.8	Uniqueness Proposition	75
8	Syntactic Operators	76
8.1	Motivation	76
8.2	Sequence-Independent Operators	76
8.2.1	Reference	76
8.2.2	Identity	77
8.2.3	Set Quantification	77
8.2.4	Composition	78
8.2.5	Example	79
8.3	Finite-Sequence Operators	79
8.3.1	Transitive Closure	79
8.3.2	Example	80
8.3.3	Transitive and Reflexive Closure	80
8.3.4	Preservation	80

9	Well-Foundedness	82
9.1	Motivation	82
9.2	Unbounded Nondeterminism	82
9.3	Well-Foundedness Approach	84
10	Logic of Specified Programs	86
10.1	Motivation	86
10.1.1	Formal System	86
10.1.2	Decomposition-Rules	87
10.1.3	Proofs	88
10.1.4	Depth	88
10.1.5	Notation	88
10.1.6	Current Set of Rules	89
10.2	Consequence-Rule	89
10.3	Pre-Rule	89
10.4	Access-Rule	90
10.5	Skip-Rule	90
10.6	Assignment-Rule	90
10.6.1	Introduction	90
10.6.2	Without Auxiliary Variables	91
10.6.3	General Case	91
10.7	Block-Rule	92
10.8	Sequential-Rule	92
10.9	If-Rule	93
10.10	While-Rule	93
10.11	Parallel-Rule	93
10.11.1	Simple Case	93
10.11.2	General Case	94
10.11.3	Generalised Parallel-Rule	94
10.12	Await-Rule	95
10.12.1	Introduction	95
10.12.2	Without Auxiliary Variables	96
10.12.3	General Case	96
10.13	Elimination-Rule	96
10.14	Effect-Rule	97
10.15	Global-Rule	97
10.16	Auxiliary-Rule	97
10.17	Introduction-Rule	97

11 Simplifying Notation	99
11.1 Motivation	99
11.2 Operations	99
11.2.1 Main Structure	99
11.2.2 Variables	100
11.2.3 Observable and Hidden Changes	101
11.2.4 Splitting of Objects	102
11.3 Assertions Occurring in the Code	102
11.3.1 Sequential Case	102
11.3.2 Hooking	103
11.3.3 Interference	103
11.3.4 Two Interpretations	104
11.4 <i>Inv</i> and <i>Dyn</i>	106
12 Dining-Philosophers	108
12.1 Task	108
12.1.1 Mutual Exclusion	108
12.1.2 Workshop	108
12.2 Development	109
12.2.1 Data Structure	109
12.2.2 Main Program	110
12.2.3 Specification of Processes	111
12.2.4 Correctness Proof	111
12.2.5 Process Decomposition	112
12.2.6 Implementing Atomic Statements	114
13 Bubble-Lattice-Sort	115
13.1 Task	115
13.2 Development	116
13.2.1 Division into Processes	116
13.2.2 More Data Structure	116
13.2.3 Dynamic Invariant	118
13.2.4 Main Structure	118
13.2.5 Specification of Processes	119
13.2.6 Correctness Proof	120
13.2.7 Process Decomposition	121
13.2.8 Implementing Atomic Statements	123

14 Set-Partition	125
14.1 Task	125
14.2 Development	126
14.2.1 Algorithm	126
14.2.2 Data Structure	127
14.2.3 First Decomposition Step	127
14.2.4 Specifying <i>Init</i>	127
14.2.5 Final Implementation of <i>Init</i>	128
14.2.6 Invariant	129
14.2.7 Dynamic Invariant	129
14.2.8 Freedom from Deadlock	130
14.2.9 Specifying <i>Small</i>	130
14.2.10 Specifying <i>Large</i>	132
14.2.11 Parallel Composition	133
14.2.12 Decomposing <i>Small</i>	134
14.2.13 Final Implementation of <i>Sml</i>	135
14.2.14 Decomposing <i>Large</i>	137
14.2.15 Final Implementation of <i>Lrg</i>	139
15 Modified System	140
15.1 Motivation	140
15.2 Modifications	140
15.2.1 Specified Programs	140
15.2.2 Assumptions	141
15.2.3 Commitments	141
15.2.4 Satisfaction	141
15.2.5 Decomposition-Rules	142
15.2.6 Soundness and Completeness	142
15.3 Advantages	143
15.3.1 Always Enabled	143
15.3.2 Global Invariants	143
16 Dekker's Algorithm	144
16.1 Task	144
16.2 Development	145
16.2.1 Data Structure	145
16.2.2 Invariant	145
16.2.3 Main Structure	146

16.2.4	Guar-Condition	147
16.2.5	Wait-Condition	147
16.2.6	Specification	148
16.2.7	Composition Proof	149
16.2.8	Get Access	149
16.2.9	Release Access	152
17	Soundness	153
17.1	Parallel-Decomposition Proposition	153
17.2	Soundness Proposition	156
18	Relative Completeness	178
18.1	Motivation	178
18.1.1	Completeness	178
18.1.2	Incompleteness of LSP	178
18.1.3	Structure of Relative Completeness Proof	179
18.2	Construction of New Assertions	180
18.2.1	Decomposition	180
18.2.2	Closed Relations	181
18.2.3	Strongest Eff-Relation	182
18.2.4	Strongest Wait-Relation	182
18.2.5	Strongest Guar-Relation	183
18.2.6	Conversion Function	183
18.3	Historic Form	184
18.3.1	Enrichment of the Global State	184
18.3.2	History Variables	185
18.3.3	Component Functions	187
18.3.4	Historic-Form Function	191
18.3.5	Historic-Form Proposition	191
18.3.6	Decomposition Propositions	192
18.4	Expressiveness Proposition	198
18.5	Relative-Completeness Propositions	206
19	Discussion	217
19.1	Motivation	217
19.2	Without Hooked Variables	217
19.2.1	Two Representations	217
19.2.2	Unary Eff-Conditions	218

19.2.3	Rely and Guar as Assertion Sets	218
19.3	Separate Pre-Conditions	219
19.4	Scope of Eff-Conditions	219
19.4.1	Three Alternatives	219
19.4.2	Stirling	220
19.4.3	First Alternative Approach	221
19.4.4	Second Alternative Approach	222
19.4.5	Third Alternative Approach	222
19.5	Multi-State Assertions	223
19.5.1	Accessing the Initial State	223
19.5.2	Accessing the Final State	223
19.5.3	Where to Stop?	224
19.5.4	Conclusion	224
19.6	Atomicity	225
19.7	Boolean Tests of If and While	225
19.8	Fairness	228
19.8.1	Introduction	228
19.8.2	Two Alternative Systems	229
19.9	Data Reification	230
19.9.1	Introduction	230
19.9.2	Jones' System	230
19.9.3	LSP	230
19.10	Nondeterminacy	231
19.10.1	Introduction	231
19.10.2	Selection	231
19.10.3	Repetition	232
19.10.4	Decomposition-Rules	233
19.11	Partial Functions	233
19.11.1	Introduction	233
19.11.2	LPF	234
19.11.3	Changes to the Operational Semantics	234
19.11.4	Modified Rules	235
19.12	CSP	236
19.13	Without Reflexivity and Transitivity Constraints	237
19.13.1	Modifications	237
19.13.2	Alternative Rules	238
19.14	Allowing Environments to Diverge	239
19.15	Wait as an Assumption, Not as a Commitment	239

19.16 Proof-Obligations on Specifications	241
19.16.1 Implementability	241
19.17 Related Work	243
19.17.1 Owicki/Gries	243
19.17.2 Jones	243
19.17.3 Soundararajan	243
19.17.4 Barringer/Kuiper/Pnueli	243
19.17.5 Stirling	244
Bibliography	245

Abstract

A syntax-directed formal system for the development of totally correct programs with respect to an unfair shared-state parallel while-language is proposed. The system can be understood as a compositional reformulation of the Owicki/Gries method for verification of parallel programs.

Auxiliary variables are used both as a specification tool to eliminate undesirable implementations, and as a verification tool to make it possible to prove that an already finished program satisfies a particular specification.

Auxiliary variables may be of any sort, and it is up to the user to define the auxiliary structure he prefers. Moreover, the auxiliary structure is only a part of the logic. This means that auxiliary variables do not have to be implemented as if they were ordinary programming variables.

The system is proved sound and relatively complete with respect to an operational semantics and employed to develop three nontrivial algorithms: the Dining-Philosophers, the Bubble-Lattice-Sort and the Set-Partition algorithms.

Finally, a related method for the development of (possibly nonterminating) programs with respect to four properties is described. This approach is then used to develop Dekker's algorithm.

Declaration

No portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Education

Ketil Stølen received a Cand. Scient. degree in Computer Science from the University of Oslo in June 1986. He worked at the Norwegian Defense Research Establishment, Department of System Analysis from October 1986 until July 1987 (part of compulsory military service). He has been a post-graduate research student at Manchester University since September 1987.

Acknowledgements

First of all I would like to thank my supervisor, Professor Cliff B. Jones, for advice, encouragement and for carefully reading through and commenting on a large number of notes and drafts.

I am also indebted to Xu Qiwen and Wojciech Penczek, who have read and commented on several chapters, and to Professor Howard Barringer for discussions on temporal logic.

A special thanks goes to Jill Jones for correcting the English in an early draft of this thesis.

Finally, I would like to thank my examiners, Professor Mathai Joseph and Professor Howard Barringer, for their comments.

Financial support has been received from the Norwegian Research Council for Science and the Humanities.

Chapter 1

Introduction and Summary

1.1 Motivation

1.1.1 Program Development

In the 1960's it became more and more clear that the existing techniques for program development could not cope with the ever growing size and complexity of software products. Naur and Floyd proposed the use of logic to formally prove that a program had its desired effect [Nau66], [Flo67]¹.

Hoare reformulated this idea as a set of syntax-directed rules for a simple programming language [Hoa69]. This gave rise to what today is called Hoare-logic.

So far the emphasis was on verifying already finished programs. In the early 1970's the attention turned towards stepwise program development. The basic idea was that instead of verifying the correctness of a program first after it had been finished, such reasoning should be part of the development process and thereby simplify the argumentation and serve as a guide to the developer. Dijkstra advocated such views already in 1965 [Dij65]. Wirth, Dahl and Hoare had similar ideas (see [Wir71], [DH72], [Hoa72a]), and they were given a formal setting in [Mil71], [Hoa72b], [Jon72].

¹It is interesting to note that Turing proposed similar strategies already in 1949 [Tur49]. Naur and Floyd were unaware of this paper.

1.1.2 Concurrency

Attempts were soon made to employ similar techniques to concurrent programs. In Hoare-logic a triple of the form

$$\{P\} z \{Q\},$$

often called a Hoare-triple, is valid if for any initial state which satisfies the pre-condition P , the program z either ends in an infinite loop or terminates in a state which satisfies the post-condition Q .

Unfortunately, in the concurrent case such triples are insufficient since they do not contain information about the other processes running in parallel. This is why the Owicki/Gries [Owi75], [OG76a], [OG76b] extension of Hoare-logic to cover parallel programs on shared data structures and much of what followed depend upon an interference-freedom proof which first can be carried out when the whole program is complete. In these approaches programs are developed according to one set of rules (closely related to the sequential rules of Hoare-logic) down to the most concrete level, then the complete programs (with their proofs) are subject to a non-interference test.

For large software products this strategy is totally unacceptable, because erroneous design decisions, taken early in the design process, may remain undetected until the implementors attempt to provide the final non-interference proof. Since, in the worst case, everything that depends upon such mistakes will have to be thrown away, much work could be wasted.

The existence of non-interference proofs is not restricted to systems for shared-state parallel programs. Several of the early CSP [Hoa78] proof methods, like for example those described in [LG81] and [AFdR80], are based on related strategies.

1.1.3 Compositionality

To facilitate top-down development of programs, a proof method should satisfy the principle of compositionality (see [dR85], [Zwi89], [KKZ89]):

- A proof method is compositional if a program's specification can be verified on the basis of the specifications of its constituent components, without knowledge of the interior program structure of those components.

Hoare-logic is compositional, while the Naur/Floyd method [Flo67] is not; nor is any of the methods for concurrent programs discussed above.

Lamport proposed a system [Lam80], [LS84] suitable for proving partial correctness of both shared-state and communication-based parallel programs. Strictly speaking, the method is compositional. However, to prove that a program satisfies a specification, it is necessary to show that each of the component programs satisfies the very same specification. This means that program decomposition does not reduce the specification's complexity and, as pointed out in [Lam85], the method actually assumes a context of a complete program.

A more suitable compositional proof system for CSP was proposed by Misra and Chandy in [MC81]. Their proof tuples are of the form $r[h]s$, where h denotes a process, and r and s are assertions which express respectively what is assumed by the process h , and what h establishes under this assumption. The method was later extended in [MCS82] to deal with a weak form of liveness. In the latter approach, the proof tuple is augmented with a third assertion characterising when the process is guaranteed to progress.

A compositional proof system for CSP was also suggested in [ZH81].

1.1.4 Rely- and Guar-Conditions

Francez and Pnueli were the first to reason about shared-state parallel programs in terms of assumptions and commitments [Fra76], [FP78]. The basic idea is:

- If the environment, by which we mean the set of processes running in parallel with the one in question, fulfills the assumptions, then the actual process is required to fulfill the commitments.

Jones employed rely- and guar-conditions [Jon81], [Jon83a], [Jon83b] in a similar style. However, while the previous approach focused essentially on program verification; Jones main concern was top-down program development. His proof tuples are of the form

$$z \text{ \underline{sat} } (P, R, G, Q),$$

where z is a program and (P, R, G, Q) is a specification consisting of four assertions P , R , G and Q . The pre-condition P and the rely-condition R constitute assumptions that the developer can make about the environment.

In return the implementation z must satisfy the guar-condition G , the post-condition Q , and terminate, when operated in an environment which fulfills the assumptions.

The pre-condition is intended to characterise a set of states to which the implementation is applicable. Any uninterrupted state transition by the environment is supposed to satisfy the rely-condition, while any atomic state transition by the implementation must satisfy the guar-condition. The rely-condition is required to be both reflexive and transitive, while the guar-condition is only constrained to be reflexive². Finally, the post-condition is required to characterise the overall effect of using the implementation in such an environment.

This method allows erroneous interference decisions to be spotted and corrected at the level where they are taken. Thus programs can be developed in a top-down style.

Unfortunately, the system cannot deal with synchronisation. Hence, programs whose algorithms depend upon some sort of delay-construct cannot be developed.

In [Sti88] programs are specified in a similar style, although the rely- and guar-conditions are represented as sets of invariants, and the post-condition is unary, not binary as in [Jon81]. The paper reformulates and generalises Owicki's approach (in [OG76a]) as a system of syntax-directed rules for a while-language extended with await- and parallel-constructs. Although, this method favours top-down development in the sense of [Jon81], it can only be used for the design of partially correct programs.

In [Sou84] shared-variable concurrency is dealt with in a CSP inspired style. The method is compositional, but rather difficult to use in practice.

1.1.5 Temporal Logic

Many methods have been suggested for the development of totally correct programs in a similar way. In most cases they have been based on temporal logic.

The first truly compositional approach to (linear time) temporal logic was due to Barringer, Kuiper and Pnueli [BK84], [BKP84]. They offered an axiomatic semantics for a shared-state parallel language. The technique was later modified to deal with a CSP-like language [BKP85]. In both cases, they

²In [Jon83b] the guar-condition is also required to be transitive.

ended up with a very general specification language and logic.

1.1.6 Compositional State Specifications

As pointed out in [Sta88], an alternative to the use of powerful temporal operators is the technique of conceptual-state specification. In such a method, the behaviour of a process with respect to a collection of program variables is specified with the help of a collection of auxiliary variables, whose values serve as an abstract representation of the internal state of the process. Auxiliary variables appearing in a specification are not intended to be implemented.

Approaches in this tradition are described in [HO80], [OL82] and [Lam85]. Although these papers employ temporal operators to express liveness properties, the operators are weaker. Instead, auxiliary variables are employed to give the necessary expressive power.

1.1.7 Without Temporal Logic

Several authors have suggested the use of modal logic to prove termination of sequential programs. Burstall [Bur74] introduced the so-called intermittent assertions method, while Knuth applied related strategies on specific examples [Knu68].

However, despite the fact that the while-construct has a highly ‘temporal’ semantics, the standard techniques for proving termination of loops, the well-foundedness approach and the loop method (see [KM75]), do not depend upon temporal logic. One may therefore ask, can these methods be extended to deal with shared-state concurrency? In [LPS81] and [MP84] invariance and well-foundedness are used to prove temporal properties of parallel programs. This is further developed in [MP83] where temporal logic is applied together with invariance and well-foundedness. However, these methods are not compositional.

Stark [Sta85] has proposed a proof technique where a program can be proved to satisfy a rely/guarantee specification $R \supset G$, given that the same program satisfies a finite collection of rely/guarantee specifications $R_i \supset G_i$.

In [AS89] Buchi automata are employed to specify temporal properties. Proof obligations are derived from the automata and shown to satisfy an implementation by devising suitable invariant assertions and variant functions.

Both these methods are a lot more general than our approach. However,

neither of them can be said to provide the conceptual simplicity of Hoare-logic.

1.1.8 Other Systems

The historical overview given above is far from complete. We believe that the mentioned approaches are amongst the most important, but there are of course many other interesting and influential papers that could and perhaps should have been mentioned.

For example nothing has been said about action-based development methods like Unity (see [CM88] and [Gol90]), nor has Back's work been discussed [Bac88], [BS90]. Other promising ideas are described in [Len82], [SMSV83], [Mos86], [Bro89], [Ame89], [PJ87], [Pan90] and [Lam90].

In [Mid89] VDM is combined with temporal logic. However, this approach is difficult to evaluate due to lack of worked out examples.

1.2 Our Approach

1.2.1 Properties

This thesis proposes a method for top-down development of totally correct programs with respect to an unfair shared-state while-language extended with await- and parallel-constructs. The system can be interpreted as a compositional reformulation of the Owicki/Gries approach [OG76a].

Because the programming language is unfair, the method cannot deal with programs whose algorithms rely upon busy waiting³. For example, the parallel composition of the two programs

³It will later be indicated how similar strategies can be employed to develop totally correct programs with respect to fair programming languages (see page 228).

```

begin
  loc v;
  v := x;
  while v do
    v := x
  od
end

```

and

```

x := false

```

is not guaranteed to terminate, because the latter may be infinitely overtaken by the former.

Another deficiency with our approach is that it cannot be used to develop nonterminating programs⁴.

1.2.2 Lack of Generality

There are several methods that can be used to develop totally correct parallel programs with respect to an unfair programming language, and since in most cases they are a lot more general than the one described in this thesis, it may be argued that our system is of limited interest.

We do not agree with this. Although, it is quite possible to employ temporal logic in its most general form to develop totally correct sequential programs, most users would prefer to apply ordinary Hoare-logic in the style of for example Z [Spi88] or VDM [Jon90]. The reason is that these approaches are designed to deal with the sequential case only, and they are therefore both simpler to use and easier to understand than a formalism powerful enough to deal with concurrency.

The same can be said with respect to the development of terminating programs versus programs that are not supposed to terminate, and regarding different fairness constraints.

Our approach can be considered as an attempt to extend the method in [Jon81] to handle synchronisation, in the same way as [Jon81] extends VDM

⁴It will shown below that the method can be modified to develop (possibly nonterminating) programs with respect to four properties (see page 140).

to deal with interfering programs.

1.2.3 Operational Semantics

The programming language is given an operational semantics in the style of [Abr79], [Plo81] and [Acz83]. A (potential) computation is defined as a possibly infinite sequence of external and internal transitions on configurations. An internal transition represents a transition by the actual program, while an external transition is due to the program's environment. If a computation is finite, then no internal transition is enabled in its final configuration. The latter constraint ensures that a program will always progress given that the program stays enabled and is not infinitely overtaken by its environment. A configuration is blocked if it is disabled and its program component is different from the empty program.

1.2.4 Auxiliary Variables

Auxiliary variables will be employed for two different purposes:

- To strengthen a specification to eliminate undesirable implementations. In this case auxiliary variables are employed as a specification tool; they are used to characterise a program that has not yet been implemented.
- To strengthen a specification to make it possible to prove that an already finished program satisfies a particular specification. Here auxiliary variables are used as a verification tool, since they are introduced to show that a given algorithm satisfies a specific property.

In [Owi75] and [Sti88], where auxiliary variables are applied only as a verification tool, auxiliary variables are first implemented as if they were ordinary programming variables, and then afterwards removed by a deduction rule specially designed for this purpose.

This is not a very satisfactory method, because in some specifications a large number of auxiliary variables is needed, and the procedure of first implementing them and then removing them is rather tedious.

The method described in this thesis is more closely related to [HO80], [Sou84], [Hoa85], where the auxiliary structure is only a part of the logic and does not appear in the programs. Nevertheless, although it is possible to define trace-related variables in our system, auxiliary variables may be of any sort, and it is up to the user to define the auxiliary structure he prefers.

1.2.5 Specified Programs

The proof tuples of our system will be called specified programs and are of the form:

$$z \text{ \underline{sat} } \omega,$$

where z stands for a program and ω is a specification.

1.2.6 Specifications

A specification is a tuple of the form

$$(\vartheta, \alpha) :: (P, R, W, G, E)$$

consisting of two sets of variables ϑ and α , and five assertions P , R , W , G and E . The glo-set ϑ is the set of global variables, while the aux-set α is the set of auxiliary variables.

The pre-condition P and the rely-condition R describe the assumptions that can be made about the environment, while the wait-condition W , the guar-condition G and the eff-condition E constrain the implementation.

The pre-, rely- and guar-conditions are identical to the similarly named assertions in [Jon81]. The eff-condition corresponds to what [Jon81] calls the post-condition and covers interference both before the first internal transition and after the last. The wait-condition is supposed to characterise the set of states in which the implementation may become blocked. The implementation is not allowed to become blocked inside the body of an await-statement.

1.2.7 Convergent Environments

An environment is called convergent if it can only perform a finite number of consecutive atomic steps. This means that when a program is executed in a convergent environment, then no computation has infinitely many external transitions unless it also has infinitely many internal transitions.

1.2.8 Divergence, Convergence and Deadlock

Given a convergent environment, an execution of a program will be said to diverge if one of its processes ends in an infinite loop, or the body of one of

its await-statements becomes blocked. Moreover, it will be said to converge if it does not diverge, and to deadlock if it converges but does not terminate (in which case it is blocked in its final configuration).

1.2.9 Satisfaction — Ignoring Auxiliary Variables

A specified program,

$$z \text{ \underline{sat} } (\vartheta, \{\}) :: (P, R, W, G, E),$$

whose set of auxiliary variables is empty, is valid, if

- z converges,
- any atomic step due to the implementation satisfies G ,
- z can only become blocked in a state which satisfies W ,
- the overall effect is characterised by E if z terminates,

whenever

- the environment is convergent,
- z is called in a state which satisfies P ,
- any uninterrupted state transition by the environment satisfies R .

1.2.10 Satisfaction — General Case

In the general case,

$$z_1 \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E)$$

is valid, if there is a program z_2 such that

$$z_2 \text{ \underline{sat} } (\vartheta \cup \alpha, \{\}) :: (P, R, W, G, E)$$

is valid, and z_1 can be obtained from z_2 by removing all auxiliary structure with respect to α . The auxiliary structure is of course required to satisfy a number of constraints.

1.2.11 Total Correctness

This means that if

$$z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E)$$

is a valid specified program, and z is executed in a convergent environment characterised by P and R , then z either deadlocks in a state which satisfies W or terminates in a state such that the overall effect is characterised by E . Thus, z is totally correct with respect to the same specification if W is equivalent to false.

1.2.12 Logic of Specified Programs

The formal system, which is called LSP (Logic of Specified Programs), consists of sixteen decomposition-rules, and a base logic which is a first-order logic with a second order extension to allow for well-foundedness proofs. Only twelve decomposition-rules are needed for top-down development. The others have been introduced to make it easier to take advantage of already finished developments.

The basic system, consisting of only twelve rules, is called LSP_B .

1.2.13 Parallel-Rule

The premises of the parallel-rule must ensure that the two component programs are compatible with respect to mutual interference. It should also be clear that the parallel composition of two programs is only guaranteed to converge if called in an environment in which both component programs are guaranteed to converge. Thus, if the component programs do not deadlock, then the following rule is sufficient:

$$\frac{\begin{array}{l} z_1 \text{ sat } (\vartheta, \alpha) :: (P, R_1, \text{false}, G \wedge R_2, E_1) \\ z_2 \text{ sat } (\vartheta, \alpha) :: (P, R_2, \text{false}, G \wedge R_1, E_2) \end{array}}{\{z_1 \parallel z_2\} \text{ sat } (\vartheta, \alpha) :: (P, R_1 \wedge R_2, \text{false}, G, E_1 \wedge E_2)}$$

To formulate the general rule, it is enough to observe that z_1 is guaranteed to be released whenever it becomes blocked in a state in which z_2 cannot become blocked or terminate. Similarly, z_2 is guaranteed to be released in any state in which z_1 cannot become blocked or terminate. Thus:

$$\begin{array}{c}
\neg(W_1 \wedge E_2) \wedge \neg(W_2 \wedge E_1) \wedge \neg(W_1 \wedge W_2) \\
z_1 \text{ sat } (\vartheta, \alpha) :: (P, R_1, W \vee W_1, G \wedge R_2, E_1) \\
z_2 \text{ sat } (\vartheta, \alpha) :: (P, R_2, W \vee W_2, G \wedge R_1, E_2) \\
\hline
\{z_1 \parallel z_2\} \text{ sat } (\vartheta, \alpha) :: (P, R_1 \wedge R_2, W, G, E_1 \wedge E_2)
\end{array}$$

1.2.14 Soundness

LSP is proved sound with respect to the operational semantics; in other words:

- for any structure π and specified program ψ ,
 - if ψ is provable in LSP given the set of all base-logic assertions, valid in π , as axioms, then ψ is valid in π .

1.2.15 Relative Completeness

Under a number of constraints on the assertion language and the set of legal structures, the converse result is also shown:

- for any structure π and specified program ψ ,
 - if ψ is valid in π , then ψ is provable in LSP_B given the set of all base-logic assertions, valid in π , as axioms.

In other words, we can prove any valid specified program ψ with respect to a structure π , given an oracle capable of deciding if a base logic assertion is valid in π or not.

1.2.16 Compositional Completeness

The completeness result, together with the fact that our decomposition-rules are independent of the program component's internal structure, implies that LSP_B is compositionally complete [Zwi89]. Thus LSP favours top-down development.

1.2.17 Adaptation Completeness

Unfortunately, our method is less suited for bottom-up reasoning. It does not satisfy the following criterion, also called adaptation completeness [Zwi89]:

- if the specification of a program, whose internal structure is unknown, implies another specification for the same program, then the proof system admits a formal deduction of that fact.

For example, since

$$z \text{ sat } (\{v\}, \{\}) :: (\text{true}, v = \overline{v}, \text{false}, v \geq \overline{v}, v = \overline{v}),$$

is valid only if z leaves v unchanged, it should be possible to deduce

$$z \text{ sat } (\{v\}, \{\}) :: (\text{true}, v = \overline{v}, \text{false}, v = \overline{v}, v = \overline{v}).$$

However, there is no way to do this given the current set of rules.

1.2.18 Achievements

A compositional proof system for the deduction of totally correct shared-state parallel programs with respect to an unfair programming language is presented.

With respect to [OG76a], the main difference is that our system is compositional, that auxiliary variables can be employed as a specification tool, and that auxiliary variables are only a part of the logic and do not have to be implemented.

With respect to [Jon81], the main difference is that our system can deal with synchronisation, and that auxiliary variables can be used both as specification and verification tools.

With respect to [Sou84], the main difference is that our system can be used to prove total correctness and that auxiliary variables can be of any sort.

With respect to [BKP84], the main difference is that our system is conceptually simpler because it is designed to deal with one particular type of parallel programs.

With respect to [Sti88], the main difference is that our method can be used to prove total correctness, that auxiliary variables can be used as a

specification tool, and that auxiliary variables are only a part of the logic and do not have to be implemented.

1.2.19 Organisation of Thesis

The thesis is organised as follows. First the assertion language and its semantics are introduced. Secondly, the programming language is defined and given an operational semantics. The next two chapters deal respectively with specifications and auxiliary variables.

Then specified programs are defined both at the syntactic level and with respect to the operational semantics. Next, a number of simplifying syntactic operators are defined, well-foundedness is discussed, the decomposition-rules are introduced, and the method is applied on three different examples; the Bubble-Lattice-Sort, the Dining-Philosophers and the Set-Partition problems.

The approach is thereafter modified to allow development of possibly nonterminating programs with respect to four properties.

Then, LSP is proved sound, and it is shown that LSP_B is relatively complete.

Finally, alternative approaches and relationships with other systems are discussed, and possible extensions are indicated.

Chapter 2

First-Order Language

2.1 Syntax

2.1.1 Restrictions

Let L be a many-sorted first-order language with equality. It is assumed that $\mathbb{B}$ is a sort in L , and that L_P , the language of Peano arithmetic on the sort $\mathbb{N}$, is contained in L . Moreover, if Σ is a sort in L , then $\text{seq of } \Sigma$ is a sort in L , and

$$\begin{aligned} \text{len } _ &: \text{seq of } \Sigma \rightarrow \mathbb{N}, \\ _(-) &: \text{seq of } \Sigma \times \mathbb{N} \rightarrow \Sigma, \\ [] &: \rightarrow \text{seq of } \Sigma, \\ [-] &: \Sigma \rightarrow \text{seq of } \Sigma, \\ _ \curvearrowright _ &: \text{seq of } \Sigma \times \text{seq of } \Sigma \rightarrow \text{seq of } \Sigma \end{aligned}$$

are function symbols in L^1 .

The notation $e : \Sigma$ will be used to state that e is an expression in L of sort Σ .

2.1.2 Variables

We will follow [Jon90] in adding hooks to variables when it is necessary to refer to an earlier state (which is not necessarily the previous state). This

¹Alternatively, we could have introduced the notion of an acceptable structure (see [Mos74]).

means that, for any unhooked variable $v : \Sigma$, there is a hooked variable $\overleftarrow{v} : \Sigma$.

To avoid unnecessary complication, the same variable cannot be assigned to more than one sort. In other words, if $v_1 : \Sigma_1$ and $v_2 : \Sigma_2$, then $\Sigma_1 \neq \Sigma_2$ implies $v_1 \neq v_2$.

Finally, it is assumed that the same variable cannot be both quantified (bound) and unquantified (free) in the same expression. For example, this means that there is no expression in L of the form

$$x = y \wedge \forall x : x > 15.$$

To simplify expressions it is assumed that $\Rightarrow$ has higher priority than $\wedge$ and $\vee$, which again have higher priority than $\neg$, which has higher priority than $=$, which has higher priority than all other function symbols in L.

This means that we can write

$$x + y = 5 \wedge y = 2 \Rightarrow x = 3,$$

instead of

$$(((x + y) = 5) \wedge (y = 2)) \Rightarrow (x = 3).$$

Expressions of sort $\mathbb{B}$ will be called assertions.

2.2 Semantics

2.2.1 Structures

The next step is to assign meanings to expressions in L.

Definition 1 A structure π is a mapping of every sort Σ in L, to a nonempty set of values $\pi(\Sigma)$, and a mapping of every function symbol

$$f : \Sigma_1 \times \dots \times \Sigma_n \rightarrow \Sigma_{n+1}$$

in L, to a total function

$$\pi(f) : \pi(\Sigma_1) \times \dots \times \pi(\Sigma_n) \rightarrow \pi(\Sigma_{n+1}).$$

In particular, it is assumed that

- $\pi(\mathbb{B}) = \{false, true\}$,
- $\mathbb{N}$ is assigned the set of natural numbers,
- L_P is given its standard interpretation,
- for any sort, the equality operator $_ = _$ is given a standard interpretation,
- if Σ is assigned the set $\pi(\Sigma)$, then $\text{seq of } \Sigma$ is assigned the set of all finite sequences of the form

$$e_1 e_2 \dots e_n,$$

where for all $1 \leq j \leq n$, e_j is an element of $\pi(\Sigma)$,

- if $\pi(s)$ is a finite sequence, then the term $\text{len } s$ denotes the number of elements in $\pi(s)$,
- if $\pi(n)$ is a natural number greater than 0 and less than or equal to the number of elements in the sequence $\pi(t)$, then the term $t(n)$ denotes the $\pi(n)$ 'th element of $\pi(t)$,
- the term $[]$ denotes the empty sequence,
- if $\pi(t)$ is an element of $\pi(\Sigma)$, then the term $[t]$ denotes the sequence in $\pi(\text{seq of } \Sigma)$ consisting of one element $\pi(t)$,
- if $\pi(t_1)$ and $\pi(t_2)$ are elements of $\pi(\text{seq of } \Sigma)$, then the term $t_1 \frown t_2$ denotes the sequence in $\pi(\text{seq of } \Sigma)$ consisting of $\pi(t_1)$ concatenated with $\pi(t_2)$.

2.2.2 Valuations and States

A valuation $\mathcal{U}$ over a structure π , is a mapping of all variables in L to values in the structure. Obviously, if $v : \Sigma$ is a variable in L , then $\mathcal{U}(v) \in \pi(\Sigma)$. Similarly, a state is a mapping of all unhooked variables to values.

Given a set of variables V and two states s_1, s_2 , then

$$s_1 \stackrel{V}{=} s_2$$

denotes that for all variables $v \in V$, $s_1(v) = s_2(v)$, while

$$s_1 \stackrel{V}{\neq} s_2$$

means that there is a variable $v \in V$, such that $s_1(v) \neq s_2(v)$.

2.2.3 Validity

Given a structure π , and a valuation $\mathcal{U}$, then the expressions in L can be assigned meanings in the usual way. In the end, every expression $t : \Sigma$ in L will denote a value $\pi_{\mathcal{U}}(t) \in \pi(\Sigma)$. Given an assertion A , then

$$\models_{\pi} A,$$

if and only if, for all valuations $\mathcal{U}$ over π , $\pi_{\mathcal{U}}(A) = \text{true}$; in other words, if and only if A is valid in π . Moreover,

$$\models A,$$

if and only if A is valid in any structure. Similarly, if (s_1, s_2) is a pair of states, then

$$(s_1, s_2) \models_{\pi} A,$$

if and only if $\pi_{\mathcal{U}}(A) = \text{true}$, where for all unhooked variables v , $\mathcal{U}(\overleftarrow{v}) = s_1(v)$ and $\mathcal{U}(v) = s_2(v)$. The first state s_1 may be omitted if A has no occurrences of hooked variables.

2.2.4 Unary and Binary Assertions

An assertion A defines the set of all pairs of states (s_1, s_2) , such that

$$(s_1, s_2) \models_{\pi} A.$$

If A has no occurrences of hooked variables, it may also be thought of as the set of all states s , such that

$$s \models_{\pi} A.$$

We will use both interpretations below. This means that assertions without occurrences of hooked variables can be given two different interpretations. To indicate the intended meaning, we will distinguish between binary assertions and unary assertions. When an assertion is binary it denotes a set of pairs of states, and when an assertion is unary it denotes a set of states. In other words, an assertion with occurrences of hooked variables is always binary, while an assertion without occurrences of hooked variables can be both binary and unary.

2.2.5 Reflexivity

A binary assertion A will be said to be reflexive, if for all states s , $(s, s) \models_{\pi} A$.

2.2.6 Transitivity

A binary assertion A will be said to be transitive, if for all states s_1, s_2, s_3 , $(s_1, s_2) \models_{\pi} A$ and $(s_2, s_3) \models_{\pi} A$ implies $(s_1, s_3) \models_{\pi} A$.

2.2.7 Respects

A binary assertion A respects a set of variables V , if for all states s_1, s_2 , $(s_1, s_2) \models_{\pi} A$ implies $s_1 \stackrel{V}{=} s_2$.

2.2.8 Substitution

Given $m + 1$ expressions $t, r_1, \dots, r_m$ and m distinct variables $v_1, \dots, v_m$, then $t(r_1/v_1, \dots, r_m/v_m)$ denotes the result of replacing any occurrence of v_j ($1 \leq j \leq m$) in t with r_j .

Chapter 3

Syntax of Programming Language

3.1 Motivation

The programming language is closely related to the language used in [OG76a], where a traditional while language is extended with a parallel-statement and an await-construct. As pointed out in [OG76a], this is a flexible but primitive language, so primitive that other methods for synchronisation such as semaphores and events can be easily implemented using it. In other words, a formal system for this language can be employed to prove the correctness of programs using other tools as well. In our approach there is only one additional construct, namely a block-statement.

3.2 Syntax

3.2.1 BNF

A *program* is a finite, nonempty list of symbols whose context-independent syntax can be characterised in the well-known BNF-notation [Bac59]:

- Given that $\langle var \rangle$ denotes an unhooked variable in L , $\langle exp \rangle$ stands for an expression in L without hooks and quantifiers, and $\langle tst \rangle$ represents an expression in L of sort $\mathbb{B}$ without hooks and quantifiers,

then any program will be of the form $\langle pg \rangle$, where

$$\begin{aligned}
 \langle pg \rangle &::= \langle sk \rangle \mid \langle as \rangle \mid \langle bl \rangle \mid \langle sc \rangle \mid \langle if \rangle \mid \\
 &\quad \langle wd \rangle \mid \langle pr \rangle \mid \langle aw \rangle \\
 \langle sk \rangle &::= \text{skip} \\
 \langle as \rangle &::= \langle var \rangle := \langle exp \rangle \\
 \langle bl \rangle &::= \text{begin loc } \langle vl \rangle ; \langle pg \rangle \text{ end} \\
 \langle vl \rangle &::= \langle var \rangle \mid \langle var \rangle, \langle vl \rangle \\
 \langle sc \rangle &::= \langle pg \rangle ; \langle pg \rangle \\
 \langle if \rangle &::= \text{if } \langle tst \rangle \text{ then } \langle pg \rangle \text{ else } \langle pg \rangle \text{ fi} \\
 \langle wd \rangle &::= \text{while } \langle tst \rangle \text{ do } \langle pg \rangle \text{ od} \\
 \langle pr \rangle &::= \{ \langle pg \rangle \parallel \langle pg \rangle \} \\
 \langle aw \rangle &::= \text{await } \langle tst \rangle \text{ do } \langle pg \rangle \text{ od}
 \end{aligned}$$

3.2.2 Additional Restrictions

The main structure of a program is characterised above. However, a syntactically correct program is also required to satisfy four supplementary constraints, namely:

- Assignments:

For any assignment-statement ($\langle as \rangle$), the variable on the left-hand side is of the same sort as the expression on the right-hand side.

- Scope of Variables:

The block-statement ($\langle bl \rangle$) allows us to declare variables. A variable is local to a program, if it is declared in the program; otherwise it is said to be global. For example,

```

begin
  loc  $x, y$ ;
   $x := 5 + v$ ;
   $y := 9 - x$ 
end

```

declares two local variables, namely x and y , while v is a global variable. To avoid complications due to name clashes, it is required that

- the same variable cannot be declared more than once in the same program,
- a local variable cannot appear outside its block.

The first constraint avoids name clashes between local variables, while the second ensures that the set of global variables is disjoint from the set of local variables¹.

- Initialisation of Local Variables:

To avoid complicating the operational semantics, it is assumed that local variables are never read before they have been initialised. This means that

```
begin
  loc  $y$ ;
   $x := y$ 
end
```

is not a syntactically correct program, because y is not initialised before its value is (read and) assigned to x (see page 46).

- Boolean Tests :

To simplify the decomposition rules and the reasoning with auxiliary variables, it is assumed that for any program of the form $\{z_1 \parallel z_2\}$ or of the form $\{z_2 \parallel z_1\}$, and any program z_3 contained in z_1 :

- if z_3 is a while- ($< wd >$) or if-statement ($< if >$), then its Boolean test ($< tst >$) can only access variables declared in z_1 (in other words, variables local to z_1).

This constraint does not reduce the number of implementable algorithms, because if b is a Boolean test with occurrences of global variables, then

¹There are well-known techniques which allow these constraints to be relaxed. See for example [Ak89].

```

if  $b$  then
     $z_1$ 
else
     $z_2$ 
fi,

```

can be rewritten as

```

begin
    loc  $v$ ;
     $v := b$ ;
    if  $v$  then
         $z_1$ 
    else
         $z_2$ 
    fi
end,

```

while a program of the form

```

while  $b$  do
     $z$ 
od,

```

can be rewritten as

```

begin
    loc  $v$ ;
     $v := b$ ;
    while  $v$  do
         $z$ ;
         $v := b$ 
    od
end.

```

This constraint will later be discussed in more detail (see page 225).

3.2.3 Further Comments

Observe that there are no extra constraints on programs that occur in the bodies of await-statements ($\langle aw \rangle$). Thus, nested await-statements are legal, and parallel-statements may occur in the body of await-statements.

Moreover, there is no restriction with respect to the number of variable occurrences on the right hand side of an assignment-statement.

3.2.4 Notation

In the usual way, a program z_1 is a subprogram of a program z_2 , if z_1 occurs in z_2 . This means that any program is a subprogram of itself.

For any program z , $hid[z]$ is the least set consisting of any variable, which is declared in z (in other words, is a local variable in z) or occurs in the Boolean test of an if- or a while-statement in z .

For any expression or program t , $var[t]$ is the least set consisting of the unhooked version of any free hooked or unhooked variable in t .

Chapter 4

Operational Semantics

4.1 Motivation

The next step is to give a semantics for our programming language. We have decided to use an operational model (transition system) in the style of [HP79] but extended to potential computations. Usually, the meaning of a program in an operational semantics is characterised by the program's execution sequences. Unfortunately, when dealing with concurrency in a compositional way, this approach is not satisfactory.

The reason is that other programs running concurrently may interfere, and this has led several authors, see [Abr79], [Plo81], [Acz83], [BKP84] and [Sti88], to suggest that information about the environment should be included in the execution sequences. Our approach is in the latter tradition, and what we will call a computation is actually a potential computation or alternatively a potential execution sequence.

4.2 Computations

4.2.1 Configurations

A configuration is a pair $\langle z, s \rangle$, where z is either a program or the symbol ϵ which will be called the empty program, and s is a state.

4.2.2 Transitions

An *internal transition* represents a transition by the actual program, while a transition due to the program's environment is called an *external transition*. We will use two binary relations on configurations, $\xrightarrow{i}$ and $\xrightarrow{e}$, to characterise respectively the set of legal internal transitions and the set of legal external transitions ¹:

Definition 2 *Given a structure π , let $\xrightarrow{e}$ be the binary relation, such that*

- $\langle z, s_1 \rangle \xrightarrow{e} \langle z, s_2 \rangle$,

and let $\xrightarrow{i}$ be the least binary relation, such that either:

- $\langle \text{skip}, s \rangle \xrightarrow{i} \langle \epsilon, s \rangle$,
- $\langle v := r, s \rangle \xrightarrow{i} \langle \epsilon, s_r^v \rangle$, where s_r^v denotes the state that is obtained from s , by mapping the variable v to the value of the term r , determined by π and s , and leaving all other maplets unchanged,
- $\langle \text{begin loc } v_1, \dots, v_n; z \text{ end}, s \rangle \xrightarrow{i} \langle z, s \rangle$,
- $\langle z_1; z_2, s_1 \rangle \xrightarrow{i} \langle z_2, s_2 \rangle$ if $\langle z_1, s_1 \rangle \xrightarrow{i} \langle \epsilon, s_2 \rangle$,
- $\langle z_1; z_2, s_1 \rangle \xrightarrow{i} \langle z_3; z_2, s_2 \rangle$ if $\langle z_1, s_1 \rangle \xrightarrow{i} \langle z_3, s_2 \rangle$ and $z_3 \neq \epsilon$,
- $\langle \text{if } b \text{ then } z_1 \text{ else } z_2 \text{ fi}, s \rangle \xrightarrow{i} \langle z_1, s \rangle$ if $s \models_{\pi} b$,
- $\langle \text{if } b \text{ then } z_1 \text{ else } z_2 \text{ fi}, s \rangle \xrightarrow{i} \langle z_2, s \rangle$ if $s \models_{\pi} \neg b$,
- $\langle \text{while } b \text{ do } z_1 \text{ od}, s \rangle \xrightarrow{i} \langle z_1; \text{while } b \text{ do } z_1 \text{ od}, s \rangle$ if $s \models_{\pi} b$,
- $\langle \text{while } b \text{ do } z_1 \text{ od}, s \rangle \xrightarrow{i} \langle \epsilon, s \rangle$ if $s \models_{\pi} \neg b$,
- $\langle \{z_1 \parallel z_2\}, s_1 \rangle \xrightarrow{i} \langle z_2, s_2 \rangle$ if $\langle z_1, s_1 \rangle \xrightarrow{i} \langle \epsilon, s_2 \rangle$,
- $\langle \{z_1 \parallel z_2\}, s_1 \rangle \xrightarrow{i} \langle z_1, s_2 \rangle$ if $\langle z_2, s_1 \rangle \xrightarrow{i} \langle \epsilon, s_2 \rangle$,

¹It may be argued that this definition could have been simplified by allowing 'programs' of the form $\epsilon; z_1$, $\{\epsilon \parallel z_1\}$ and $\{z_1 \parallel \epsilon\}$. However, we find our approach 'cleaner', and also easier to work with in the definitions and proofs below.

- $\langle \{z_1 \parallel z_2\}, s_1 \rangle \xrightarrow{i} \langle \{z_3 \parallel z_2\}, s_2 \rangle$ if $\langle z_1, s_1 \rangle \xrightarrow{i} \langle z_3, s_2 \rangle$ and $z_3 \neq \epsilon$,
- $\langle \{z_1 \parallel z_2\}, s_1 \rangle \xrightarrow{i} \langle \{z_1 \parallel z_3\}, s_2 \rangle$ if $\langle z_2, s_1 \rangle \xrightarrow{i} \langle z_3, s_2 \rangle$ and $z_3 \neq \epsilon$,
- $\langle \text{await } b \text{ do } z_1 \text{ od}, s_1 \rangle \xrightarrow{i} \langle z_n, s_n \rangle$ if $s_1 \models_{\pi} b$, and
 - there is a list of configurations $\langle z_2, s_2 \rangle, \langle z_3, s_3 \rangle, \dots, \langle z_{n_1}, s_{n_1} \rangle$, such that for all $1 < k \leq n$, $\langle z_{k_1}, s_{k_1} \rangle \xrightarrow{i} \langle z_k, s_k \rangle$ and $z_n = \epsilon$,
- $\langle \text{await } b \text{ do } z_1 \text{ od}, s_1 \rangle \xrightarrow{i} \langle \text{await } b \text{ do } z_1 \text{ od}, s_1 \rangle$ if $s_1 \models_{\pi} b$, and
 - there is an infinite list of configurations $\langle z_2, s_2 \rangle, \langle z_3, s_3 \rangle, \dots, \langle z_n, s_n \rangle, \dots$, such that for all $k > 1$, $\langle z_{k_1}, s_{k_1} \rangle \xrightarrow{i} \langle z_k, s_k \rangle$, or
 - there is a finite list of configurations $\langle z_2, s_2 \rangle, \langle z_3, s_3 \rangle, \dots, \langle z_n, s_n \rangle$, where $z_n \neq \epsilon$, there is no configuration $\langle z_{n+1}, s_{n+1} \rangle$ such that $\langle z_n, s_n \rangle \xrightarrow{i} \langle z_{n+1}, s_{n+1} \rangle$, and for all $1 < k \leq n$, $\langle z_{k_1}, s_{k_1} \rangle \xrightarrow{i} \langle z_k, s_k \rangle$.

The assignment-statement and Boolean tests are here interpreted as atomic. We will later discuss this constraint in more detail (see page 225).

There are two different internal transitions defined for the await-statement. The first deals with the case when the await-statement's body terminates. The second models that the await-statement's body either ends in an infinite loop or becomes blocked. This topic will be discussed in greater detail below (see page 231). The different program constructs are all deterministic (although the parallel-construct has a nondeterministic behaviour). Furthermore, all functions are assumed to be total. We will later explain how the system can be extended to handle both nondeterministic statements (see page 231) and partial functions (see page 233).

4.2.3 No Fairness

No fairness constraint will be assumed. This means that a program can be infinitely overtaken by its environment. Thus, given that z_1 denotes the program


```

begin
  loc v;
  v := b;
  while  $\neg v$  do
    v := b
  od
end,

```

while z_2 represents

```

b := true,

```

then the program $\{z_1 \parallel z_2\}$ is not guaranteed to terminate because z_2 can be infinitely overtaken by z_1 . In other words, our system cannot be used to develop programs whose algorithms depend upon busy waiting as in this example.

4.2.4 Naive Attempt

As explained above, because there is no fairness constraint, a program can be infinitely overtaken by its environment. One might therefore think that it is sufficient to define a computation as a possibly infinite sequence of the form:

$$\langle z_1, s_1 \rangle \xrightarrow{a_1}_\pi \langle z_2, s_2 \rangle \xrightarrow{a_2}_\pi \dots \xrightarrow{a_{k_1}}_\pi \langle z_k, s_k \rangle \xrightarrow{a_k}_\pi \dots ,$$

where each arrow stands for an atomic transition, either internal or external, and for all $j \geq 1$, s_j is the state (with respect to the structure π) after j_1 transitions, while z_j either equals ϵ , in which case the actual program has terminated, or is a program describing what is left to be executed.

Unfortunately, this is too weak; for two reasons: we need both a progress property and a way to constrain the environment's access to hidden variables.

4.2.5 Progress

Nobody doubts that the sequential program

$x := 1$

(given its usual semantics) eventually will terminate. The reason is that any sensible sequential programming language satisfies the following progress property,

- if something can happen then eventually something will happen.

In the concurrent case, with respect to an unfair programming language, a slightly different progress property is required; namely that if the actual program is ‘enabled’ in the current configuration, then eventually a transition, either internal or external, will take place.

Otherwise, sequences of the form

$$\begin{aligned}
 & \langle \text{skip}, s \rangle \\
 & \langle \text{skip}, s_1 \rangle \xrightarrow{e} \langle \text{skip}, s_2 \rangle \xrightarrow{e} \langle \text{skip}, s_3 \rangle \xrightarrow{e} \langle \text{skip}, s_4 \rangle \\
 & \langle v_1 := t_1; v_2 := t_2, s_1 \rangle \xrightarrow{e} \langle v_1 := t_1; v_2 := t_2, s_2 \rangle \\
 & \quad \xrightarrow{i} \langle v_2 := t_2, s_3 \rangle
 \end{aligned}$$

are correct ‘computations’, and no program is totally correct, since this would allow programs to ‘stop executing’ without being disabled or infinitely overtaken by the environment.

To give a more accurate description of the progress property, and thereby define what we mean by a computation, we must first characterise when a configuration is enabled:

Definition 3 *A configuration c_1 is enabled if there is a configuration c_2 , such that $c_1 \xrightarrow{i} c_2$. If a configuration is not enabled, it will be said to be disabled. Finally, a configuration is blocked, if it is disabled, and its program component is different from the empty program.*

The progress property can then be formulated as follows: a computation is either infinite or the final configuration is disabled.

4.2.6 Hidden Variables

The second reason why the ‘definition’ of a computation given above is insufficient is that the environment must be constrained from updating hidden variables. For example, if z_1 denotes the program

```

    if  $v = 0$  then
      skip
    fi,

```

then there is no program of the form $\{z_1 \parallel z_2\}$, because v is not declared in z_1 (see page 37). To take advantage of this when formulating the decomposition-rules, another restriction on computations is needed. The following concept is useful:

Definition 4 *An external transition*

$$\langle z, s_1 \rangle \xrightarrow{e} \langle z, s_2 \rangle$$

respects a set of variables V , if and only if $s_1 \stackrel{V}{=} s_2$.

Then, if z denotes the program

```

begin
  loc  $y$ ;
   $y := v$ ;
  while  $x < 100$  do
     $x := x + y$ 
  od
end,

```

it is clear that $hid[z] = \{x, y\}$ (hid is defined on page 39). Since no program running in parallel with z can change the value of x or y , the required constraint with respect to this example is of course that any external transition in a computation of z must respect $hid[z]$.

4.2.7 Summary

A computation can now be defined as below:

Definition 5 A computation is a possibly infinite sequence of the form

$$\langle z_1, s_1 \rangle \xrightarrow{a_1}_{\pi} \langle z_2, s_2 \rangle \xrightarrow{a_2}_{\pi} \dots \xrightarrow{a_{k1}}_{\pi} \langle z_k, s_k \rangle \xrightarrow{a_k}_{\pi} \dots ,$$

such that

- $a_j \in \{e, i\}$ and $\langle z_j, s_j \rangle \xrightarrow{a_j}_\pi \langle z_{j+1}, s_{j+1} \rangle$,
- any external transition respects $hid[z_1]$,
- the sequence is either infinite or the final configuration is disabled.

4.2.8 Why Local Variables Must be Initialised

One of the constraints on programs is that local variables must be initialised before they are read. To see why, it is enough to observe that without this constraint and given the semantics above, the ‘program’

```

while  $v > 0$  do
  begin
    loc  $x$ ;
     $x := x_1$ ;
     $v := x$ 
  end
od

```

is guaranteed to terminate. The problem is that the operational semantics ‘remembers’ the value of x from the previous iteration. Thus if the initial value of x is m , the loop will terminate after m iterations.

One way to deal with this is to weaken the second constraint in definition 5, and only constrain the environment to leave local variables unchanged while they are ‘active’.

A second alternative is to define some sort of initialisation convention at semantic level. However, to avoid complicating the operational semantics, we are instead insisting that local variables must be initialised before they are read.

4.2.9 Useful Notation

Given a computation σ , then $\tau(\sigma)$, $\delta(\sigma)$ and $\lambda(\sigma)$ are the obvious projection functions to sequences of possibly empty programs, states and transition labels (actions), while $\tau(\sigma_j)$, $\delta(\sigma_j)$, $\lambda(\sigma_j)$ and σ_j denote respectively the j ’th

program, the j 'th state, the j 'th transition label and the j 'th configuration. Furthermore, $\sigma(j, \dots, k)$ represents

$$< \tau(\sigma_j), \delta(\sigma_j) > \xrightarrow{\lambda(\sigma_j)}_{\pi} \dots \xrightarrow{\lambda(\sigma_{k_1})}_{\pi} < \tau(\sigma_k), \delta(\sigma_k) >,$$

while $\sigma(j, \dots, \infty)$ stands for

$$< \tau(\sigma_j), \delta(\sigma_j) > \xrightarrow{\lambda(\sigma_j)}_{\pi} \dots \xrightarrow{\lambda(\sigma_{k_1})}_{\pi} < \tau(\sigma_k), \delta(\sigma_k) > \xrightarrow{\lambda(\sigma_k)}_{\pi} \dots .$$

If σ is infinite, then $len(\sigma) = \infty$, otherwise $len(\sigma)$ is equal to its number of configurations.

Finally, $\sigma_j \stackrel{V}{=} \sigma'_k$ means that $\tau(\sigma_j) = \tau(\sigma'_k)$ and $s(\sigma_j) \stackrel{V}{=} s(\sigma'_k)$, while $\sigma_j \stackrel{V}{\neq} \sigma'_k$ means that $\tau(\sigma_j) \neq \tau(\sigma'_k)$ or $s(\sigma_j) \stackrel{V}{\neq} s(\sigma'_k)$.

4.2.10 Meanings

As explained above, the meaning of a program is characterised by its set of computations. We will use a special notation to refer to this set:

Definition 6 *Given a structure π , and a program z , let $cp_{\pi}[z]$ be the set of all computations σ in π such that $\tau(\sigma_1) = z$.*

4.2.11 Uninterrupted Transitions

By an uninterrupted state transition by the environment we mean the overall effect of a finite number of external transitions not interrupted by any internal transition. Similarly, an uninterrupted state transition by the implementation is the overall effect of a finite number of internal transitions not interrupted by any external transition.

4.3 Composition

4.3.1 Motivation

A computation σ' of a program z_1 will be said to be compatible with a computation σ'' of a program z_2 , if they are of the same length, have identical

sequences of states, and are not performing internal transitions at the same time.

For example, let z_1 denote the program

$$v := v + t_1; z_3,$$

where z_3 is an atomic statement, and assume that z_2 represents

$$v := v + t_2$$

then

$$\langle z_1, s_1 \rangle \xrightarrow{i} \langle z_3, s_2 \rangle \xrightarrow{i} \langle \epsilon, s_3 \rangle$$

is a computation of z_1 which is not compatible with any computation of z_2 , because each finite computation of z_2 has exactly one internal transition and the computation above has no external transitions.

On the other hand, the computation

$$\langle z_1, s_1 \rangle \xrightarrow{i} \langle z_3, s_2 \rangle \xrightarrow{e} \langle z_3, s_3 \rangle \xrightarrow{e} \langle z_3, s_4 \rangle \xrightarrow{i} \langle \epsilon, s_5 \rangle$$

of z_1 , and

$$\langle z_2, s_1 \rangle \xrightarrow{e} \langle z_2, s_2 \rangle \xrightarrow{i} \langle \epsilon, s_3 \rangle \xrightarrow{e} \langle \epsilon, s_4 \rangle \xrightarrow{e} \langle \epsilon, s_5 \rangle$$

of z_2 , are compatible. Furthermore, they can be ‘composed’ into a computation

$$\begin{aligned} & \langle \{z_1 \parallel z_2\}, s_1 \rangle \xrightarrow{i} \langle \{z_3 \parallel z_2\}, s_2 \rangle \xrightarrow{i} \langle z_3, s_3 \rangle \\ & \xrightarrow{e} \langle z_3, s_4 \rangle \xrightarrow{i} \langle \epsilon, s_5 \rangle \end{aligned}$$

of $\{z_1 \parallel z_2\}$, by ‘composing’ the program part of each configuration pair, and making a transition internal if and only if one of the two component transitions are internal.

4.3.2 Summary

Definition 7 *Given a computation σ' of z_1 and a computation σ'' of z_2 , then σ' and σ'' are compatible, also written $\sigma' \bullet \sigma''$ if and only if*

- $len(\sigma') = len(\sigma'')$,
- $\delta(\sigma') = \delta(\sigma'')$,
- for all $1 \leq j \leq len(\sigma')$, $\lambda(\sigma'_j) = \lambda(\sigma''_j)$ implies $\lambda(\sigma'_j) = e$.

By joining together the compatible computations as described above, the set of computations representing the parallel composition of two programs can be generated from the two sets of computations characterising the subprograms.

4.4 Decomposition

4.4.1 Motivation

We have already shown how the execution sequences of two component programs can be used to determine the computations of their parallel composition. The next step is to indicate a method for decomposing computations.

Given that z_1 and z_2 denote the same programs as above, then

$$\begin{aligned} & \langle \{z_1 \parallel z_2\}, s_1 \rangle \xrightarrow{e} \langle \{z_1 \parallel z_2\}, s_2 \rangle \xrightarrow{i} \langle z_1, s_3 \rangle \xrightarrow{e} \langle z_1, s_4 \rangle \\ & \xrightarrow{e} \langle z_1, s_5 \rangle \xrightarrow{i} \langle z_3, s_6 \rangle \xrightarrow{i} \langle \epsilon, s_7 \rangle \end{aligned}$$

is a computation of $\{z_1 \parallel z_2\}$, while

$$\begin{aligned} & \langle z_1, s_1 \rangle \xrightarrow{e} \langle z_1, s_2 \rangle \xrightarrow{e} \langle z_1, s_3 \rangle \xrightarrow{e} \langle z_1, s_4 \rangle \xrightarrow{e} \langle z_1, s_5 \rangle \\ & \xrightarrow{i} \langle z_3, s_6 \rangle \xrightarrow{i} \langle \epsilon, s_7 \rangle \end{aligned}$$

is a computation of z_1 , and

$$\begin{aligned} & \langle z_2, s_1 \rangle \xrightarrow{e} \langle z_2, s_2 \rangle \xrightarrow{i} \langle \epsilon, s_3 \rangle \xrightarrow{e} \langle \epsilon, s_4 \rangle \xrightarrow{e} \langle \epsilon, s_5 \rangle \\ & \xrightarrow{e} \langle \epsilon, s_6 \rangle \xrightarrow{e} \langle \epsilon, s_7 \rangle \end{aligned}$$

is a computation of z_2 . In other words, a computation σ of $\{z_1 \parallel z_2\}$ can be decomposed into two computations of σ 's length and with σ 's sequence

of states. Each configuration in the computation of z_1 (respectively z_2) gets ‘what is left of’ of z_1 (respectively z_2) in the corresponding configuration of the initial computation. Moreover, an external transition is mapped into two external transitions, while an internal transition is split into one internal and one external transition.

4.4.2 Summary

Definition 8 *Given three computations, σ , σ' and σ'' , then $\sigma' \bullet \sigma'' \leftarrow \sigma$, if and only if*

- $\sigma' \bullet \sigma''$,
- $\text{len}(\sigma) = \text{len}(\sigma')$,
- $\delta(\sigma) = \delta(\sigma')$,
- for all $1 \leq j \leq \text{len}(\sigma)$,
 - $\tau(\sigma'_j) = \epsilon$ implies $\tau(\sigma_j) = \tau(\sigma''_j)$,
 - $\tau(\sigma''_j) = \epsilon$ implies $\tau(\sigma_j) = \tau(\sigma'_j)$,
 - $\tau(\sigma'_j) \neq \epsilon$ and $\tau(\sigma''_j) \neq \epsilon$ implies $\tau(\sigma_j) = \{\tau(\sigma'_j) \parallel \tau(\sigma''_j)\}$,
 - $\lambda(\sigma_j) = e$ if and only if $\lambda(\sigma'_j) = e$ and $\lambda(\sigma''_j) = e$.

4.5 Composition Proposition

Proposition 1 *Given a computation σ' of z_1 and a computation σ'' of z_2 , such that $\sigma' \bullet \sigma''$, then there is a computation σ of $\{z_1 \parallel z_2\}$ such that $\sigma' \bullet \sigma'' \leftarrow \sigma$.*

Proof: We will first give an algorithm for the construction of σ . Let $\tau(\sigma_1) = \{z_1 \parallel z_2\}$, and $\delta(\sigma_1) = \delta(\sigma'_1)$. Moreover, iteratively, in ascending order, for all $2 \leq k \leq \text{len}(\sigma')$:

- let $\delta(\sigma_k) = \delta(\sigma'_k)$;
- if $\lambda(\sigma'_{k_1}) = \lambda(\sigma''_{k_1}) = e$ then
 - let $\lambda(\sigma_{k_1}) = e$

- else
 - let $\lambda(\sigma_{k_1}) = i$;
- end if;
- if $\tau(\sigma'_k) = \epsilon$
 - let $\tau(\sigma_k) = \tau(\sigma''_k)$
- else if $\tau(\sigma''_k) = \epsilon$
 - let $\tau(\sigma_k) = \tau(\sigma'_k)$
- else
 - let $\tau(\sigma_k) = \{\tau(\sigma'_k) \parallel \tau(\sigma''_k)\}$
- end if;

It can easily be shown by induction that each transition step in σ is an element of either $\xrightarrow{i}$ or $\xrightarrow{\epsilon}$. Therefore, since

- $len(\sigma) = len(\sigma') = len(\sigma'')$,
- $len(\sigma') \neq \infty$ implies that $\sigma'_{len(\sigma')}$ and $\sigma''_{len(\sigma')}$ are both disabled, which again implies that $\sigma_{len(\sigma)}$ is disabled,
- any external transition in σ is also an external transition in both σ' and σ'' , and must therefore respect $hid[z_1] \cup hid[z_2]$,

it follows that σ is a computation.

end of proof

4.6 Decomposition Proposition

Proposition 2 *Given a program $\{z_1 \parallel z_2\}$ and one of its computations σ , then there are two computations σ' and σ'' of respectively z_1 and z_2 , such that*

$$\sigma' \bullet \sigma'' \leftarrow \sigma.$$

Proof: We will first give an algorithm for the construction of σ_1 and σ_2 . Let $\tau(\sigma'_1) = z_1$, $\tau(\sigma''_1) = z_2$, and $\delta(\sigma'_1) = \delta(\sigma''_1) = \delta(\sigma_1)$. Moreover, iteratively, in ascending order, for all $2 \leq k \leq \text{len}(\sigma)$:

- let $\delta(\sigma'_k) = \delta(\sigma''_k) = \delta(\sigma_k)$;
- if $\lambda(\sigma_{k_1}) = e$, let
 - $\tau(\sigma'_k) = \tau(\sigma'_{k_1})$; $\tau(\sigma''_k) = \tau(\sigma''_{k_1})$;
 - $\lambda(\sigma'_{k_1}) = e$; $\lambda(\sigma''_{k_1}) = e$;
- else if $\tau(\sigma'_{k_1}) = \epsilon$, let
 - $\tau(\sigma'_k) = \epsilon$; $\tau(\sigma''_k) = \tau(\sigma_k)$;
 - $\lambda(\sigma'_{k_1}) = e$; $\lambda(\sigma''_{k_1}) = i$;
- else if $\tau(\sigma''_{k_1}) = \epsilon$, let
 - $\tau(\sigma'_k) = \tau(\sigma_k)$; $\tau(\sigma''_k) = \epsilon$;
 - $\lambda(\sigma'_{k_1}) = i$; $\lambda(\sigma''_{k_1}) = e$;
- else if $\tau(\sigma_k) = \tau(\sigma''_{k_1})$, let
 - $\tau(\sigma'_k) = \epsilon$; $\tau(\sigma''_k) = \tau(\sigma''_{k_1})$;
 - $\lambda(\sigma'_{k_1}) = i$; $\lambda(\sigma''_{k_1}) = e$;
- else if $\tau(\sigma_k) = \tau(\sigma'_{k_1})$, let
 - $\tau(\sigma'_k) = \tau(\sigma'_{k_1})$; $\tau(\sigma''_k) = \epsilon$;
 - $\lambda(\sigma'_{k_1}) = e$; $\lambda(\sigma''_{k_1}) = i$;
- else if $\tau(\sigma_k)$ is of the form $\{\tau \parallel \tau(\sigma''_{k_1})\}$ and $\langle \tau(\sigma'_{k_1}), s_{k_1} \rangle \xrightarrow{i} \langle \tau, s_k \rangle$, let
 - $\tau(\sigma'_k) = \tau$; $\tau(\sigma''_k) = \tau(\sigma''_{k_1})$;
 - $\lambda(\sigma'_{k_1}) = i$; $\lambda(\sigma''_{k_1}) = e$;
- else if $\tau(\sigma_k)$ is of the form $\{\tau(\sigma'_{k_1}) \parallel \tau\}$ and $\langle \tau(\sigma''_{k_1}), s_{k_1} \rangle \xrightarrow{i} \langle \tau, s_k \rangle$, let

- $\tau(\sigma'_k) = \tau(\sigma'_{k_1}); \tau(\sigma''_k) = \tau;$
- $\lambda(\sigma'_{k_1}) = e; \lambda(\sigma''_{k_1}) = i;$

- end if;

It can easily be shown by induction that each transition step in both σ' and σ'' is an element of either $\xrightarrow{i}$ or $\xrightarrow{e}$. Moreover, since

- $len(\sigma') = len(\sigma'') = len(\sigma),$
- $len(\sigma) \neq \infty$ and $\sigma_{len(\sigma)}$ is disabled implies that both $\sigma'_{len(\sigma')}$ and $\sigma''_{len(\sigma')}$ are disabled,
- any external transition in σ respects $hid[\{z_1 \parallel z_2\}],$

it follows that both σ' and σ'' are computations.

end of proof

Chapter 5

Specifications

5.1 Motivation

The object of this chapter is first of all to characterise what we mean by a specification at syntactic level, and secondly to indicate its intended semantics.

5.2 Syntax

5.2.1 Tuple

A specification is a tuple of the form

$$(\vartheta, \alpha) :: (P, R, W, G, E).$$

The pre-condition P , the rely-condition R and the guar-condition G are identical to the similarly named assertions in [Jon81], while the eff-condition E corresponds to what [Jon81] calls the post-condition. Moreover, W will be called the wait-condition, while ϑ and α are finite sets of variables, called respectively the glo-set and the aux-set.

5.2.2 Assertions

The pre- and wait-conditions are unary assertions. In other words, they are without occurrences of hooked variables. The rely-, guar- and eff-conditions

are binary assertions and may therefore refer to both hooked and unhooked variables, in which case the hooked variables refer to the ‘older’ state.

5.2.3 Variable Sets

The glo-set is assumed to contain the global implementable variables, while the elements of the aux-set will be called auxiliary variables. The auxiliary variables are not implementable (with respect to the actual specification).

Since the elements of the two sets belong to the same syntactical category, the sets are required to be disjoint. Moreover, to make sure that the assertions are only constraining variables from the two sets, the unhooked version of any free variable that appears in the specification is required to be an element of one of the two sets.

5.3 Semantics

5.3.1 Assumptions

A specification states a number of assumptions about the environment.

- First of all, it is assumed that the environment can only perform a finite number of consecutive atomic steps. Such an environment is said to be convergent. Thus, when a program is executed in a convergent environment, then no computation has infinitely many external transitions unless it also has infinitely many internal transitions.

This is not a fairness requirement on the programming language, because it does not constrain the implementation of a specification. If for example a parallel-statement $\{z_1 \parallel z_2\}$ occurs in the implementation then this assumption does not influence whether or not z_1 is infinitely overtaken by z_2 .

Moreover, the assumption can be removed (it is explained how on page 239). The only disadvantage would be that we no longer can use our formal system to prove that a program terminates, but only that a program terminates when it is not infinitely overtaken by the environment.

- Another more important assumption is that the initial state satisfies the pre-condition.

- Furthermore, it is also assumed that any external transition satisfies the rely-condition. For example, given the rely-condition

$$x \leq \overleftarrow{x} \wedge y = \overleftarrow{y},$$

then it is assumed that the environment will never change the value of y . Moreover, if the environment assigns a new value to x , then this value will be less than or equal to the variable's previous value.

To make it possible for the implementor to assume that any uninterrupted interference between two internal transitions satisfies the rely-condition, the rely-condition is required to be reflexive because the environment may not interfere at all, and transitive to cover the case when the environment interferes more than once. This constraint has also a simplifying effect on some of the decomposition-rules. (It is shown on page 237 how the reflexivity and transitivity requirements can be removed.)

5.3.2 Commitments

A specification is of course not only stating assumptions about the environment, but also commitments to the implementation. To simplify their formulation we will introduce some new notation: Given a convergent environment, a computation σ will be said to terminate if $\text{len}(\sigma) \neq \infty$ and $\tau(\sigma_{\text{len}(\sigma)}) = \epsilon$, to deadlock if $\text{len}(\sigma) \neq \infty$ and $\tau(\sigma_{\text{len}(\sigma)}) \neq \epsilon$, to diverge if $\text{len}(\sigma) = \infty$, and to converge if σ does not diverge.

- Given an environment which satisfies the assumptions, then an implementation is required to converge.
- Moreover, the wait-condition is supposed to describe the set of states in which the implementation may become blocked. The implementation is not allowed to become blocked inside the body of an await-statement. Since a computation which deadlocks is finite (given a convergent environment), it is enough to insist that the final state of any deadlocking computation satisfies the wait-condition.
- Any internal transition is required to satisfy the guar-condition. The guar-condition is constrained to be reflexive because this has a simpli-

fying effect on some of the decomposition-rules (it is explained on page 237 how this constraint can be removed). However, the guar-condition is not required to be transitive. The reason is that it from time to time is necessary to require a certain maximum granularity. For example, if we want to specify a process which simulates a certain behaviour in a particular environment. (*Phil(l)* on page 111 is one example.)

As mentioned above, the rely-condition is assumed to characterise the overall effect of any (finite) uninterrupted sequence of external transitions. The guar-condition is closely related to the rely-condition since it determines what the environment can rely on with respect to the implementation: any (finite) uninterrupted sequence of internal transitions is characterised by the transitive closure of the guar-condition; in other words, by the strongest transitive assertions implied by the guar-condition.

- Finally, the eff-condition is intended to characterise the overall effect when the implementation terminates. External transitions both before the first internal transition and after the last are included. This means that given the rely-condition $x \geq \overline{x}$, the strongest eff-condition for the program `skip` is $x \geq \overline{x}$.

5.4 Summary

The constraints on a specification are restated in a more formal notation below:

Definition 9 *A specification is a tuple of the form*

$$(\vartheta, \alpha) :: (P, R, W, G, E),$$

where R , G and E are binary assertions, P and W are unary assertions, ϑ and α are finite, disjoint sets of variables, such that:

- R is reflexive and transitive,
- G is reflexive,
- the unhooked version of any free variable occurring in (P, R, W, G, E) is an element of $\vartheta \cup \alpha$.

5.5 An Example

Assume we want to specify a program that reads 10 natural numbers from a variable *In* and adds them to an initially empty bag *Rcv*. Only one number can be stored in *In* at a time, and it is the environment's job to provide a new number when the old one has been read.

There is no other restriction on when the program may be applied given that the constraint imposed by the rely-condition is satisfied, so it is enough if the pre-condition constrains *Rcv* to be empty.

Moreover, the environment is assumed not to change the value of *Rcv*. In other words, *Rcv* will have 10 elements when (and if) the program terminates. The eff-condition can be used to express this requirement.

To inform the environment that a number has been read, the program switches on the flag *Flag*. The environment signals that a new number has been loaded into *In* by changing *Flag* to false.

It is assumed that the program updates the variables *Flag* and *Rcv* in the same atomic step. Thus the guar-condition must express that the program either leaves the state as it is or changes *Flag* from false to true and adds *In* to *Rcv* without changing the value of *In*.

Furthermore, it is clear that the program will only wait if *Flag* is true and the size of *Rcv* is less than 10, which gives us the wait-condition.

In what way values are loaded into *In* is unspecified. For example, the environment may change the value of *Flag* any number of times between each addition of a new value to *Rcv*. Thus, the only constraint on the environment is that it cannot change the value of *Rcv*.

Therefore, given that $\#Rcv$ denotes the number of elements in *Rcv*, we end up with the following specification:

glo $\{Rcv, Flag, In\}$
 aux $\{\}$
 pre $Rcv = \{\}$
 rely $Rcv = \overline{Rcv}$
 wait $Flag \wedge \#Rcv < 10$
 guar $In = \overline{In} \wedge$
 $((\overline{Flag} = Flag \wedge Rcv = \overline{Rcv}) \vee$
 $(\neg \overline{Flag} \wedge Flag \wedge Rcv = \overline{Rcv} \cup \{In\}))$
 eff $\#Rcv = 10.$

Chapter 6

Auxiliary Variables

6.1 Motivation

Auxiliary variables were early recognised as helpful tools to reason about concurrency, see [BH73], [Lau73] or [Cli73]. Since then they have occurred in many shapes and forms. For example, in CSP related systems, [Hoa85], [Sou84] or [PJ87], auxiliary variables appear in the logic as nonimplementable traces. Another related concept is what [Lam83] calls state functions.

In [Owi75] and [Sti88], auxiliary variables are first implemented as if they were ordinary programming variables, and then afterwards removed by a deduction rule specially designed for this purpose. This is not a very satisfactory method, because in some specifications a large number of auxiliary variables are needed, and the procedure of first implementing them and then removing them is rather tedious.

Our approach is more in the style of [Hoa85], where the auxiliary structure is only a part of the logic and does not appear in the programs. Nevertheless, although it is possible to define trace-related variables in our system, auxiliary variables may be of any sort, and it is up to the user to define the auxiliary structure he prefers.

We will use auxiliary variables for two different purposes:

- To strengthen a specification to eliminate undesirable implementations. In this case auxiliary variables are used as a specification tool; they are employed to characterise a program that has not yet been implemented.
- To strengthen a specification to make it possible to prove that an already finished program satisfies a particular specification. Here auxil-

iary variables are used as a verification tool, since they are introduced to show that a given algorithm satisfies a specific property.

6.2 As a Specification Tool

6.2.1 In Isolation

We will first discuss an example where it is necessary to use auxiliary variables as a specification tool. Assume we want to specify a program that adds a new element to a global buffer called *Buff*. If we are satisfied with an implementation that can only run in isolation, then this can easily be expressed as follows:

glo	$\{Buff\}$
aux	$\{\}$
pre	true
rely	$Buff = \overline{Buff}$
wait	false
guar	$Buff = \overline{Buff} \vee Buff = [A] \curvearrowright \overline{Buff}$
eff	$Buff = [A] \curvearrowright \overline{Buff}$

The pre-condition allows the implementation to be applied in any state. Moreover, the rely-condition restricts the environment from changing the value of *Buff*, thus we may use the eff-condition to express the desired property. Finally, the guar-condition specifies that the concatenation step takes place in isolation, while the falsity of the wait-condition requires the implementation to terminate.

6.2.2 Allowing Interference

However, if the environment is allowed to make any changes to *Buff*, the task of formulating a specification becomes more difficult. Observe that we still consider the actual concatenation step to be atomic; the only difference from above is that the environment may now interfere immediately before and (or) after the concatenation takes place.

Thus, since there are no restrictions on the ways the environment can change *Buff*, and because external transitions, both before the first inter-

nal transition and after the last, are included in the eff-condition, the eff-condition must allow anything to happen.

This means that the eff-condition is no longer of much use. Thus, we are left with the guar-condition as our only hope to pin down the intended meaning. The specification

glo	$\{Buff\}$
aux	$\{\}$
pre	true
rely	true
wait	false
guar	$Buff = \overline{Buff} \vee Buff = [A] \curvearrowright \overline{Buff}$
eff	true

is almost sufficient. The only problem is that there is no restriction on the number of times the operation is allowed to add A to $Buff$. Hence, for example, the statement

skip

is one possible implementation, while

$$\begin{aligned} Buff &:= [A] \curvearrowright Buff; \\ Buff &:= [A] \curvearrowright Buff \end{aligned}$$

is another correct implementation.

6.2.3 Introducing Auxiliary Variables

One solution is to introduce a Boolean auxiliary variable called *Done*, and use *Done* as a flag to indicate whether A has been added to $Buff$ or not. Then the program can be specified as follows:

glo	$\{Buff\}$
aux	$\{Done\}$
pre	$\neg Done$
rely	$Done = \overline{Done}$
wait	false
guar	$(Buff = \overline{Buff} \wedge Done = \overline{Done}) \vee$ $(Buff = [A] \curvearrowright \overline{Buff} \wedge \neg Done \wedge Done)$
eff	$Done$

Since

- the environment cannot change the value of $Done$,
- A can only be added to $Buff$ in a state where $Done$ is false,
- the concatenation transition changes $Done$ from false to true,
- the operation is not allowed to change $Done$ from true to false,

the pre- and eff-conditions imply that A is added to $Buff$ once and only once.

6.3 As a Verification Tool

6.3.1 Encoding

We will now try to explain what we mean by using auxiliary variables as a verification tool. At any point during the execution of a program, the set of possible internal transitions is a function of the global state, the local state and the program counter.

In the tradition of [Jon81] specifications will only constrain the global state, and it is up to the implementor to choose the local data structure. Thus, to add auxiliary structure and use this to encode information about the program counter and the local state into the global state, will in many cases be of crucial importance.

6.3.2 Too Weak Specifications

For example, without auxiliary structure, the strongest possible guar-condition for the program

$$\begin{array}{l} v := v + 1; \\ v := v + 2, \end{array}$$

is

$$v = \underline{v} \vee v = \underline{v} + 1 \vee v = \underline{v} + 2.$$

Moreover,

glo	$\{v\}$
aux	$\{\}$
pre	true
rely	$v \geq \underline{v}$
wait	false
guar	$v = \underline{v} \vee v = \underline{v} + 1 \vee v = \underline{v} + 2$
eff	$v \geq \underline{v} + 3$

is the specification that gives the best possible characterisation given the actual assumptions about the environment. Furthermore, this tuple is also the strongest possible specification of the program

$$\begin{array}{l} v := v + 2; \\ v := v + 1 \end{array}$$

with respect to the same constraints on the environment.

Let z_1 denote the first program and z_2 the second. If we constrain the overall environment to leave v unchanged, it is clear that the parallel composition of z_1 and z_2 satisfies:

glo	$\{v\}$
aux	$\{\}$
pre	true
rely	$v = \underline{v}$
wait	false
guar	true
eff	$v = \underline{v} + 6.$

Unfortunately, there is no way to deduce this from the information in the two component specifications. The problem is of course that the rely-conditions do not give enough information. Thus, unless the expressibility is increased, any compositional system based on our specifications will be hopelessly incomplete.

6.3.3 Introducing Auxiliary Variables

One way to increase the expressibility is to introduce auxiliary variables. Let p_1 be a variable that records the overall effect of the updates to v in z_1 , while p_2 characterises the overall change to v due to z_2 . Clearly, it is required that z_2 has no effect on p_1 , and also that z_1 cannot change the value of p_2 .

The specification of z_1 can then be rewritten as

glo	$\{v\}$
aux	$\{p_1, p_2\}$
pre	true
rely	$p_1 = \underline{p_1} \wedge$ $v - \underline{v} = p_2 - \underline{p_2}$
wait	false
guar	$p_2 = \underline{p_2} \wedge$ $v - \underline{v} = p_1 - \underline{p_1}$
eff	$v = \underline{v} + 3 + (p_2 - \underline{p_2}) \wedge p_1 - \underline{p_1} = 3,$

while z_2 fulfills

glo	$\{v\}$
aux	$\{p_1, p_2\}$
pre	true
rely	$p_2 = \underline{p_2} \wedge$ $v - \underline{v} = p_1 - \underline{p_1}$
wait	false
guar	$p_1 = \underline{p_1} \wedge$ $v - \underline{v} = p_2 - \underline{p_2}$
eff	$v = \underline{v} + (p_1 - \underline{p_1}) + 3 \wedge p_2 - \underline{p_2} = 3.$

It follows easily from these two specifications that, if the overall environment is restricted from changing v , then the overall effect of the parallel composition of z_1 and z_2 is characterised by $v = \underline{v} + 6$.

In other words, we have given an informal description of how auxiliary structure can be introduced to prove that an already finished implementation satisfies a given specification.

6.3.4 Level of Detail

Not surprisingly, auxiliary variables can be used to specify the global effect of a program to any level of detail. For example, the specification of z_1 can easily be strengthened to

glo	$\{v\}$
aux	$\{p_1, p_2\}$
pre	$p_1 = 0$
rely	$p_1 = \underline{p_1} \wedge$ $v - \underline{v} = p_2 - \underline{p_2}$
wait	false
guar	$p_2 = \underline{p_2} \wedge$ $((p_1 = \underline{p_1} \wedge v = \underline{v}) \vee$ $(\underline{p_1} = 0 \wedge p_1 = 1 \wedge v = \underline{v} + 1) \vee$ $(\underline{p_1} = 1 \wedge p_1 = 3 \wedge v = \underline{v} + 2))$
eff	$v = \underline{v} + 3 + (p_2 - \underline{p_2}) \wedge p_1 = 3,$

which not only specifies exactly how the state can be altered by z_1 , but also gives the order of the updates.

Chapter 7

Specified Programs

7.1 Motivation

A formal system, whose well-formed formulas are called *specified programs*, will be defined below. The aim of this chapter is to characterise the latter concept; both at syntactic level, and with respect to the operational semantics.

Satisfaction will be defined in two steps; first for the situation when the set of auxiliary variables is empty. This definition is then extended to characterise satisfaction in the general case.

7.2 Syntax

7.2.1 Two Components

Since the object of our formal system is to prove that a program satisfies a specification, a specified program is a tuple of the form

$$z \text{ sat } \omega,$$

where z is a program and ω is a specification. To avoid unnecessary complications when formulating the block-rule, it is required that none of the program's local variables occurs in the specification. To ensure that only programming variables are implemented, any global variable occurring in the program must be included in the specification's glo-set.

7.2.2 Summary

The constraints are summed up in the definition below:

Definition 10 *A specified program is a tuple of the form*

$$z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E),$$

where z is a program and $(\vartheta, \alpha) :: (P, R, W, G, E)$ is a specification, such that

- no local variable of z occurs in $\alpha \cup \vartheta$,
- any global variable of z is an element of ϑ .

The set of all specified programs will be denoted SP .

7.3 Satisfaction — Ignoring Auxiliary Variables

7.3.1 External

What does it mean for a program to satisfy a specification when the set of auxiliary variables is empty? First of all, as indicated in the informal description, a specification invites the implementor to assume a number of things about the environment. One assumption is that an implementation is called only in a state which satisfies the pre-condition. Moreover, any uninterrupted state transition by the environment is supposed to satisfy the rely-condition, and it is also assumed that the environment is convergent¹: Thus, we may restrict our attention to computations characterised by $ext_\pi[P, R]$:

Definition 11 *Given a pre-condition P , a rely-condition R , and a structure π , then $ext_\pi[P, R]$ denotes the set of all computations σ in π , such that:*

- $\delta(\sigma_1) \models_\pi P$,
- for all $1 \leq j < \text{len}(\sigma)$, if $\lambda(\sigma_j) = e$ then $(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_\pi R$,
- if $\text{len}(\sigma) = \infty$, then for all $j \geq 1$, there is a $k \geq j$, such that $\lambda(\sigma_k) = i$.

¹This assumption can be removed. See page 239.

The need for the first constraint should be obvious. The second guarantees that any external transition satisfies the rely-condition. Since the rely-condition is transitive, this means that the rely-condition is satisfied by any uninterrupted state transition by the environment. Finally, the third condition ensures that the environment is convergent.

Clearly, for any program z and state s which satisfies the pre-condition, there are infinitely many computations

$$\sigma \in \text{ext}_\pi[P, R] \cap \text{cp}_\pi[z]$$

whose initial configuration is $\langle z, s \rangle$ (cp_π is defined on page 47).

7.3.2 Internal

We have already described the assumptions the implementor can make about the environment. Moreover, for a given structure π , the set of all possible computations of a program z , with respect to a convergent environment characterised by a pre-condition P and a rely-condition R , is

$$\text{ext}_\pi[P, R] \cap \text{cp}_\pi[z].$$

The next step is to formulate what the implementor must provide in return. First of all, the implementation is constrained to converge when operated in an environment which fulfills the assumptions. Secondly, any internal transition is required to satisfy the guar-condition. Thirdly, the implementor must make sure that the implementation can only become blocked in a state which satisfies the wait-condition, and that the overall effect satisfies the eff-condition if the implementation terminates. The requirements on the implementation are summed up below:

Definition 12 *Given a wait-condition W , a guar-condition G , an eff-condition E , and a structure π , then $\text{int}_\pi[W, G, E]$ denotes the set of all computations σ in π , such that:*

- $\text{len}(\sigma) \neq \infty$,
- for all $1 \leq j < \text{len}(\sigma)$, if $\lambda(\sigma_j) = i$ then $(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_\pi G$,
- if $\tau(\sigma_{\text{len}(\sigma)}) \neq \epsilon$ then $\delta(\sigma_{\text{len}(\sigma)}) \models_\pi W$,

- if $\tau(\sigma_{len(\sigma)}) = \epsilon$ then $(\delta(\sigma_1), \delta(\sigma_{len(\sigma)})) \models_{\pi} E$.

The first constraint guarantees that the implementation either terminates or deadlocks, while the second ensures that any internal transition satisfies the guar-condition. The third constraint requires the final state of a deadlocking computation to satisfy W , which means that the implementation can only become blocked in a state which satisfies the wait-condition. Finally, the fourth condition makes sure that the overall effect of a terminating computation satisfies the eff-condition.

7.3.3 Summary

Thus, if the set of auxiliary variables is empty, satisfaction can be defined as below:

Definition 13 *Given a specified program $z \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$ and a structure π , then*

- $\models_{\pi} z \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$

if and only if

- $ext_{\pi}[P, R] \cap cp_{\pi}[z] \subseteq int_{\pi}[W, G, E]$.

7.4 Satisfaction — General Case

7.4.1 Existence Check

In program-development systems (like for example [OG76a] and [Sti88]) employing the two-step strategy when reasoning with auxiliary variables, the object of the first step, where auxiliary variables are implemented as if they were ordinary programming variables, is to prove that such an extension really exists; in other words, that the actual program can be extended with auxiliary structure in such a way that the specification's requirements on both auxiliary variables and ordinary programming variables are fulfilled.

7.4.2 Removals

In our approach, this existence requirement is incorporated in the semantics. To characterise it, it is necessary to define what we mean by a removal:

Definition 14 *A removal is an expression of the form*

$$z_1 \xrightarrow{(\vartheta, \alpha)} z_2,$$

where z_1 and z_2 are programs, ϑ and α are finite, disjoint (possibly empty) sets of variables such that no element of α occurs in z_2 , and a variable of z_2 is an element of ϑ if and only if it is global. (Observe, that this does not mean that ϑ cannot have occurrences of variables that do not occur in z_2 .)

Informally, a removal is valid, written $\models z_1 \xrightarrow{(\vartheta, \alpha)} z_2$, if z_1 can be obtained from z_2 by adding auxiliary structure related to ϑ and α .

7.4.3 Restrictions

There are of course a number of constraints on the auxiliary structure:

- First of all, to make sure that the auxiliary structure has no influence on the algorithm, auxiliary variables must be restricted from occurring in the Boolean tests of while-, if- and await-statements. Furthermore, they cannot appear on the right-hand side of an assignment, unless the variable on the left-hand side is auxiliary.
- Moreover, since we want to be able to remove some auxiliary variables from a specified program without having to remove all the auxiliary variables, it is important that they do not depend upon each other. This means that if an auxiliary variable v occurs on the left-hand side of an assignment-statement, the only auxiliary variable that may occur on the right-hand side is v . In other words, to eliminate all occurrences of an auxiliary variable from a program, it is enough to remove all assignment-statements with this variable on the left-hand side. However, an assignment to an auxiliary variable may have any number of programming variables on the right-hand side.
- Finally, since auxiliary variables will only be employed to record information about state changes and synchronisation, auxiliary variables are

only allowed to be updated in connection with await- and assignment-statements.

7.4.4 Summary

What we mean by a valid removal is characterised more formally below:

Definition 15 *Given a removal $z_1 \xrightarrow{(\vartheta, \alpha)} z_2$, then*

$$\models z_1 \xrightarrow{(\vartheta, \alpha)} z_2,$$

if and only if z_1 can be obtained from z_2 by substituting

- *a statement of the form*

```

await true do
   $a_1 := u_1;$ 
   $\vdots$ 
   $a_n := u_n;$ 
   $v := r$ 
od,
```

where for all $1 \leq j, k \leq n$, $\text{var}[u_j] \subseteq \vartheta \cup \{a_j\}$, $a_j \in \alpha$, and $j \neq k$ implies $a_j \neq a_k$, for each occurrence of an assignment-statement of the form

$v := r,$

which does not occur in the body of an await-statement,

- *a statement of the form*

```

await b do
  z';
  a1 := u1;
  ⋮
  an := un
od,

```

where $\models z' \xrightarrow{(\vartheta, \alpha)} z$, for all $1 \leq j, k \leq n$, $\text{var}[u_j] \subseteq \vartheta \cup \{a_j\}$, $a_j \in \alpha$, and $j \neq k$ implies $a_j \neq a_k$, for each occurrence of an await-statement of the form

```

await b do
  z
od,

```

which does not occur in the body of another await-statement.

We will use $\models z_1 \hookrightarrow z_2$ to denote that there are sets of variables ϑ and α , such that $\models z_1 \xrightarrow{(\vartheta, \alpha)} z_2$.

7.4.5 Example

For example, if z_2 and z_1 denote the programs

```

x := x + y;
await b do
  skip
od;
x := x + y

```

and

```

await true do
   $a_1 := a_1 + x;$ 
   $x := x + y$ 
od;
await  $b$  do
  skip;
   $a_1 := a_1 + x$ 
od;
await true do
   $a_2 := x;$ 
   $x := x + y$ 
od,

```

then $\models z_1 \hookrightarrow z_2$.

7.4.6 General Case

It is now straightforward to extend definition 13 on page 70 to the general case:

Definition 16 *Given a specified program $z_1 \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E)$ and a structure π , then*

- $\models_{\pi} z_1 \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E)$

if and only if

- *there is a program z_2 such that*
 - $\models z_2 \xrightarrow{(\vartheta, \alpha)} z_1$,
 - $\text{ext}_{\pi}[P, R] \cap \text{cp}_{\pi}[z_2] \subseteq \text{int}_{\pi}[W, G, E]$.

7.4.7 Auxiliary Form

The following concept will be useful in the relative completeness proof:

Definition 17 *A specified program of the form*

$$z_1 \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$$

will be said to be of auxiliary form, if there is a program z_2 such that $\models z_1 \hookrightarrow z_2$.

7.4.8 Uniqueness Proposition

Proposition 3 *Given that*

$$\begin{aligned} &\models z_1 \hookrightarrow z_2, \\ &\models z_1 \hookrightarrow z_3, \end{aligned}$$

then $z_2 = z_3$.

Proof: Follows easily from definition 15 on page 72.
end of proof

Chapter 8

Syntactic Operators

8.1 Motivation

By a syntactic operator we mean a function, which returns an expression of the first-order language, given a finite number of expressions as arguments. Syntactic operators are not themselves symbols of the first-order language, nor may they appear in programs.

The aim of this chapter is to introduce a number of syntactic operators to reduce the size and complexity of assertions in L. We will distinguish between two different types of operators:

- those which can be expressed independently of the sequence concept,
- those whose expressibility depends upon the existence of finite sequences.

8.2 Sequence-Independent Operators

8.2.1 Reference

There are occasions when it is necessary to add hooks to all free unhooked variables in an expression. To express this, we will follow the convention suggested in [Jon90], where $\overleftarrow{r}$ denotes the result of hooking all free unhooked variables in the expression r .

8.2.2 Identity

In most cases a program can only change a rather small subset of the global state. Thus, the rely- and guar-conditions will often be of the form

$$A \wedge v_1 = \overleftarrow{v_1} \wedge \dots \wedge v_n = \overleftarrow{v_n},$$

where A constrains the part of the state that can be changed.

To simplify such binary assertions, we have found it useful to introduce a syntactic operator I_β , which constrains all variables not in the set β to remain unchanged. Hence when it occurs in a specified program of the form

$$z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E)$$

(or in the context of a glo-set ϑ and a aux-set α), it denotes the binary assertion

$$v_1 = \overleftarrow{v_1} \wedge \dots \wedge v_n = \overleftarrow{v_n},$$

where

$$(\vartheta \cup \alpha) \setminus \beta = \{v_1, v_2, \dots, v_n\}.$$

This allows us to shorten the assertion above to

$$A \wedge I_\beta.$$

When the set of changeable variables is empty, we will write I instead of $I_\emptyset$.

8.2.3 Set Quantification

To make it easier to deal with sets of variables, we will use

$$\begin{aligned} \forall V : A, \\ \exists V : A, \\ \forall \overleftarrow{V} : A, \\ \exists \overleftarrow{V} : A, \end{aligned}$$

where V is a set of unhooked variables $\{v_1, \dots, v_n\}$, to denote respectively

$$\begin{aligned}
& \forall v_1, \dots, v_n : A, \\
& \exists v_1, \dots, v_n : A, \\
& \forall \overleftarrow{v}_1, \dots, \overleftarrow{v}_n : A, \\
& \exists \overleftarrow{v}_1, \dots, \overleftarrow{v}_n : A.
\end{aligned}$$

8.2.4 Composition

Since we are dealing with binary assertions, a composition operator on binary assertions, corresponding to relational composition, is useful. Given two binary assertions A and B , then for any structure π , the assertion $A|B$ denotes the least binary relation T on the set of states, such that for all states s_1, s_2, s_3 :

- $(s_1, s_2) \models_{\pi} A$ and $(s_2, s_3) \models_{\pi} B$ implies that $(s_1, s_3) \in T$.

The proposition below shows that $A|B$ is always expressible in L .

Proposition 4 *Given two binary assertions A and B , then $A|B$ is expressible in L .*

Proof: Assume the list

$$v_1 : \Sigma_1, v_2 : \Sigma_2, \dots, v_n : \Sigma_n$$

consists of the unhooked versions of all free variables in A or B , with their respective sorts. Let

$$v'_1 : \Sigma_1, v'_2 : \Sigma_2, \dots, v'_n : \Sigma_n$$

be a list of ‘new’ variables. Then $A|B$ can be defined to be the assertion:

$$\begin{aligned}
& \exists v'_1, \dots, v'_n : \\
& A(v'_1/v_1, \dots, v'_n/v_n) \wedge B(v'_1/\overleftarrow{v}_1, \dots, v'_n/\overleftarrow{v}_n).
\end{aligned}$$

end of proof

8.2.5 Example

The sequential composition of the two assertions

$$\begin{aligned}\overleftarrow{v}_1 &= 5 \wedge v_1 = 77 \wedge v_2 \geq \overleftarrow{v}_2, \\ \overleftarrow{v}_1 &= 77 \wedge v_1 = 5 \wedge v_2 = \overleftarrow{v}_2 + 5,\end{aligned}$$

is for example characterised by

$$\begin{aligned}\exists v'_1, v'_2 : \overleftarrow{v}_1 &= 5 \wedge v'_1 = 77 \wedge v'_2 \geq \overleftarrow{v}_2 \wedge \\ v'_1 &= 77 \wedge v_1 = 5 \wedge v_2 = v'_2 + 5,\end{aligned}$$

which simplifies to

$$\begin{aligned}\exists v'_1, v'_2 : \overleftarrow{v}_1 &= 5 \wedge v'_1 = 77 \wedge v'_2 \geq \overleftarrow{v}_2 \wedge \\ v_1 &= 5 \wedge v_2 = v'_2 + 5.\end{aligned}$$

8.3 Finite-Sequence Operators

8.3.1 Transitive Closure

A transitive closure operator on binary assertions is also needed. Given a binary assertion A , then for any structure π , the assertion $A^\dagger$ denotes the least binary relation T on the set of states, such that for all states s_1, s_2, s_3 :

- $(s_1, s_2) \models_\pi A$ implies $(s_1, s_2) \in T$,
- $(s_1, s_2) \in T$ and $(s_2, s_3) \in T$, implies $(s_1, s_3) \in T$.

The proposition below shows that $A^\dagger$ is always expressible in L .

Proposition 5 *Given a binary assertion A , then $A^\dagger$ is expressible in L .*

Proof: Assume the list

$$v_1 : \Sigma_1, v_2 : \Sigma_2, \dots, v_n : \Sigma_n$$

consists of the unhooked versions of all free variables in A , with their respective sorts. Let

$$h_1 : \text{seq of } \Sigma_1, h_2 : \text{seq of } \Sigma_2, \dots, h_n : \text{seq of } \Sigma_n, m : \mathbb{N}, j : \mathbb{N}$$

be a list of ‘new’ variables. Then $A^\dagger$ can be defined to be the assertion:

$$\begin{aligned} \exists h_1, \dots, h_n, m : \\ \overleftarrow{v}_1 = h_1(1) \wedge \dots \wedge \overleftarrow{v}_n = h_n(1) \wedge \\ v_1 = h_1(m) \wedge \dots \wedge v_n = h_n(m) \wedge \\ (\forall j : 1 \leq j < m \Rightarrow \\ A(h_1(j)/\overleftarrow{v}_1, \dots, h_n(j)/\overleftarrow{v}_n, h_1(j+1)/v_1, \\ \dots, h_n(j+1)/v_n)). \end{aligned}$$

end of proof

8.3.2 Example

The transitive closure of the assertion

$$v_1 = \overleftarrow{v}_1 + 5 \wedge v_2 \geq \overleftarrow{v}_2$$

is for example characterised by

$$\begin{aligned} \exists h_1, h_2, m : \\ \overleftarrow{v}_1 = h_1(1) \wedge \overleftarrow{v}_2 = h_2(1) \wedge \\ v_1 = h_1(m) \wedge v_2 = h_2(m) \wedge \\ (\forall j : 1 \leq j < m \Rightarrow \\ h_1(j+1) = h_1(j) + 5 \wedge h_2(j+1) \geq h_2(j)). \end{aligned}$$

8.3.3 Transitive and Reflexive Closure

It is now easy to define a transitive and reflexive closure operator. Given a binary assertion A , then for any structure π , the assertion A^* stands for $(A \vee I)^\dagger$.

8.3.4 Preservation

Given a unary assertion A and a binary assertion B , in the completeness proof it will often be necessary to define a binary assertion A^B which characterises

the least set of states T , such that for all states s_1, s_2 :

- $s_1 \models_\pi A$ implies $s_1 \in T$,
- $s_1 \in T$ and $(s_1, s_2) \models_\pi B$ implies $s_2 \in T$.

Even this operator is expressible in L.

Proposition 6 *Given an unary assertion A , and a binary assertion B , then A^B is expressible in L.*

Proof: Assume the list

$$v_1 : \Sigma_1, v_2 : \Sigma_2, \dots, v_n : \Sigma_n$$

consists of the unhooked version of all free variables in A , with their respective sorts. Then A^B can be defined to be the assertion

$$\exists \overleftarrow{v_1}, \dots, \overleftarrow{v_n} : \overleftarrow{A} \wedge B^*.$$

end of proof

Chapter 9

Well-Foundedness

9.1 Motivation

Many strategies have been proposed to prove termination of loops. In [KM75] four different methods are employed on a number of examples. Two of them will be discussed here, namely Floyd's [Flo67] well-foundedness approach and the loop method first suggested in [ELW74].

The loop method has some obvious advantages. First of all, the method is easy to use and of the four methods discussed in [KM75], it is recommended as the one which can most easily be integrated into an automatic verification system.

Secondly, this approach has the additional advantage of giving an upper bound on the number of possible iterations for a specific initial state. In other words, the method can also be used to estimate the loop's time consumption.

For that reason, some authors distinguish between proving strong termination, in which case an upper bound on the number of iterations is required, and proving weak termination when it is enough that the loop eventually terminates.

9.2 Unbounded Nondeterminism

Unfortunately, with respect to LSP there is generally no upper bound on the number of iterations for a specific initial state. To see that, assume we want to prove that the program z :


```

while  $x \geq 0$  do
   $y := \text{true};$ 
   $v := v - 1;$ 
   $x := v$ 
od,

```

satisfies the specification

glo	$\{v, x, y\}$
aux	$\{\}$
pre	true
rely	$\frac{}{y} \Rightarrow y \wedge v = \frac{}{v}$
wait	false
guar	true
eff	true.

Given that the initial state s satisfies $s(x) \geq 0$ and $s(y)=\text{false}$, then although the while-statement is guaranteed eventually to terminate, there is no upper bound on the number of times it may iterate, since between the initial state and the first execution of $y := \text{true}$, the environment is free to update v as it likes. (Remember that since $x \in \text{hid}[z]$, no external transition can change the value of x .)

The reason for this is of course that the rely-condition gives rise to unbounded nondeterminism (see [Dij76], [Apt84] for a detailed discussion). The program above is ‘equivalent’ to a sequential program, namely

```

 $y, v := R(y, v);$ 
while  $x \geq 0$  do
   $y, v := R(y, v);$ 
   $y := \text{true};$ 
   $y, v := R(y, v);$ 
   $v := v - 1;$ 
   $y, v := R(y, v);$ 
   $x := v;$ 
   $y, v := R(y, v);$ 
od;
 $y, v := R(y, v)$ 

```

where $y, v := R(y, v)$ is a random assignment constrained by

$$\overline{y} \Rightarrow y \wedge v = \overline{v}.$$

9.3 Well-Foundedness Approach

We will therefore employ the well-foundedness approach. This means that our system can only be used to prove weak termination. Given an environment, the basic idea is:

- if we can show that the loop's body terminates in a state which satisfies an unary assertion A for any initial state which satisfies A and the Boolean test, and B is a binary assertion characterising the effect of the loop's body with respect to the same environment and the same restrictions on the initial state, then the loop terminates if and only if, there is no infinite sequence of states $s_1 s_2 \dots s_n \dots$, such that for all $j \geq 1$,

$$(s_j, s_{j+1}) \models_{\pi} B.$$

This property is a well-known mathematical concept. A number of authors have proposed (second order) extensions of the first-order logic to make it possible to express and reason about well-foundedness. See for example [HP73],

[dB80], [Fra86], [AP86], [SdRG89]. Unfortunately, none of these approaches consider unbounded well-foundedness in the most general case¹.

We will not attempt to deal with this problem here. Instead it will be assumed that there is a conservative (second order) extension L_{wf} of L , which, for any binary assertion A in L , allows us to formulate an assertion, denoted by $\mathbf{wf} A$, which is valid in a structure π if and only if A is well-founded in π .

We will use Tr_π to denote the set of all assertions in L_{wf} which are valid in π .

¹In [AP86] Apt and Plotkin propose a version of μ -calculus which handles countable (unbounded) nondeterminism. If all sorts are restricted to be countable, it should be possible to define a similar logic for our system.

Chapter 10

Logic of Specified Programs

10.1 Motivation

10.1.1 Formal System

The object of this chapter is to introduce a formal system, called LSP (Logic of Specified Programs), which consists of sixteen decomposition-rules for specified programs, axioms and deduction-rules for L_{wf} , and all valid removals as axioms. The decomposition-rules are split into two categories:

- The rules needed for top-down development of programs; namely the consequence-, pre-, access-, skip-, assignment-, block-, sequential-, if-, while-, parallel-, await- and elimination-rules.

These rules will be called the basic rules, and it will later be shown that the basic rules are sufficient to prove relative completeness. Thus, since the decomposition-rules are compositional, it follows that LSP_B is compositionally complete (see [KKZ89], [Zwi89] for a more detailed discussion).

- The rules introduced to simplify the development and make it possible to take advantage of already finished designs; namely the effect-, global-, auxiliary- and introduction-rules.

These rules will be called the adaptation rules. Unfortunately, this set of rules is not strong enough to ensure adaptation completeness (see [Zwi89])¹, namely that

¹It is not difficult to find a stronger set, but we do not know how to formulate a set

- if the specification of a program, whose internal structure is unknown, implies another specification for the same program, then the proof system admits a formal deduction of that fact.

For example, since

$$z \text{ sat } (\{v\}, \{\}) :: (\text{true}, v = \overleftarrow{v}, \text{false}, v \geq \overleftarrow{v}, v = \overleftarrow{v}),$$

is valid only if z leaves v unchanged, it should be possible to deduce

$$z \text{ sat } (\{v\}, \{\}) :: (\text{true}, v = \overleftarrow{v}, \text{false}, v = \overleftarrow{v}, v = \overleftarrow{v}).$$

However, there is no way to do this given the current set of rules.

Axioms and deduction-rules for first-order systems can be found in almost any text-book on logic, and this will therefore not be discussed here.

10.1.2 Decomposition-Rules

The decomposition-rules are all of the form

$$\frac{\begin{array}{c} Prem_1 \\ \vdots \\ Prem_n \end{array}}{Concl},$$

where the conclusion, $Concl$, is a schema characterising a specified program, and the j 'th premise, $Prem_j$, is a schema denoting either a specified program, a formula in L_{wf} or a removal.

Moreover, a decomposition-rule should be interpreted as follows:

- given a structure π and $n+1$ matching specified programs formulas and removals², then the conclusion is valid in π , if each of the n premises are valid in π .

which is adaptation complete.

² n premises + conclusion = $n + 1$.

10.1.3 Proofs

A proof in LSP is a finite tree where

- each node represents a removal, a formula of L_{wf} or a specified program,
- all leaves are axioms,
- a node is either a leaf or can be deduced from its immediate sons by one of the deduction- or decomposition-rules.

10.1.4 Depth

A proof's depth is 0 if the root is a formula of L_{wf} or a removal, 1 if the root is a specified program without sons; otherwise $n+1$, where n is the maximum depth of its immediate subproofs.

10.1.5 Notation

LSP restricted to the set of basic rules will be denoted LSP_B . No proof in LSP_B has a removal in its tree. This means that auxiliary variables are only a part of the logic and do not have to be implemented.

Given a specified program ψ , we will use

$$\vdash \psi$$

to denote that ψ is provable in LSP,

$$\vdash_B \psi$$

to denote that ψ is provable in LSP_B , and

$$\begin{aligned} Tr_\pi &\vdash \psi, \\ Tr_\pi &\vdash_B \psi \end{aligned}$$

to denote that ψ is provable in LSP, respectively LSP_B , given the assertions in Tr_π as axioms.

10.1.6 Current Set of Rules

There are many ways to formulate the different decomposition-rules, and we are not claiming the given set is the best possible for all applications. However, since this set is both sound and relatively complete, it can always be used to deduce new more user friendly decomposition-rules.

10.2 Consequence-Rule

The consequence-rule is probably the easiest to understand. Basically, the pre- and rely-conditions can be strengthened, because any program that behaves correctly given the original pre- and rely-conditions, will also behave correctly when they are strengthened.

Similarly, we may weaken the wait-, guar- and eff-conditions, because any program that satisfies the stronger wait-, guar- and eff-conditions, will also satisfy the new ones:

$$\begin{array}{l}
 z \text{ sat } (\vartheta, \alpha) :: (P_1, R_1, W_1, G_1, E_1) \\
 P_2 \Rightarrow P_1 \\
 R_2 \Rightarrow R_1 \\
 W_1 \Rightarrow W_2 \\
 G_1 \Rightarrow G_2 \\
 E_1 \Rightarrow E_2 \\
 \hline
 z \text{ sat } (\vartheta, \alpha) :: (P_2, R_2, W_2, G_2, E_2)
 \end{array}$$

10.3 Pre-Rule

The pre-rule is also straightforward. If the actual program is employed in a state which does not satisfy the pre-condition, there are no constraints on its behaviour. Thus, we may restrict the eff-condition to transitions from states which satisfy the pre-condition:

$$\begin{array}{l}
 z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E) \\
 \hline
 z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, \overline{P} \wedge E)
 \end{array}$$

10.4 Access-Rule

Since for any program z , the environment respects $hid[z] \cap \vartheta$, a rule which allows us to weaken the rely-condition is also needed:

$$\frac{z \text{ sat } (\vartheta, \alpha) :: (P, R \wedge v = \overline{v}, W, G, E)}{z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E)}$$

where $v \in hid[z] \cap \vartheta$

10.5 Skip-Rule

The skip-statement generates only one internal transition. Since this transition leaves the state unchanged and since the guar-condition is reflexive, it is clear that any internal transition satisfies the guar-condition. Moreover, since each transition due to the environment is assumed to satisfy the rely-condition, which is both reflexive and transitive, it follows that the overall effect of the skip-statement satisfies the rely-condition:

$$\text{skip sat } (\vartheta, \alpha) :: (P, R, W, G, R)$$

10.6 Assignment-Rule

10.6.1 Introduction

The assignment-rule is more complicated. Although there is only one internal transition, we must allow auxiliary variables to be updated in the same atomic step. The reason is of course that due to the definition of the satisfaction relation, the execution of a statement z_1 of the form

$$v := r$$

corresponds to the execution of a statement z_2 of the form


```

await true do
   $a_1 := u_1;$ 
   $\vdots$ 
   $a_m := u_m;$ 
   $v := r$ 
od;

```

where $\models z_2 \hookrightarrow z_1$. In other words, z_1 may be extended with auxiliary structure in such a way that the auxiliary variables $a_1, \dots, a_m$ are updated in the same atomic step as v .

To simplify the discussion, we will first deal with the situation where the set of auxiliary variables is empty, and thereafter show how the rule can be extended to cover the general case.

10.6.2 Without Auxiliary Variables

Since the assignment-statement is interpreted as atomic in our operational semantics, there is only one internal transition. Clearly, this transition must be shown to satisfy the guar-condition.

Moreover, the assignment can only take place in a state that can be reached from a state which satisfies the pre-condition by a finite number of external transitions, and since the rely-condition is both transitive and reflexive, the following rule is sufficient:

$$\frac{\overline{P^R} \wedge v = \overline{r} \wedge I_{\{v\}} \Rightarrow G \wedge E}{v := r \text{ sat } (\varnothing, \{\}) :: (P, R, W, G, R|E|R)}$$

10.6.3 General Case

It is now straightforward to extend the rule to handle the general situation. The only real difference is that the premise in this case must also guarantee that the assignment-statement can be extended with auxiliary structure in such a way that the specified changes to both the auxiliary variables and the programming variables will indeed take place:

$$\frac{\overline{P^R} \wedge v = \overleftarrow{r} \wedge I_{\{v\} \cup \alpha} \wedge \bigwedge_{a \in \alpha} a = \overleftarrow{u_a} \Rightarrow G \wedge E}{v := r \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, R|E|R)}$$

where for all $a \in \alpha$, $\text{var}[u_a] \subseteq \vartheta \cup \{a\}$

10.7 Block-Rule

The block-rule is also easy to understand. When employed in a top-down style, the set of global programming variables is extended with the new ones. Moreover, since no program running in parallel can access these variables, we may restrict the environment to leave them unchanged:

$$\frac{z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R \wedge \bigwedge_{j=1}^n v_j = \overleftarrow{v_j}, W, G, E)}{\text{begin loc } v_1, \dots, v_n; z \text{ end \underline{sat} } (\vartheta \setminus \bigcup_{j=1}^n \{v_j\}, \alpha) :: (P, R, W, G, E)}$$

It follows from the constraints on specified programs that the variables

$$v_1, \dots, v_n$$

do not occur in the specification

$$(\vartheta \setminus \bigcup_{j=1}^n \{v_j\}, \alpha) :: (P, R, W, G, E).$$

10.8 Sequential-Rule

The sequential-rule is straightforward. Basically, the first component's effect-condition must imply the second component's pre-condition. This explains why P_2 occurs in the first statement's effect-condition:

$$\frac{z_1 \text{ \underline{sat} } (\vartheta, \alpha) :: (P_1, R, W, G, P_2 \wedge E_1) \quad z_2 \text{ \underline{sat} } (\vartheta, \alpha) :: (P_2, R, W, G, E_2)}{z_1; z_2 \text{ \underline{sat} } (\vartheta, \alpha) :: (P_1, R, W, G, E_1|E_2)}$$

10.9 If-Rule

In this case there is not much to explain. Basically, because the if-statement's Boolean test has been restricted to have no occurrences of variables accessible by the environment, the truth value of the Boolean test is maintained by the environment. Thus, the following rule is sufficient:

$$\frac{\begin{array}{l} z_1 \text{ sat } (\vartheta, \alpha) :: (P \wedge b, R, W, G, E) \\ z_2 \text{ sat } (\vartheta, \alpha) :: (P \wedge \neg b, R, W, G, E) \end{array}}{\text{if } b \text{ then } z_1 \text{ else } z_2 \text{ fi sat } (\vartheta, \alpha) :: (P, R, W, G, E)}$$

10.10 While-Rule

In the same way as for the if-rule, interference before the first evaluation of the Boolean test has no influence on the tests outcome. Moreover, the falsity of the Boolean test will be preserved after the while-statement terminates. As explained above, we are content with proving weak (conditional) termination, thus it is enough to show that an assertion characterising the effect of the loop's body in the actual environment is well-founded when considered as a binary relation on states. We can therefore formulate a rule fairly similar to the while-rule in [Jon86] (if we ignore that the latter is not dealing with rely-, wait- and guar-conditions):

$$\frac{\begin{array}{l} \text{wf } Z \\ z \text{ sat } (\vartheta, \alpha) :: (P \wedge b, R, W, G, P \wedge Z) \end{array}}{\text{while } b \text{ do } z \text{ od sat } (\vartheta, \alpha) :: (P, R, W, G, (Z^\dagger \vee R) \wedge \neg b)}$$

The only real difference is that, because of possible interference from the environment, the I in the conclusion's eff-condition has been replaced by R .

10.11 Parallel-Rule

10.11.1 Simple Case

We will first design a rule for the case when the two component programs do not deadlock (given their respective environments). Obviously, the premises

must make sure that the component programs are compatible with respect to mutual interference.

It should also be clear that the parallel composition of two programs is only guaranteed to converge if called in an environment in which both component programs are guaranteed to converge. Finally, the statements overall effect is also the overall effect of both component programs:

$$\frac{z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_1, \text{false}, G \wedge R_2, E_1) \quad z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_2, \text{false}, G \wedge R_1, E_2)}{\{z_1 \parallel z_2\} \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_1 \wedge R_2, \text{false}, G, E_1 \wedge E_2)}$$

10.11.2 General Case

To formulate the general rule, it is enough to observe that z_1 is guaranteed to be released whenever it becomes blocked in a state in which z_2 cannot become blocked or terminate. Similarly, z_2 is guaranteed to be released in any state in which z_1 cannot become blocked or terminate:

$$\frac{\neg(W_1 \wedge E_2) \wedge \neg(W_2 \wedge E_1) \wedge \neg(W_1 \wedge W_2) \quad z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_1, W \vee W_1, G \wedge R_2, E_1) \quad z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_2, W \vee W_2, G \wedge R_1, E_2)}{\{z_1 \parallel z_2\} \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_1 \wedge R_2, W, G, E_1 \wedge E_2)}$$

One possible alternative is:

$$\frac{\neg(W_1 \wedge E_2) \wedge \neg(W_2 \wedge E_1) \wedge \neg(W_1 \wedge W_2) \quad z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, (R \vee G_2)^\dagger, W \vee W_1, G \wedge G_1, E_1) \quad z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P, (R \vee G_1)^\dagger, W \vee W_2, G \wedge G_2, E_2)}{\{z_1 \parallel z_2\} \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E_1 \wedge E_2)}$$

10.11.3 Generalised Parallel-Rule

The parallel-rule can also be generalised to deal with more than two processes:

$$\frac{\neg(W_j \wedge \bigwedge_{k=1, k \neq j}^m (W_k \vee E_k))_{1 \leq j \leq m} \quad z_j \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_j, W \vee W_j, G \wedge \bigwedge_{k=1, k \neq j}^m R_k, E_j)_{1 \leq j \leq m}}{\{z_1 \parallel \dots \parallel z_m\} \underline{\text{sat}} (\vartheta, \alpha) :: (P, \bigwedge_{j=1}^m R_j, W, G, \bigwedge_{j=1}^m E_j)}$$

Obviously, $\{z_1 \parallel \dots \parallel z_m\}$ denotes any program that can be obtained from $\{z_1 \parallel \dots \parallel z_m\}$ by adding curly brackets. This rule can be deduced from the rules in LSP_B , and it will from now on be referred to as the generalised parallel-rule.

10.12 Await-Rule

10.12.1 Introduction

The await-rule is closely related to the assignment-rule; there is only one internal transition, and auxiliary variables are allowed to be updated in the same atomic step. Thus, the execution of a statement z_1 of the form

```
await b do
  z
od
```

corresponds to the execution of a statement z_2 of the form

```
await b do
  z';
  a_1 := u_1;
  ⋮
  a_m := u_m
od,
```

where $\models z_2 \hookrightarrow z_1$ and $\models z' \hookrightarrow z$.

To simplify the discussion, we will first deal with the situation where the set of auxiliary variables is empty, and thereafter show how the rule can be extended to cover the general case.

10.12.2 Without Auxiliary Variables

Clearly, any state which does not satisfy the Boolean test and can be reached from a state which satisfies the pre-condition by a finite number of external transitions must satisfy the wait-condition.

The operational semantics guarantees that the environment cannot interfere with the body of an await-statement. Moreover, the statement's body is required to terminate for any state which satisfies the Boolean test and can be reached from a state which satisfies the pre-condition by a finite number of external transitions. Finally, the effect of the await-statement's body is required to satisfy the overall guar-condition:

$$\frac{P^R \wedge \neg b \Rightarrow W \quad z \text{ sat } (\vartheta, \{\}) :: (P^R \wedge b, I, \text{false}, \text{true}, G \wedge E)}{\text{await } b \text{ do } z \text{ od } \text{sat } (\vartheta, \{\}) :: (P, R, W, G, R|E|R)}$$

10.12.3 General Case

It is now straightforward to formulate the rule for the general case:

$$\frac{P^R \wedge \neg b \Rightarrow W \quad E_1 | (\bigwedge_{a \in \alpha} a = \overleftarrow{u}_a \wedge I_\alpha) \Rightarrow G \wedge E_2 \quad z \text{ sat } (\vartheta, \alpha) :: (P^R \wedge b, I, \text{false}, \text{true}, E_1)}{\text{await } b \text{ do } z \text{ od } \text{sat } (\vartheta, \alpha) :: (P, R, W, G, R|E_2|R)}$$

where for all $a \in \alpha$, $\text{var}[u_a] \subseteq \vartheta \cup \{a\}$

10.13 Elimination-Rule

We also need a rule that allows us to eliminate auxiliary structure from the specification:

$$\frac{z \text{ sat } (\vartheta, \alpha \cup \{a\}) :: (P, R, W, G, E)}{z \text{ sat } (\vartheta, \alpha \setminus \{a\}) :: (\exists a : P, \forall \overleftarrow{a} : \exists a : R, W, G, E)}$$

Since the conclusion is required to be a specified programs, it follows that W , R and E have no occurrences of a .

10.14 Effect-Rule

It is also true that any state reachable from the pre-condition can be reached by the transitive closure of the rely- and guar-conditions. The eff-condition can therefore be strengthened with their closure:

$$\frac{z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E)}{z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E \wedge (R \vee G)^\dagger)}$$

10.15 Global-Rule

The global-rule allows us to introduce a new variable.

$$\frac{z \text{ sat } (\vartheta \setminus \{v\}, \alpha) :: (P, R, W, G, E)}{z \text{ sat } (\vartheta \cup \{v\}, \alpha) :: (P, R, W, G \wedge v = \overline{v}, E)}$$

The premise and the constraints on specified programs imply that the variable does not occur in the actual program. Thus no internal transition can change its value.

10.16 Auxiliary-Rule

Similarly, the auxiliary-rule allows us to introduce a new auxiliary variable:

$$\frac{z \text{ sat } (\vartheta, \alpha \setminus \{a\}) :: (P, R, W, G, E)}{z \text{ sat } (\vartheta, \alpha \cup \{a\}) :: (P, R, W, G \wedge a = \overline{a}, E)}$$

10.17 Introduction-Rule

Although our logic allows us to reason about auxiliary variables without having to implement them, the following rule can be useful if we have already

developed a program with respect to a specification and later want to use the same algorithm in another connection:

$$\frac{z_1 \xrightarrow{(\vartheta \setminus \alpha, \alpha)} z_2 \quad z_1 \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)}{z_2 \text{ sat } (\vartheta \setminus \alpha, \alpha) :: (P, R, W, G, E)}$$

Chapter 11

Simplifying Notation

11.1 Motivation

Before applying LSP to the Dining-Philosophers, the Bubble-Lattice-Sort and the Set-Partition algorithms, we will introduce some simplifying notation. The new concepts are introduced in a rather informal style. However, based on the way they are used in the different examples, it should not be difficult to grasp their intended meaning¹.

11.2 Operations

11.2.1 Main Structure

First of all, to make specifications more readable, a VDM-related [Jon90] operation concept of the form

¹The reason why we have not given a formal semantics is that we are currently experimenting with this notation, and it is still subject to change.

<i>Name(In) Out</i>	
glo	<i>g_dcl</i>
aux	<i>a_dcl</i>
pre	<i>P</i>
rely	<i>R</i>
wait	<i>W</i>
guar	<i>G</i>
eff	<i>E</i>

has been found helpful. Not surprisingly, *Name* is the name of the operation, *In* is the list of input parameters, while *Out* is the list of output parameters. Moreover, global variables are declared in *g_dcl*, while *a_dcl* is used to declare auxiliary variables. Finally, *P*, *R*, *W*, *G* and *E* denote respectively the pre-, rely-, wait-, guar- and eff-conditions.

11.2.2 Variables

The set of input parameters is constrained to be disjoint from the set of output parameters. Syntactically the parameters are unhooked variables. Similarly, the set of global variables is required to be disjoint from the set of auxiliary variables, and no global or auxiliary variable can occur as a parameter.

The five conditions may contain occurrences of global variables that are not declared in *g_dcl*. However, no such variable is allowed to appear in the final code implementing the operation. In other words, no global variable may occur in the implementation unless it is listed in *g_dcl*.

Auxiliary variables that will be ‘updated’ by the implementation must be listed in *a_dcl*. Variables local to the implementation are not allowed to appear in the specification.

If a free variable is hooked in the rely-, guar- or eff-conditions, then it is either auxiliary or global. If it is clear from *g_dcl* and *a_dcl* that a particular variable cannot be changed by the operation, its environment or both, then this will not be restated in the operation’s rely-, guar- and eff-conditions.

For example, if *v* is a global variable not mentioned in *g_dcl*, then it is clear that any internal transition will satisfy $v = \overleftarrow{v}$, but to keep the specifications as simple as possible this does not have to be restated in the guar-condition,

although it may be added as an extra conjunct when proofs are undertaken.

11.2.3 Observable and Hidden Changes

In VDM, which covers only non-interfering programs, it is indicated in the declaration whether the operation has write access to a variable, only read access or no access at all. If an operation has no write access to a variable, clearly its value will be left unchanged.

Unfortunately, when dealing with concurrency the situation is a lot more complicated. First of all, other programs running in parallel may interfere. Moreover, because of the await-statement, there can be changes to the global structure due to the implementation that are hidden from the environment and vice versa.

If a particular global variable is not updated outside await-statements, and the variable's value upon entry to an await-statement is always equal to its value when the await-statement terminates, then changes to this variable are not observable from the outside.

Therefore, instead of declaring global and auxiliary variables according to whether they are writable or only readable, we will distinguish variables that can be changed in an observable way from those that cannot. Both global variables and auxiliary variables will be declared according to the following convention. Let v be a variable of sort T , then:

- **ioeo** $v : T$ — (internal observe, external observe.)

Means that the operation can change the value of v in such a way that this can be observed by the environment. Similarly, the operation can observe changes to v that are due to the environment.

- **ioeh** $v : T$ — (internal observe, external hide.)

Means that the operation can change the value of v in such a way that this can be observed by the environment, while any changes to v due to the environment are hidden from the operation.

- **iheo** $v : T$ — (internal hide, external observe.)

Means that any changes to v due to the operation are hidden from the environment, while the operation can observe updates to v caused by the environment.

- **ih** $v : T$ — (internal hide, external hide.)

Means that any changes to v due to the operation are hidden from the environment, and the operation cannot observe any of the environment's changes to v .

11.2.4 Splitting of Objects

We have also found it convenient to split up objects consisting of more than one memory location and declare them separately. For example, given a global variable

$$A : \text{array } \{1, \dots, 9\} \text{ to } \mathbb{B},$$

and assuming we are specifying an operation which needs **ioeo**-access to $A(8)$, but never accesses the rest of A , then

$$\text{ioeo } A(8) : \mathbb{B},$$

will be the only reference to A in g_dcl .

11.3 Assertions Occurring in the Code

11.3.1 Sequential Case

When proving properties of programs it is often useful to insert assertions into the code. For example, the sequential program

```
{true}
while  $x > 0$  do
  { $x > 0$ }
   $x := x - 1$ 
od
{ $x \leq 0$ }
```

has three such assertions. The first and last characterise respectively the initial and final states, while the one in the middle describes the state each

time the program counter is ‘situated’ between the Boolean test and the assignment-statement.

We will insert assertions in a similar style. However, because the assertions may have occurrences of hooked variables, because the environment may interfere, and because of the way our operational model is related to a more intuitive interpretation of parallel programs, it is necessary to discuss the meaning of such assertions in more detail.

11.3.2 Hooking

Assertions occurring in the code will have occurrences of hooked variables when this is convenient. The hooked variables are supposed to refer to the initial state with respect to the particular piece of code in which the assertion occur. For example, in the annotated program:

```

{true}
 $x := x + 5;$ 
 $\{x = \overline{x} + 5\}$ 
 $x := x + 3;$ 
 $\{x = \overline{x} + 8\}$ 
 $x := x + 2$ 
 $\{x = \overline{x} + 10\},$ 

```

the second assertion states that whenever the program counter is situated between the first and the second assignment-statement, the difference between the current value of x and the initial value of x is five, the third assertion states that whenever the program counter is situated between the second and the third assignment-statement, the difference between the current value of x and the initial value of x is eight, while the fourth assertion states that whenever the program counter is situated after the third assignment-statement, the difference between the current value of x and the initial value of x is ten.

11.3.3 Interference

It is important to realise that assertions inserted into the code are supposed to be preserved by the actual rely-condition. In the example above it is implicitly assumed that the environment leaves x unchanged. If we change the rely-condition to

$$x \geq \underline{x},$$

we end up with the following annotated program:

```

{true}
x := x + 5;
{x ≥ x + 5}
x := x + 3;
{x ≥ x + 8}
x := x + 2
{x ≥ x + 10}.

```

11.3.4 Two Interpretations

Await-statements play an important rôle in the examples below. Actually, they can be used in two different ways:

- First of all, an await-statement can be employed to guarantee mutual exclusion in cases where this is necessary to ensure the correctness of the actual algorithm. In other words, cases when the replacement of the await-statement with the await-statement's body would result in a different 'behaviour'.
- Secondly, an await-statement can be applied as an abstraction tool; to simplify the development and verification of an algorithm. In this case, the replacement of the await-statement with the await-statement's body 'has no real effect' on the algorithm's 'behaviour'. The 'amount of concurrency' is not reduced by using the await-statement as an abstraction tool. The only disadvantage is the unnecessary hiding of variables that are already 'hidden' by some sort of mutual exclusion algorithm in the actual program.

Some readers may find the second alternative a bit surprising. After all, in the operational model, the execution of an await-statement is represented by one internal transition, while the second option above seems to imply that the await-statement is non-atomic. Actually, by a similar argument, they might claim that:

- Since any computation consists of an infinite sequence of nonoverlapping external and internal transitions, there can be no real concurrency at all. In other words, that our operational model is inadequate.

But this is a misunderstanding. An atomic step should not be interpreted as something immediate, without extension in time, nor should it be understood as a program that must be run in isolation, i.e. until it has terminated no other process is allowed to progress.

Instead, one should consider an atomic step as a program, which, while it is executing, has some sort of device to deny other processes access to the memory locations it is accessing. This way, although the environment may change any of the remaining memory locations, no other process can interfere, and the atomic transition can be reasoned about as a sequential program.

Moreover, since the sets of memory locations accessed by two concurrent atomic transitions are disjoint, it follows that they may just as well be considered to have taken place in sequential order without any overlap in time. This motivates our choice of operational model.

Unfortunately, this distinction between the operational model and a more intuitive interpretation, where atomic steps may overlap in time, can easily lead to confusion when it comes to interpreting assertions occurring in the code.

If one think in terms of the operational model, which is what we recommend, assertions occurring in the code should be interpreted as follows:

- An assertion occurring in the code is true whenever the program counter reaches the assertion's address in the actual program.

This corresponds to how the assertions are interpreted in the example above. On the other hand, if one insists on thinking in terms of the more intuitive model, assertions occurring in the code should be interpreted as follows:

- An assertion occurring in the code is true whenever the program counter reaches the assertion's address, and all the variables occurring in the assertion are accessible, i.e. not occupied by an atomic transition currently under execution.

Not surprisingly, the pre-, rely-, wait-, guar- and eff-conditions can be given two different interpretations in a similar way.

11.4 *Inv* and *Dyn*

In the examples we have found it useful to formulate two special assertions; a global invariant denoted by *Inv* and a dynamic invariant called *Dyn*².

The global invariant is a unary assertion which is true initially and thereafter preserved by both the rely- and guar-conditions.

The dynamic invariant is a binary assertion which is supposed to model any finite sequence of consecutive rely and guar steps. For that reason *Dyn* is required to be both reflexive and transitive. This means that any state transition which satisfies the rely-condition, will also satisfy *Dyn*. The same is of course also true for the guar-condition.

Unfortunately, *Inv* and *Dyn* are often false inside await-statements, which means that we cannot automatically interpret

glo	ϑ
aux	α
pre	P
rely	R
wait	W
guar	G
eff	E

as

glo	ϑ
aux	α
pre	$P \wedge Inv$
rely	$R \wedge Dyn \wedge (\overleftarrow{Inv} \Rightarrow Inv)$
wait	$W \wedge Inv$
guar	$G \wedge Dyn \wedge (\overleftarrow{Inv} \Rightarrow Inv)$
eff	$E \wedge Inv \wedge Dyn,$

since we also want to specify subprograms of await-statements. Instead it will be stated explicitly in each particular specification whether *Inv* and *Dyn* are valid or not.

It may be argued that *Inv* and *Dyn* should have arguments indicating

²Jones arrived at a similar conclusion in [Jon81], although his definition of a dynamic invariant is slightly different from ours.

which part of the global state they affect. We have ignored this because of the large number of arguments needed in some cases³.

³The number of arguments could probably have been reduced if we had introduced records, in the style of for example VDM, to structure the global state.

Chapter 12

Dining-Philosophers

12.1 Task

12.1.1 Mutual Exclusion

Two statements are *mutually exclusive* if they cannot be executed simultaneously. Obviously, if the two statements occur in the same process, then they are mutually exclusive unless one of them is a substatement of the other. If the statements are taken from two different processes, on the other hand, it is often difficult to decide if they are mutually exclusive or not.

If the sets of memory locations accessed by the two statements overlap, one way to secure mutual exclusion is to place them in the bodies of two different await-statements. If the sets of memory locations are disjoint, mutual exclusion is normally not needed.

Unfortunately, from time to time problems arise where the programmer must do the synchronisation himself and provide his own code to secure mutual exclusion. Such a situation will be discussed in the example below.

12.1.2 Workshop

Assume M ($M \geq 3$) philosophers are attending a workshop. During this workshop each philosopher is scheduled to eat Q times¹. The rest of the time, he is supposed to spend thinking or discussing with his colleagues. The

¹Both M and Q are constants — not variables.

food is served on a round table. Each philosopher has his own plate, which means that there are M plates.

The dish is spaghetti. Unfortunately the spaghetti is so long and tangled that a philosopher needs two forks to eat it.

By some mistake there are only M forks on the table; one between each plate. Moreover, no philosopher would consider using any other fork than the ones to his immediate left or right.

It is also assumed that when a philosopher grabs his forks, he grabs both of them at the same time. In other words, a philosopher holds either two forks or no forks at all. This means that two neighbours cannot eat at the same time.

Our task is to write a program for each philosopher that provides this synchronisation. The algorithm presented here is closely related to the one discussed in [OG76a].

12.2 Development

12.2.1 Data Structure

Let

$$D_{(n,m)} = \{n, n+1, \dots, m\}.$$

We will employ one global array

$$Frk : \text{array } D_{(0,M-1)} \text{ to } D_{(0,2)},$$

and two auxiliary arrays

$$\begin{aligned} Num &: \text{array } D_{(0,M-1)} \text{ to } D_{(0,Q)}, \\ Eating &: \text{array } D_{(0,M-1)} \text{ to } D_{(0,1)}. \end{aligned}$$

To avoid making a special case of the ‘first’ and the ‘last’ philosopher, we will apply arithmetic modulo M to access the ‘previous’ and the ‘next’ philosopher. The symbol $\ominus$ will be used to denote subtraction modulo M , while $\oplus$ stands for addition modulo M .

This means that we would like the arrays to be defined on $D_{(0,M-1)}$, and that the process $Phil(j-1)$ simulates the j ’th philosopher. Moreover, at

any time, $Frk(j)$ equals the number of forks available to philosopher $j + 1$, while $Num(j)$ gives the number of times he has eaten.

Finally, since we want to prove that our implementation satisfies the desired mutual exclusion property, namely that two neighbours cannot eat at the same time, we will use $Eating(j)$ to indicate if philosopher $j + 1$ is holding his forks or not.

If philosopher $j+1$ is eating then $Eating(j) = 1$, otherwise $Eating(j) = 0$. The reason why the range of $Eating$ is $D_{(0,1)}$ and not $\mathbb{B}$, is that this allows us to formulate the following invariant

$$\forall j \in D_{(0,M-1)} Frk(j) = 2 - (Eating(j \oplus 1) + Eating(j \ominus 1)).$$

Moreover, since a philosopher can only eat if he holds both his forks, we will also insist on the truth of

$$\forall j \in D_{(0,M-1)} Eating(j) = 1 \Rightarrow Frk(j) = 2.$$

From now on, Inv will denote the conjunction of these two assertions.

12.2.2 Main Program

This means that the main program will be of the form:

```

{true}

begin
  loc  $Frk$ ;
   $Init()$ ;

  { $\forall j \in D_{(0,M-1)} Eating(j) = 0 \wedge Num(j) = 0 \wedge Inv$ }

  { $Phil(0) \parallel Phil(1) \parallel \dots \parallel Phil(M-1)$ }
end

{ $\forall j \in D_{(0,M-1)} Eating(j) = 0 \wedge Num(j) = Q \wedge Inv$ }.

```

12.2.3 Specification of Processes

Since $Phil(l)$ only waits if $Phil(l \ominus 1)$ or $Phil(l \oplus 1)$ are eating, it follows that:

$$\begin{array}{ll}
Phil(l : D_{(0, M-1)}) & \\
\text{glo } \text{ioeo } Frk(l \ominus 1) : D_{(0, 2)}, & \\
& Frk(l \oplus 1) : D_{(0, 2)}, \\
& \text{iheo } Frk(l) : D_{(0, 2)}, \\
\text{aux } \text{ioeh } Eating(l) : D_{(0, 1)}, & \\
& \text{ioeh } Num(l) : D_{(0, Q)} \\
\\
\text{pre } \underline{Eating(l) = 0 \wedge Num(l) = 0 \wedge Inv} & \\
\text{rely } \underline{Inv \Rightarrow Inv} & \\
\text{wait } (\underline{Eating(l \ominus 1) = 1 \vee Eating(l \oplus 1) = 1}) \wedge \underline{Inv} & \\
\text{guar } ((\underline{Eating(l) = Eating(l)} \wedge \underline{Num(l) = Num(l)}) \vee & \\
& (\underline{Eating(l) = 0 \wedge Eating(l) = 1 \wedge Num(l) = Num(l)}) \vee \\
& (\underline{Eating(l) = 1 \wedge Eating(l) = 0 \wedge Num(l) = Num(l) + 1})) \wedge \\
& \underline{(Inv \Rightarrow Inv)} \\
\text{eff } Eating(l) = 0 \wedge Num(l) = Q \wedge Inv. &
\end{array}$$

12.2.4 Correctness Proof

We will now show that the parallel composition of these M processes has the desired effect. Firstly, observe that the constraints imposed on the global memory locations by the declarations in the respective processes are satisfied. Secondly, since it is clear that

$$\begin{aligned}
& \vdash Eating(l \oplus 1) = 1 \wedge Inv \Rightarrow \\
& \quad \neg((Eating(l \oplus 1) = 0 \wedge Num(l + 1) = Q \wedge Inv) \vee \\
& \quad ((Eating(l) = 1 \vee Eating(l \oplus 2) = 1) \wedge Inv)),
\end{aligned}$$

$$\begin{aligned}
& \vdash Eating(l \ominus 1) = 1 \wedge Inv \Rightarrow \\
& \quad \neg((Eating(l \ominus 1) = 0 \wedge Num(l + 1) = Q \wedge Inv) \vee \\
& \quad ((Eating(l) = 1 \vee Eating(l \ominus 2) = 1) \wedge Inv)),
\end{aligned}$$

it follows that whenever $Phil(l)$ is blocked, then there is at least one other process which is enabled. Thus, it can be deduced by the consequence- and generalised parallel-rules that the parallel-statement satisfies:

$$\begin{array}{ll}
\text{glo} & \{Frk\} \\
\text{aux} & \{Num, Eating\} \\
\text{pre} & \forall j \in D_{(0, M-1)} Eating(j) = 0 \wedge Num(j) = 0 \wedge Inv \\
\text{rely} & I \\
\text{wait} & \text{false} \\
\text{guar} & \overline{Inv} \Rightarrow Inv \\
\text{eff} & \forall j \in D_{(0, M-1)} Eating(j) = 0 \wedge Num(j) = Q \wedge Inv.
\end{array}$$

12.2.5 Process Decomposition

The final step is to implement $Phil(l)$. Assume that $Eat(l)$ and $Study(l)$ are operations whose guar-condition implies that

$$Num = \overline{Num} \wedge Frk = \overline{Frk} \wedge Eating = \overline{Eating}.$$

Then, since a philosopher will only have to wait if one of his two neighbours is eating, it follows that $Phil(l)$ should be of the following form:

```

{Num(l) = 0 ∧ Eating(l) = 0 ∧ Inv}

begin
  loc j;
  j := 0;
  while j < Q do

    {Num(l) < Q ∧ Eating(l) = 0 ∧ Inv}

    await Frk(l) = 2 do GrabFrks(l) od;

    {Num(l) < Q ∧ Eating(l) = 1 ∧
     Frk(l ⊖ 1) < 2 ∧ Frk(l ⊕ 1) < 2 ∧ Inv}

    Eat(l);

    {Num(l) < Q ∧ Eating(l) = 1 ∧
     Frk(l ⊖ 1) < 2 ∧ Frk(l ⊕ 1) < 2 ∧ Inv}

    await true do DropFrks(l) od;

    {Num(l) ≤ Q ∧ Eating(l) = 0 ∧ Inv}

    Think(l);
    j := j + 1;

    {Num(l) ≤ Q ∧ Eating(l) = 0 ∧ Inv}

  od
end

{Num(l) = Q ∧ Eating(l) = 0 ∧ Inv}.

```

12.2.6 Implementing Atomic Statements

The final step is to implement the atomic statements $GrabFrks(l)$ and $DropFrks(l)$. The first one must satisfy:

```

glo    {Frk}
aux    {Num, Eating}
pre    Num(l) < Q ∧ Eating(l) = 0 ∧ Frk(l) = 2 ∧ Inv
rely    I
wait    false
guar    true
eff    Num(l) < Q ∧ Eating(l) = 1 ∧ Frk(l) = 2 ∧
        Frk(l ⊖ 1) =  $\overline{Frk(l ⊖ 1) - 1}$  ∧
        Frk(l ⊕ 1) =  $\overline{Frk(l ⊕ 1) - 1} ∧ Inv$ .

```

Thus,

$$Frk(l ⊖ 1) := Frk(l ⊖ 1) - 1;$$

$$Frk(l ⊕ 1) := Frk(l ⊕ 1) - 1$$

is a valid implementation. Similarly, $DropFrks(l)$ requires that

```

glo    {Frk}
aux    {Num, Eating}
pre    Num(l) < Q ∧ Eating(l) = 1 ∧ Frk(l) = 2 ∧
        Frk(l ⊖ 1) < 2 ∧ Frk(l ⊕ 1) < 2 ∧ Inv
rely    I
wait    false
guar    true
eff    Num(l) < Q ∧ Eating(l) = 0 ∧ Frk(l) = 2 ∧
        Frk(l ⊖ 1) =  $\overline{Frk(l ⊖ 1) + 1}$  ∧
        Frk(l ⊕ 1) =  $\overline{Frk(l ⊕ 1) + 1} ∧ Inv$ ,

```

which is satisfied by

$$Frk(l ⊖ 1) := Frk(l ⊖ 1) + 1;$$

$$Frk(l ⊕ 1) := Frk(l ⊕ 1) + 1.$$

Chapter 13

Bubble-Lattice-Sort

13.1 Task

We will now employ LSP to develop Bubble-Lattice-Sort, a sorting algorithm which relies heavily upon synchronisation between $M + 1$ processes¹. A number of proof techniques are used to verify a related algorithm in [Bar85].

Let

$$D_{(n,m)} = \{n, n + 1, \dots, m\}.$$

Our job is to implement the following operation specification

```
Sort()
ioeh  B  :  array  $D_{(1,M)}$  to Value

    pre   true
    rely  true
    wait  false
    guar  true
    eff    $\forall j \in D_{(1,M-1)} \cdot B(j) \geq B(j+1) \wedge$ 
           $\forall v \in \textit{Value}.$ 
           $\{j \in D_{(1,M)} \mid \overleftarrow{B}(j) = v\} = \{j \in D_{(1,M)} \mid B(j) = v\}.$ 
```

The operation *Sort* sorts *B* into decreasing order, given that the environment

¹*M* is a constant — not a variable.

does not change the value of B (in a way observable to $Sort$) while $Sort$ is being executed². Moreover, $Sort$ does not rely upon any help from the environment to terminate.

13.2 Development

13.2.1 Division into Processes

There are of course many ways to implement this operation. As already indicated, we will use an algorithm called Bubble-Lattice-Sort³, which is related to the well-known sequential Bubble-Sort algorithm.

Bubble-Lattice-Sort consists of $M + 1$ processes:

$$\{Send() \parallel Bubble(1) \parallel \dots \parallel Bubble(M)\}.$$

$Send()$ reads the next value of B to be bubbled (from left to right) and passes it on to $Bubble(1)$, which, when it has received two values, will transfer the smallest to $Bubble(2)$ (or one of them if they are equal). This is thereafter repeated each time $Bubble(1)$ receives a value from $Send()$.

$Bubble(2)$ feeds $Bubble(3)$ in a similar style and so on. The algorithm terminates when (for the first time) $Bubble(M)$ receives a value from $Bubble(M-1)$.

The sending and receiving is synchronised in such a way that no process is ‘storing’ more than two values at a time.

13.2.2 More Data Structure

To implement the algorithm we will employ two global arrays

$$\begin{aligned} Feed &: \text{array } D_{(1,M+1)} \text{ to } Value, \\ Empty &: \text{array } D_{(1,M+1)} \text{ to } \mathbb{B}. \end{aligned}$$

²Observe that the rely-condition actually is ‘equivalent’ to $B = \overleftarrow{B}$ since the ioeh-declaration is constraining B from being changed (in an observable way) by the environment.

³The origin of this problem is a description given by T. C. Chen to C. B. Jones (in a *Heueriger!*) of how lattices of magnetic bubbles can be used to sort values in linear time. Chen’s claim, in 1975, was that there was no algebra in which this could be verified.

The following auxiliary array will also be useful

$Num : \text{array } D_{(0,M+1)} \text{ to } D_{(0,M)}.$

$Send()$ can only access the memory locations $Empty(1)$, $Feed(1)$, $Num(0)$ and B . $Bubble(l)$, on the other hand, is restricted to accessing $Empty(l)$, $Empty(l+1)$, $Feed(l)$, $Feed(l+1)$, $Num(l)$ and $B(l)$.

$Bubble(j)$ is ready to receive a value when $Empty(j)$ is true, otherwise $Empty(j)$ is false. When $Empty(j)$ is false and $Empty(j+1)$ is true, then no other process will try to access any of the memory locations accessible by $Bubble(j)$.

$Num(j)$ is equal to the number of values bubbled by $Bubble(j)$ (the number of values read by $Send()$ if $j = 0$), while $Feed(j)$ will be used by $Bubble(j-1)$ to pass on a new value to $Bubble(j)$ (from $Send()$ to $Bubble(1)$ if $j = 1$).

A value passed on to $Bubble(j)$ will be sent directly to $Bubble(j+1)$, if the value is less than or equal to $B(j)$, otherwise the value stored in $B(j)$ will be passed on to $Bubble(j+1)$, and the new value will be assigned to $B(j)$. Moreover, because the internal steps of one particular bubbling may be hidden completely inside an await-statement, we can insist on the truth of

$$\forall j \in D_{(1,M-1)} \cdot (Num(j+1) \neq 0 \Rightarrow B(j) \geq B(j+1)) \wedge (\neg Empty(j+1) \Rightarrow Feed(j+1) \leq B(j)).$$

In other words, this condition can be used as an invariant.

Moreover, to ensure that the different processes stay in step, we would also like

$$\begin{aligned} & (Empty(1) \Rightarrow Num(0) = Num(1)) \wedge \\ & (\neg Empty(1) \Rightarrow Num(0) = Num(1) + 1) \wedge \\ & \forall j \in D_{(1,M)} \cdot (\neg Empty(j+1) \Rightarrow Num(j+1) + 2 = Num(j)) \wedge \\ & (Empty(j+1) \Rightarrow Num(j+1) + 1 = Num(j) \vee \\ & \quad Num(j) = Num(j+1) = 0) \end{aligned}$$

to be an invariant. From now on, we will use Inv to denote the conjunction of these two conditions.

The reason why $M+1$ occurs in the domains of the new arrays, is that we

then avoid making a special case of implementing *Bubble*(M). Furthermore, this also has a simplifying effect on some of the proofs.

13.2.3 Dynamic Invariant

The next step is to characterise the way in which the $M + 1$ processes may interact. Since any externally observable state change by any of the processes is assumed to maintain the invariant, this has to some extent already been done. What remains is to make sure that the bag of values occurring in B is preserved. The condition below is sufficient:

$$\begin{aligned} \forall v \in Value. \\ \{j \in D_{(1,M)} \mid \overleftarrow{B}(j) = v \wedge \overleftarrow{Num}(j) > 0\} + \\ \{j \in D_{(1,M)} \mid \overleftarrow{Feed}(j) = v \wedge \neg \overleftarrow{Empty}(j)\} + \\ \{j \in D_{(1,M)} \mid \overleftarrow{B}(j) = v \wedge j > \overleftarrow{Num}(0)\} = \\ \{j \in D_{(1,M)} \mid B(j) = v \wedge Num(j) > 0\} + \\ \{j \in D_{(1,M)} \mid Feed(j) = v \wedge \neg Empty(j)\} + \\ \{j \in D_{(1,M)} \mid B(j) = v \wedge j > Num(0)\}. \end{aligned}$$

The reason why this constraint is not included in the invariant *Inv* is of course that it refers to a previous state. Instead, it can be defined as a dynamic invariant, and we will use *Dyn* to denote this formula. In other words, in any specification where we need to state the condition above, we will instead write *Dyn*.

13.2.4 Main Structure

This means that *Sort* will be of the form:

```

{true}

loc Feed : array  $D_{(1,M+1)}$  to Value;
loc Empty : array  $D_{(1,M+1)}$  to  $\mathbb{B}$ ;
Init();

{Num(0) = 0  $\wedge$  Inv  $\wedge$  Dyn}

{Send()  $\parallel$  Bubble(1)  $\parallel$  ...  $\parallel$  Bubble(M)}

{Num(0) = M  $\wedge$  ( $\forall j \in D_{(1,M+1)} \cdot$  Num(j) = M - j + 1)  $\wedge$ 
  Inv  $\wedge$  Dyn}.

```

Observe that

$$\vdash \text{Num}(0) = 0 \wedge \text{Inv} \Rightarrow \forall j \in D_{(1,M)} \cdot \text{Num}(j) = 0 \wedge \text{Empty}(j).$$

13.2.5 Specification of Processes

Based on the discussion above, it is now straightforward to specify the component processes:

```

Send()
glo  ioeo  Empty(1) :  $\mathbb{B}$ ,
      ioeh  Feed(1)  : Value,
aux  ioeh  Num(0)   :  $D_{(0,M)}$ 

pre   Num(0) = 0  $\wedge$  Inv
rely  Dyn  $\wedge$  ( $\overline{\text{Inv}} \Rightarrow \text{Inv}$ )
wait  Num(0) < M  $\wedge$   $\neg \text{Empty}(1) \wedge$  Inv
guar  Dyn  $\wedge$  ( $\overline{\text{Inv}} \Rightarrow \text{Inv}$ )
eff   Num(0) = M  $\wedge$  Dyn  $\wedge$  Inv,

```

$Bubble(l : D_{(1,M)})$
 glo ioeo $Empty(l) : \mathbb{B},$
 $Empty(l+1) : \mathbb{B},$
 i heo $Feed(l) : Value,$
 ioeh $Feed(l+1) : Value,$
 $B(l) : Value,$
 aux ioeh $Num(l) : D_{(0,M)}$

 pre $Num(l) = 0 \wedge Inv$
 rely $Dyn \wedge \overline{Inv} \Rightarrow Inv$
 wait $Num(l) < M - l + 1 \wedge (Empty(l) \vee \neg Empty(l+1)) \wedge Inv$
 guar $Dyn \wedge \overline{Inv} \Rightarrow Inv$
 eff $Num(l) = M - l + 1 \wedge Dyn \wedge Inv,$

13.2.6 Correctness Proof

The next step is to prove that the parallel composition of these $M + 1$ processes has the desired effect. Firstly, observe that the constraints imposed on the global memory locations by the declarations in the respective operations are satisfied. Secondly, since

$$\begin{aligned}
 &\vdash Num(0) < M \wedge \neg Empty(1) \wedge Inv \wedge \\
 &\quad ((Num(1) = M \wedge Dyn \wedge Inv) \vee \\
 &\quad (Num(1) < M \wedge (Empty(1) \vee \neg Empty(2)) \wedge Inv)) \Rightarrow \\
 &\quad Num(1) < M \wedge \neg Empty(2) \wedge Inv
 \end{aligned}$$

$$\begin{aligned}
 &\vdash \forall l \in D_{(1,M-1)} \cdot Num(l) < M - l + 1 \wedge \neg Empty(l+1) \wedge Inv \wedge \\
 &\quad ((Num(l+1) = M - l \wedge Dyn \wedge Inv) \vee \\
 &\quad (Num(l+1) < M - l \wedge (Empty(l+1) \vee \neg Empty(l+2)) \wedge Inv)) \Rightarrow \\
 &\quad Num(l+1) < M - l \wedge \neg Empty(l+2) \wedge Inv
 \end{aligned}$$

$$\vdash \neg (Num(M) < 1 \wedge \neg Empty(M+1) \wedge Inv),$$

it follows that whenever $Send()$ becomes blocked, there is at least one other

process which is enabled. In a similar way it can be shown that for all $1 \leq l \leq M$, $Bubble(l)$ cannot become blocked in a state in which no other process is enabled. Thus, it follows by the consequence- and generalised parallel-rules that the parallel-statement satisfies:

glo	$\{B, Empty, Feed\}$
aux	$\{Num\}$
pre	$Num(0) = 0 \wedge Inv$
rely	I
wait	false
guar	$\overleftarrow{Dyn} \wedge (\overleftarrow{Inv} \Rightarrow Inv)$
eff	$\forall j \in D_{(1,M)} \cdot Num(j) = M - j + 1 \wedge Num(0) = M \wedge \overleftarrow{Dyn} \wedge Inv.$

But then, since it is clear that

$$\begin{aligned}
& \vdash \overleftarrow{Num}(0) = 0 \wedge \overleftarrow{Inv} \wedge \\
& \quad \forall j \in \{1, \dots, M\} \cdot Num(j) = M - j + 1 \wedge \\
& \quad \quad Num(0) = M \wedge \overleftarrow{Dyn} \wedge Inv \Rightarrow \\
& \quad \quad \forall j \in D_{(1,M-1)} \cdot B(j) \geq B(j+1) \wedge \\
& \quad \quad \forall v \in Value. \\
& \quad \quad \{j \in D_{(1,M)} \mid \overleftarrow{B}(j) = v\} = \{j \in D_{(1,M)} \mid B(j) = v\},
\end{aligned}$$

we have proved that the parallel composition of the $M + 1$ processes has the desired effect. Moreover, the proof is only relying upon assumptions about the environment allowed by the specification of *Sort*.

13.2.7 Process Decomposition

The final step is to implement the specified operations. As should be pretty obvious by now, we will have to rely upon the await-statement as an abstraction tool. Let us postpone implementation of the await-statement's bodies, and first design the synchronisation structure.

Since $Send()$ is supposed to read M different values in B , and $Bubble(1)$ is ready to receive a value when $Empty(1)$ is true, it follows that $Send()$ should be of the form:

```

{Num(0) = 0 ∧ Inv}

begin
  loc j;
  j := 0;
  while j < M do

    {Num(0) < M ∧ Inv ∧ Dyn}

    j := j + 1;
    await Empty(1) do S(j) od
  od
end

{Num(0) = M ∧ Inv ∧ Dyn}.

```

Moreover, since $Bubble(l)$ will receive $M - l + 1$ values, $Empty(l)$ is false if and only if an unread value is stored in $Feed(l)$, and $Bubble(l+1)$ is ready to receive a new value if $Empty(l+1)$ is true, it follows that $Bubble(l)$ should be of the form:


```

{Num(l) = 0 ∧ Inv}

await ¬Empty(l) do Fb(l) od;
begin
  loc j;
  j := 1;
  while j < M - l + 1 do

    {Num(l) < M - l + 1 ∧ Inv ∧ Dyn}

    await ¬Empty(l) ∧ Empty(l + 1) do Rb(l) od;
    j := j + 1
  od
end

{Num(l) = M - l + 1 ∧ Inv ∧ Dyn}.

```

13.2.8 Implementing Atomic Statements

To finish the implementation of $Send()$, it is necessary to substitute code for $S(j)$ and verify that the final product has the desired properties. The latter follows easily if $S(j)$ satisfies:

```

glo    {B, Empty, Feed}
aux    {Num}
pre    Num(0) < M ∧ Empty(1) ∧ Inv
rely    I
wait    false
guar    true
eff    Num(0) =  $\overleftarrow{Num}(0) + 1 \wedge \neg Empty(1) \wedge Inv \wedge Dyn$ .

```

Thus,

```

Feed(1) := B(j);
Empty(1) := false

```

is all that is required.

Similarly, to finish the code of $Bubble(l)$ we must implement $Fb(l)$ and $Rb(l)$ and make sure the specified properties are satisfied. Again, if $Fb(l)$ satisfies:

```

glo    {B, Empty, Feed}
aux    {Num}
pre    Num(l) = 0  $\wedge$   $\neg$ Empty(l)  $\wedge$  Inv
rely    I
wait    false
guar    true
eff    Num(l) = 1  $\wedge$  Empty(l)  $\wedge$  Inv  $\wedge$  Dyn,

```

and $Rb(l)$ satisfies:

```

glo    {B, Empty, Feed}
aux    {Num}
pre    Num(l) < M - l + 1  $\wedge$   $\neg$ Empty(l)  $\wedge$  Empty(l + 1)  $\wedge$  Inv
rely    I
wait    false
guar    true
eff    Num(l) =  $\overleftarrow{Num}(l)$  + 1  $\wedge$  Empty(l)  $\wedge$   $\neg$ Empty(l + 1)  $\wedge$ 
      Inv  $\wedge$  Dyn,

```

it is clear that the implementation is correct, so

```

B(l) := Feed(l);
Empty(l) := true,

```

and

```

Feed(l + 1) := min {Feed(l), B(l)};
B(l) := max {Feed(l), B(l)};
Empty(l) := true;
Empty(l + 1) := false

```

are what is missing.

Chapter 14

Set-Partition

14.1 Task

Given two non-empty, disjoint sets of integers, S and L ; our task is to develop a program which terminates in a state where the maximum element of S is less than the minimum element of L . The sizes of the two sets must remain unchanged. Moreover, after termination, the union of S and L is required to equal the union of their initial values.

This informal specification can be translated into a more mathematical notation:

```
SetPart()
glo  ioeh   $S$  : set of  $\mathbb{N}$ ,
         $L$  : set of  $\mathbb{N}$ 
aux

pre   $S \cap L = \{\} \wedge S \neq \{\} \wedge L \neq \{\}$ 
rely true
wait false
guar true
eff   $\max(S) \leq \min(L) \wedge \#S = \# \overline{S} \wedge$ 
       $\#L = \# \overline{L} \wedge S \cup L = \overline{S} \cup \overline{L}.$ 
```

The declarations of S and L allow us to assume that the environment will leave S and L unchanged. Moreover, from the wait-condition it follows that

SetPart is required to terminate. Finally, the guar-condition allows us to change S and L as we like. The rest should be clear from the informal specification.

14.2 Development

14.2.1 Algorithm

Our implementation is inspired by [Dij82]. In [Bar85] a number of methods are used to verify a related algorithm¹.

The algorithm employs two processes called respectively *Small* and *Large*. The basic idea is as follows:

- The process *Small* starts by finding the maximum element of S . This integer is sent on to *Large* and then subtracted from S .

The task of *Large* is to add the received integer to L , and thereafter send the minimum element of L (which by then contains the integer just received from *Small*) back to *Small* and remove it from L .

The process *Small* adds the element sent from *Large* to S . Then, if the maximum of S equals the integer just received from *Large*, it follows that the maximum of S is less than the minimum of L and the algorithm terminates. Otherwise, the whole procedure is repeated.

Since the difference between the maximum of S and the minimum of L is decreased at each iteration, it follows that the program will eventually terminate.

¹The algorithm used in [Bar85], where verification methods for both CSP and shared-state programs are compared, is understandably a direct translation from CSP. It may be argued that to make it easier to compare our approach with those discussed by Barringer, we should have developed the same algorithm.

However, as pointed out in [Bar85], the shared state version may deadlock, and since we want to prove total correctness, a modification is necessary. Moreover, the algorithm in [Bar85] employs two flags, one for each CSP-channel. We decided to replace these two flags with one, because this results in a simpler and more natural development and therefore gives a more correct impression of what can be achieved in LSP with respect to this particular synchronisation problem.

14.2.2 Data Structure

The variables *Max* and *Min* simulate respectively ‘the channel’ from *Small* to *Large* and ‘the channel’ from *Large* to *Small*.

To secure that the two processes stay in step, the Boolean variable *Flag* is introduced. When *Small* switches on *Flag*, it means that *Large* may read the next value from *Max*, and when *Large* makes *Flag* false, it signals that *Min* is ready to be read by *Small*.

The adding, finding the maximum and sending section of *Small* is mutually exclusive with the adding, finding the minimum and sending section of *Large*.

The only thing the process *Small* is allowed to do while *Flag* is true, is to remove from *S* the integer it just sent to *Large*. Similarly, when *Flag* is false, *Large* is only allowed to remove the element it just sent to *Small*.

14.2.3 First Decomposition Step

Our implementation will be of the form:

```
begin
  loc Max, Min, Flag;
  Init();
  {Small() || Large()}
end.
```

The task of *Init* is of course to initialise the local state.

14.2.4 Specifying *Init*

To make it easier to formulate and reason about properties satisfied by the concurrent part of our implementation, we will use *Init* to simulate the first iteration of the algorithm; in other words, to perform the first interchange of values. This means that:

```

Init()
glo  ioeh      S  : set of  $\mathbb{N}$ ,
                L  : set of  $\mathbb{N}$ ,
                Min :  $\mathbb{N}$ ,
                Max :  $\mathbb{N}$ ,
                Flag :  $\mathbb{B}$ 

aux

pre  S  $\cap$  L =  $\{\}$   $\wedge$  S  $\neq$   $\{\}$   $\wedge$  L  $\neq$   $\{\}$ 
rely true
wait false
guar true
eff  #S = # $\overleftarrow{S}$   $\wedge$ 
      #L = # $\overleftarrow{L}$   $\wedge$ 
      S  $\cup$  L  $\cup$  {Max, Min} =  $\overleftarrow{S}$   $\cup$   $\overleftarrow{L}$   $\wedge$ 
       $\neg$ Flag  $\wedge$  Min  $\leq$  Max  $\wedge$ 
      Max = max(S)  $\wedge$  Min < min(L)  $\wedge$ 
      S  $\cap$  L =  $\{\}$ .

```

Basically, this operation simulates ‘the sending’ of one element in both directions. Thus, the next process to transfer a value is *Small*, which explains the restriction on *Flag*. Moreover, *Small* has already determined the ‘new’ maximum of *S*.

14.2.5 Final Implementation of *Init*

The implementation of *Init* is not very challenging. The program below is obviously sufficient:

```

Max := max(S);
L := L  $\cup$  {Max};
S := S - {Max};
Min := min(L);
S := S  $\cup$  {Min};
L := L - {Min};
Flag := false;
Max := max(S).

```

14.2.6 Invariant

We will now characterise a few properties that will be invariantly true for the concurrent part of the implementation. Since for both processes the previously sent element is removed before the actual process starts to look for a new integer to send, it is clear that:

$$S \cap L \subseteq \{Min, Max\}.$$

Moreover, because *Large* will return the integer just received if the maximum of *S* is less than the minimum of *L*, it is also true that:

$$Min \leq Max.$$

We will use *Inv* to denote the conjunction of these two assertions. *Inv* is obviously implied by the eff-condition of *Init*.

14.2.7 Dynamic Invariant

To ensure maintenance of the original integers we will require that any state transition must satisfy:

$$S \cup L \cup \{Max, Min\} = \overleftarrow{S} \cup \overleftarrow{L} \cup \{\overleftarrow{Max}, \overleftarrow{Min}\}.$$

This is of course not enough on its own; however, if we insist that the conjunction of the eff-conditions of the two processes implies that $\{Max, Min\} \subseteq S \cup L$, it follows easily from the eff-condition of *Init* that the desired maintenance property is satisfied by the overall program.

Moreover, since the first interchange of elements has already taken place in *Init*, it is clear that any transition by either *Small* or *Large* will satisfy:

$$Max \leq \overleftarrow{Max} \wedge Min \geq \overleftarrow{Min}.$$

To prove freedom from deadlock an auxiliary Boolean variable *TrmS* is needed. The idea is that *TrmS* is switched on when *Small* leaves its critical section for the last time, i.e. in the case that *Max* equals *Min*. To show that *Flag* is true when *Small* has terminated, and that *TrmS* is true when *Large* has terminated, we must insist that:

$$\overline{\dot{TrmS}} \wedge \overline{\dot{Flag}} \Rightarrow TrmS \wedge Flag.$$

From now on we will use *Dyn* to denote the conjunction of these three assertions.

14.2.8 Freedom from Deadlock

The process *Small* can only become blocked if it wants to enter its critical section, i.e. if *Flag* is true and *TrmS* is still false. Hence, *Small* will only wait in a state which satisfies:

$$Flag \wedge \neg TrmS \wedge Inv.$$

Similarly, it is clear that *Large* will only be held back in a state characterised by:

$$\neg Flag \wedge Inv.$$

The conjunction of these two assertions is obviously inconsistent. Moreover, if we insist that the eff-condition of *Small* implies that *Flag* is true, while the eff-condition of *Large* implies that *TrmS* is switched on, it follows that the parallel composition of *Small* and *Large* will never deadlock.

14.2.9 Specifying *Small*

From the discussion above it is clear that *Small* does not need write access to *L* and *Min*. Similarly, *Large* will never have to change the value of *S*, *Max* and *TrmS*.

To secure mutual exclusion the environment must maintain the falsity of *Flag*, while *Small* in return must guarantee never to make *Flag* false.

Moreover, the only possible change of state due to the environment while *Flag* is false is that *Min* is removed from *L*. Similarly, the only thing *Small* is allowed to do while *Flag* is true is to remove *Max* from *S*.

Furthermore, to prove that the number of elements in *S*, when *Small* terminates, equals the set's initial size, any internal transition must satisfy:

$$\overleftarrow{\neg Flag} \wedge Flag \Rightarrow Max = Min \vee (Max \notin L \wedge Min \in S),$$

and for similar reasons when an integer is sent in the other direction, it is necessary that any external transition satisfies:

$$\overleftarrow{Flag} \wedge \neg Flag \Rightarrow Max = Min \vee (Min \notin S \wedge Max \in L).$$

Finally, to ensure that *TrmS* is switched on if and only if *Small* has left its critical section for the last time, any internal transition must satisfy:

$$\overleftarrow{\neg TrmS} \wedge TrmS \Leftrightarrow \overleftarrow{\neg Flag} \wedge Flag \wedge Max = Min.$$

Thus, in a more formal notation:

Small()

glo ioeo $Flag$: $\mathbb{B}$,
 ioeh S : set of $\mathbb{N}$,
 Max : $\mathbb{N}$,
 iheo L : set of $\mathbb{N}$,
 Min : $\mathbb{N}$

aux ioeh $TrmS$: $\mathbb{B}$

pre $Max = \max(S) \wedge \neg Flag \wedge Max \notin L \wedge$
 $\neg TrmS \wedge Inv$
 rely $(\neg \overline{Flag} \Rightarrow \neg Flag \wedge Min = \overline{Min} \wedge (L \cup \{Min\} = \overline{L} \vee L = \overline{L})) \wedge$
 $(\overline{Flag} \wedge \neg Flag \Rightarrow Max = Min \vee (Min \notin S \wedge Max \in L)) \wedge$
 $Dyn \wedge (\overline{Inv} \Rightarrow Inv)$
 wait $\overline{Flag} \wedge \neg TrmS \wedge Inv$
 guar $(\overline{Flag} \Rightarrow Flag \wedge Max = \overline{Max} \wedge (S \cup \{Max\} = \overline{S} \vee S = \overline{S})) \wedge$
 $(\neg \overline{Flag} \wedge Flag \Rightarrow Max = Min \vee (Max \notin L \wedge Min \in S)) \wedge$
 $(\neg TrmS \wedge TrmS \Leftrightarrow \neg \overline{Flag} \wedge Flag \wedge Max = Min) \wedge$
 $Dyn \wedge (\overline{Inv} \Rightarrow Inv)$
 eff $\#S = \#\overline{S} \wedge Max = \max(S) \wedge$
 $Max = Min \wedge Flag \wedge$
 $Dyn \wedge Inv.$

14.2.10 Specifying *Large*

The specification of *Large* is very similar:

Large()

glo ioeo $Flag : \mathbb{B}$,
 ioeh $L : \text{set of } \mathbb{N}$,
 $Min : \mathbb{N}$,
 iheo $S : \text{set of } \mathbb{N}$,
 $Max : \mathbb{N}$

aux iheo $TrmS : \mathbb{B}$

pre $\overline{Min} < \min(L) \wedge Max \notin L \wedge \neg TrmS \wedge \neg Flag \wedge Inv$
 rely $(\overline{Flag} \Rightarrow Flag \wedge Max = \overline{Max} \wedge (S \cup \{Max\} = \overline{S} \vee S = \overline{S})) \wedge$
 $(\neg \overline{Flag} \wedge Flag \Rightarrow Max = Min \vee (Max \notin L \wedge Min \in S)) \wedge$
 $(\neg \overline{TrmS} \wedge TrmS \Leftrightarrow \neg \overline{Flag} \wedge Flag \wedge Max = Min) \wedge$
 $Dyn \wedge (\overline{Inv} \Rightarrow Inv)$
 wait $\neg \overline{Flag} \wedge Inv$
 guar $(\overline{\neg Flag} \Rightarrow \neg Flag \wedge Min = \overline{Min} \wedge (L \cup \{Min\} = \overline{L} \vee L = \overline{L})) \wedge$
 $(\overline{Flag} \wedge \neg \overline{Flag} \Rightarrow Max = Min \vee (Min \notin S \wedge Max \in L)) \wedge$
 $Dyn \wedge (\overline{Inv} \Rightarrow Inv)$
 eff $\#L = \#\overline{L} \wedge Min < \min(L) \wedge$
 $Max = Min \wedge TrmS \wedge Dyn \wedge Inv.$

14.2.11 Parallel Composition

Since the wait-conditions of both operations are inconsistent with the other operation's wait- and eff-conditions, it follows by the parallel- and consequence-rules that the concurrent part of our implementation satisfies:

```

glo    {L, S, Flag, Min, Max}
aux    {TrmS}
pre    Max = max(S) ∧ ¬Flag ∧ ¬TrmS ∧ Max ∉ L ∧
       Min < min(L) ∧ Inv
rely    I
wait    false
guar    true
eff     #L = # $\overline{L}$  ∧ Min < min(L) ∧
       #S = # $\overline{S}$  ∧ Max = max(S) ∧
       Max = Min ∧ Dyn ∧ Inv,

```

which together with *Init* gives the desired overall effect.

14.2.12 Decomposing *Small*

How can we best decompose *Small*? Obviously, the while-construct is needed. One possible strategy is the following:

```

begin
  loc VS;
  VS := (Max ≠ Min);
  while VS do
    Sml();
    VS := (Max ≠ Min)
  od;
  Flag := true
end.

```

The obvious termination expression is:

$$Max - Min.$$

Thus, since

$$\begin{aligned} \overline{Min} < \overline{Max} \wedge (Max < \overline{Max} \vee Max = Min) \wedge \\ Dyn \Rightarrow Max - Min < \overline{Max} - \overline{Min}, \end{aligned}$$

and

$$Inv \Rightarrow Max - Min \geq 0,$$

it follows that the loop terminates and that the specification of *Small* is satisfied, if we can prove that *Sml* is characterised by:

$$\begin{array}{ll}
Sml() & \\
\text{glo} \quad \text{ioeo} \quad Flag & : \mathbb{B}, \\
\quad \text{ioeh} \quad S & : \text{set of } \mathbb{N}, \\
\quad \quad Max & : \mathbb{N}, \\
\quad \text{iheo} \quad L & : \text{set of } \mathbb{N}, \\
\quad \quad Min & : \mathbb{N} \\
\\
\text{aux} \quad \text{ioeh} \quad TrmS & : \mathbb{B} \\
\\
\text{pre} \quad Max = max(S) \wedge \neg Flag \wedge & \\
\quad \neg TrmS \wedge Min < Max \wedge Max \notin L \wedge Inv & \\
\text{rely} \quad (\neg Flag \Rightarrow \neg Flag \wedge Min = \overline{Min} \wedge (L \cup \{Min\} = \overline{L} \vee L = \overline{L})) \wedge & \\
\quad (Flag \wedge \neg Flag \Rightarrow Max = Min \vee (Min \notin S \wedge Max \in L)) \wedge & \\
\quad Dyn \wedge (Inv \Rightarrow Inv) & \\
\text{wait} \quad Flag \wedge \neg TrmS \wedge Inv & \\
\text{guar} \quad (Flag \Rightarrow Flag \wedge Max = \overline{Max} \wedge (S \cup \{Max\} = \overline{S} \vee S = \overline{S})) \wedge & \\
\quad (\neg Flag \wedge Flag \Rightarrow Max = Min \vee (Max \notin L \wedge Min \in S)) \wedge & \\
\quad (\neg TrmS \wedge TrmS \Leftrightarrow \neg Flag \wedge Flag \wedge Max = Min) \wedge & \\
\quad Dyn \wedge (Inv \Rightarrow Inv) & \\
\text{eff} \quad \#S = \#\overline{S} \wedge Max = max(S) \wedge & \\
\quad \neg Flag \wedge \neg TrmS \wedge & \\
\quad ((Max < \overline{Max} \wedge Max \notin L) \vee Max = Min) \wedge & \\
\quad Dyn \wedge Inv. &
\end{array}$$

14.2.13 Final Implementation of *Sml*

It can be shown that this property is satisfied by the program below:

$$\{Max = \max(S) \wedge Min < Max \wedge Max \notin L \wedge \\ \neg TrmS \wedge \neg Flag \wedge Dyn \wedge Inv\}$$

Flag := true;

$$\{Max = \max(S) \wedge \#S = \# \overleftarrow{S} \wedge Max = \overleftarrow{Max} \wedge \neg TrmS \wedge \\ (\neg Flag \Rightarrow Max = Min \vee (Min \notin S \wedge Max \in L)) \wedge Dyn \wedge Inv\}$$

S := *S* \ {*Max*};

$$\{(S \neq \{\}) \Rightarrow Max > \max(S)) \wedge Max = \overleftarrow{Max} \wedge \#S = \# \overleftarrow{S} - 1 \wedge \\ (\neg Flag \Rightarrow Max = Min \vee (Min \notin S \wedge Max \in L)) \wedge \\ \neg TrmS \wedge Dyn \wedge Inv\}$$

await $\neg Flag$ do
 skip
od;

$$\{(S \neq \{\}) \Rightarrow Max > \max(S)) \wedge \#S = \# \overleftarrow{S} - 1 \wedge \\ (Max = Min \vee (Min \notin S \wedge Max \in L)) \wedge \\ Max = \overleftarrow{Max} \wedge \neg TrmS \wedge Dyn \wedge \neg Flag \wedge Inv\}$$

S := *S* ∪ {*Min*};

$$\{((Max > \max(S) \wedge Max \in L) \vee (Max = \max(S) \wedge Max = Min)) \wedge \\ \#S = \# \overleftarrow{S} \wedge Max = \overleftarrow{Max} \wedge \\ Min \in S \wedge \neg TrmS \wedge Dyn \wedge \neg Flag \wedge Inv\}$$

Max := *max*(*S*);

$$\{Max = \max(S) \wedge ((Max < \overleftarrow{Max} \wedge Max \notin L) \vee Max = Min) \wedge \\ \#S = \# \overleftarrow{S} \wedge \neg TrmS \wedge Dyn \wedge \neg Flag \wedge Inv\}.$$

14.2.14 Decomposing *Large*

What remains is to decompose *Large*. The main structure is given below:

```

begin
  loc  $V_L$ ;
  await Flag do
    skip
  od;
   $V_L := (Max \neq Min)$ ;
  while  $V_L$  do
    Lrg();
     $V_L := (Max \neq Min)$ 
  od
end.

```

Again, the termination expression is

$$Max - Min.$$

Moreover, since

$$\overline{\overline{Min}} < \overline{\overline{Max}} \wedge Min > \overline{\overline{Min}} \wedge Dyn \Rightarrow Max - Min < \overline{\overline{Max}} - \overline{\overline{Min}},$$

and

$$Inv \Rightarrow Max - Min \geq 0,$$

it is enough to show that *Lrg* satisfies:

$Lrg()$

glo ioeo $Flag$: $\mathbb{B}$,
 ioeh L : set of $\mathbb{N}$,
 Min : $\mathbb{N}$,
 iheo S : set of $\mathbb{N}$,
 Max : $\mathbb{N}$

aux iheo $TrmS$: $\mathbb{B}$

pre $Min < \min(L) \wedge Max \notin L \wedge Min \in S \wedge Flag \wedge$
 $\neg TrmS \wedge Min < Max \wedge Inv$
 rely $(\overline{Flag} \Rightarrow Flag \wedge Max = \overline{Max} \wedge (S \cup \{Max\} = \overline{S} \vee S = \overline{S})) \wedge$
 $(\neg \overline{Flag} \wedge Flag \Rightarrow Max = Min \vee (Max \notin L \wedge Min \in S)) \wedge$
 $(\neg TrmS \wedge TrmS \Leftrightarrow \neg \overline{Flag} \wedge Flag \wedge Max = Min) \wedge$
 $Dyn \wedge (\overline{Inv} \Rightarrow Inv)$
 wait $\neg \overline{Flag} \wedge Inv$
 guar $(\neg \overline{Flag} \Rightarrow \neg Flag \wedge Min = \overline{Min} \wedge (L \cup \{Min\} = \overline{L} \vee L = \overline{L})) \wedge$
 $(Flag \wedge \neg Flag \Rightarrow Max = Min \vee (Min \notin S \wedge Max \in L)) \wedge$
 $Dyn \wedge (\overline{Inv} \Rightarrow Inv)$
 eff $\#L = \#\overline{L} \wedge Min < \min(L) \wedge Flag \wedge Min > \overline{Min} \wedge$
 $(Max = Min \vee Max \notin L) \wedge (Max = Min \Leftrightarrow TrmS) \wedge$
 $Dyn \wedge Inv.$

14.2.15 Final Implementation of Lrg

$$\{Min < \min(L) \wedge Max \notin L \wedge Min \in S \wedge Min < Max \wedge \\ Flag \wedge \neg TrmS \wedge Dyn \wedge Inv\}$$

$$L := L \cup \{Max\};$$

$$\{\#L = \# \overleftarrow{L} + 1 \wedge Flag \wedge \neg TrmS \wedge Min = \overleftarrow{Min} \wedge Min \in S \wedge \\ Min < \min(L) \wedge Max \in L \wedge Dyn \wedge Inv\}$$

$$Min := \min(L);$$

$$\{Min = \min(L) \wedge \#L = \# \overleftarrow{L} + 1 \wedge Flag \wedge \neg TrmS \wedge \\ (Max = Min \vee Min \notin S) \wedge Min > \overleftarrow{Min} \wedge Dyn \wedge Inv\}$$

$$Flag := \text{false};$$

$$\{Min = \min(L) \wedge \#L = \# \overleftarrow{L} + 1 \wedge \\ Min > \overleftarrow{Min} \wedge (Flag \Rightarrow (Max = Min \Leftrightarrow TermS)) \wedge \\ (Flag \Rightarrow Max = Min \vee (Max \notin L \wedge Min \in S)) \wedge Dyn \wedge Inv\}$$

$$L := L - \{Min\};$$

$$\{Min < \min(L) \wedge \#L = \# \overleftarrow{L} \wedge \\ Min > \overleftarrow{Min} \wedge (Flag \Rightarrow (Max = Min \Leftrightarrow TermS)) \wedge \\ (Flag \Rightarrow Max = Min \vee (Max \notin L \wedge Min \in S)) \wedge Dyn \wedge Inv\}$$

await $Flag$ do

skip

od;

$$\{Min < \min(L) \wedge \#L = \# \overleftarrow{L} \wedge \\ Min > \overleftarrow{Min} \wedge Flag \wedge (Max = Min \vee (Max \notin L \wedge Min \in S)) \wedge \\ (Max = Min \Leftrightarrow TrmS) \wedge Dyn \wedge Inv\}.$$

Chapter 15

Modified System

15.1 Motivation

LSP can only be used to develop programs which are intended to converge. There is no obvious way the system can be modified to handle more general liveness (see [AS85]) constraints in a similar style. However, if we restrict ourselves to the following four characteristics:

- the overall effect — if the implementation terminates,
- the effect of any atomic step due to the implementation,
- the set of states in which the implementation can become blocked,
- that the body of an await-statement terminates whenever it is executed,

LSP can easily be adapted. LSP_S , which is what the new system is called, depends upon LSP to prove termination of await-bodies. The rest of the system is described below.

15.2 Modifications

15.2.1 Specified Programs

A specified program is of the form

$$z \text{ sat } [\vartheta, \alpha] :: [P, R, W, G, E].$$

Syntactically, the only difference from above is that square brackets are used instead of curly brackets. Both specifications and specified programs are required to satisfy the same constraints as earlier. However, their interpretations are different.

15.2.2 Assumptions

The pre-condition is still assumed to denote a set of initial states to which the implementation is applicable, while the rely-condition as before is supposed to characterise any uninterrupted state transition by the environment.

Thus, the pre- and rely-conditions constitute assumptions which the developer can make about the environment.

15.2.3 Commitments

Any state in which the implementation can become blocked is required to satisfy the wait-condition. However, the implementation is not allowed to become blocked inside the body of an await-statement. Moreover, any internal transition is constrained to satisfy the guar-condition, while the overall effect, if the implementation terminates, must satisfy the eff-condition.

15.2.4 Satisfaction

This means that the definition of satisfaction (see page 74) can be carried over from earlier, given that ext_π and int_π (see pages 68 and 69) are redefined as below:

Definition 18 *Given a pre-condition P , a rely-condition R , and a structure π , then $ext_\pi[P, R]$ denotes the set of all computations σ in π , such that:*

- $\delta(\sigma_1) \models_\pi P$,
- for all $1 \leq j < len(\sigma)$, if $\lambda(\sigma_j) = e$ then $(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_\pi R$.

Definition 19 *Given a wait-condition W , a guar-condition G , an eff-condition E , and a structure π , then $int_\pi[W, G, E]$ denotes the set of all computations σ in π , such that:*

- for all $1 \leq j \leq len(\sigma)$, if σ_j is blocked then $\delta(\sigma_j) \models_\pi W$,

- for all $1 \leq j < \text{len}(\sigma)$, if $\lambda(\sigma_j) = i$ then $(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_{\pi} G$ and $\tau(\sigma_j) \neq \tau(\sigma_{j+1})$,
- for all $1 \leq j \leq \text{len}(\sigma)$, if $\tau(\sigma_j) = \epsilon$ then $(\delta(\sigma_1), \delta(\sigma_j)) \models_{\pi} E$.

The second conjunct of *int*'s second condition restricts the bodies of await-statements to terminate. (Remember that the only internal transition which leaves the program component unchanged is the one which models that the execution of an await-statement's body either ends in an infinite loop or becomes blocked. See page 42.)

15.2.5 Decomposition-Rules

With two exceptions the decomposition-rules are identical to the decomposition-rules for LSP (although it is of course necessary to substitute square brackets for curly brackets). The first exception is the while-rule. Since the statement is no longer required to terminate, the first premise may be removed:

$$\frac{z \text{ sat } [\vartheta, \alpha] :: [P \wedge b, R, W, G, P \wedge Z]}{\text{while } b \text{ do } z \text{ od } \text{ sat } [\vartheta, \alpha] :: [P, R, W, G, (Z^{\dagger} \vee R) \wedge \neg b]}$$

Secondly, since it is necessary to use LSP to prove total correctness of the await-statement's body, the await-rule is of the form:

$$\frac{\begin{array}{l} P^R \wedge \neg b \Rightarrow W \\ E_1 | (\bigwedge_{a \in \alpha} a = \overleftarrow{u}_a \wedge I_{\alpha}) \Rightarrow G \wedge E_2 \\ z \text{ sat } (\vartheta, \alpha) :: (P^R \wedge b, I, \text{false}, \text{true}, E_1) \end{array}}{\text{await } b \text{ do } z \text{ od } \text{ sat } [\vartheta, \alpha] :: [P, R, W, G, R | E_2 | R]}$$

where for all $a \in \alpha$, $\text{var}[u_a] \subseteq \vartheta \cup \{a\}$

15.2.6 Soundness and Completeness

LSP_S is sound and satisfies that same relative-completeness criterion as LSP. The proofs are straightforward modifications of the proofs for LSP.

15.3 Advantages

15.3.1 Always Enabled

LSP_S allows us to prove that a program is always enabled. If for example

$$\begin{aligned} \vdash_S z_1 \underline{\text{sat}} [\vartheta, \alpha] &:: [\text{true}, \text{true}, b, \text{true}, \text{false}], \\ \vdash_S z_2 \underline{\text{sat}} [\vartheta, \alpha] &:: [\text{true}, \text{true}, \neg b, \text{true}, \text{false}], \end{aligned}$$

we may deduce that

$$\vdash_S \{z_1 \parallel z_2\} \underline{\text{sat}} [\vartheta, \alpha] :: [\text{true}, \text{true}, \text{false}, \text{true}, \text{false}],$$

which means that $\{z_1 \parallel z_2\}$ will never become blocked.

15.3.2 Global Invariants

When developing nonterminating programs it is often useful to state a global invariant. As should be clear from the previous examples, LSP_S is well suited to deal with invariants. If for example

$$\vdash_S z \underline{\text{sat}} [\vartheta, \alpha] :: [P, R, W, G, E],$$

then to prove that the assertion T is a global invariant, it is enough to show that

$$\begin{aligned} \vdash P &\Rightarrow T, \\ \vdash \overline{T} \wedge (G \vee R) &\Rightarrow T. \end{aligned}$$

Chapter 16

Dekker's Algorithm

16.1 Task

We will now use LSP_S to develop a program with respect to the four properties discussed above. Consider the following problem:

- Two processes $P(0)$ and $P(1)$ are executing in an infinite loop. Both processes consist of two sections; a critical section and a uncritical section. The executions of the two critical sections are not allowed to overlap. Our job is to find an implementation.

Given that the sets of memory locations accessed by the critical sections are not disjoint, then this can easily be achieved by placing the two critical sections inside the bodies of two await-statements. Thus, to increase the challenge, we will add an extra constraint:

- Each process can only ‘hide’ at most one memory location at the same time.

Exactly what the critical and uncritical sections are doing is not known. However, we will assume that their access is restricted to a global data structure Glb of an unspecified sort T .

16.2 Development

16.2.1 Data Structure

We will employ Dekker's algorithm [Dij68] to deal with this mutual exclusion problem. Let

$$D_{(0,1)} = \{0, 1\}.$$

To do the basic synchronisation we will employ an array Ok , defined with $D_{(0,1)}$ as domain and $\mathbb{B}$ as range, and a variable $Turn$ of sort $D_{(0,1)}$. These memory locations will only be accessed by the processes $P(0)$ and $P(1)$. Thus we will restrict our attention to a program of the form:

```
begin
  loc  $Ok, Turn$ ;
   $\{P(0) \parallel P(1)\}$ 
end.
```

To simplify the presentation, we will use arithmetics modulo 2; $\ominus$ denotes subtraction modulo 2, while $\oplus$ stands for addition modulo 2.

16.2.2 Invariant

To prove mutual exclusion it is useful to introduce an auxiliary array $Crit$ that maps $D_{(0,1)}$ to $\mathbb{B}$; $Crit(l)$ is true whenever the process $P(l)$ is inside its critical section. Thus

$$\neg(Crit(l) \wedge Crit(l \oplus 1))$$

is an invariant.

Whenever the process $P(l)$ is ready to let the other process $P(l \oplus 1)$ enter its critical section, then $Ok(l)$ is true. To secure mutual exclusion we must therefore insist that

$$Crit(l) \Rightarrow \neg Ok(l).$$

We will use Inv to denote the conjunction of these two assertions.

16.2.3 Main Structure

The process $P(l)$ can therefore be structured as below:

```

{ $Ok(l) \wedge \neg Crit(l) \wedge Inv$ }

while true do

    { $Ok(l) \wedge \neg Crit(l) \wedge Inv$ }

     $GetAcc(l)$ ;

    { $\neg Ok(l) \wedge Crit(l) \wedge \neg Crit(l \oplus 1) \wedge Inv$ }

     $DoCrit(l)$ ;

    { $\neg Ok(l) \wedge Crit(l) \wedge \neg Crit(l \oplus 1) \wedge Inv$ }

     $RlsAcc(l)$ ;

    { $Ok(l) \wedge \neg Crit(l) \wedge Inv$ }

     $DoUnCrit(l)$ 

    { $Ok(l) \wedge \neg Crit(l) \wedge Inv$ }

od

{false}.

```

Obviously, $DoCrit(l)$ and $DoUnCrit(l)$ are the critical and uncritical sections. Moreover, the object of $GetAcc(l)$ is to find the right moment to let the process $P(l)$ enter its critical section, while $RlsAcc(l)$ makes it possible for the process $P(l \oplus 1)$ to start on its critical section.

16.2.4 Guar-Condition

The critical and uncritical sections are assumed to leave the synchronisation structure unchanged, thus the atomic state changes of $DoCrit(l)$ and $DoUnCrit(l)$ must imply:

$$Ok = \overleftarrow{Ok} \wedge Turn = \overleftarrow{Turn} \wedge Crit = \overleftarrow{Crit}.$$

Furthermore, because the only task of $GetAcc(l)$ and $RlsAcc(l)$ is to secure mutual exclusion, their internal transitions must leave Glb unchanged.

Since the process $P(l)$ is only allowed to enter its critical section if $Ok(l \oplus 1)$ is true, and because $P(l)$ cannot change the value of $Ok(l \oplus 1)$, it follows that any state transition due to $P(l)$ must satisfy:

$$\neg \overleftarrow{Crit}(l) \wedge Crit(l) \Rightarrow Ok(l \oplus 1).$$

To avoid one process infinitely overtaking the other, $Turn$ will be used to provide fairness in cases where both processes want to enter their respective critical sections.

If both processes want to enter their critical sections and $Turn = l$, then $P(l)$ is the last process to have completed its critical section, while $Turn = l \oplus 1$ implies that $P(l \oplus 1)$ is the most recent process to have finished its critical section. Thus, it is enough if $P(l)$ updates $Turn$ once per iteration, namely immediately after it has finished its critical section. Furthermore, it is clear that the guar-condition of $P(l)$ must imply:

$$\overleftarrow{Turn} = l \Rightarrow Turn = l.$$

16.2.5 Wait-Condition

If both processes want to enter their critical sections, in which case

$$\neg Ok(l) \wedge \neg Ok(l \oplus 1)$$

is true, we can use $Turn$ to hold back the process that last finished its critical section, namely by changing its Ok location to true, and let it wait until the

other process changes the value of $Turn$. This means that the process $P(l)$ may have to wait in a state which satisfies:

$$Ok(l) \wedge Turn = l.$$

Moreover, the process $P(l)$ may also have to be held back if it wants to enter its critical section in a state that satisfies:

$$Turn = l \oplus 1 \wedge \neg Ok(l) \wedge \neg Ok(l \oplus 1).$$

In this case, the process $P(l \oplus 1)$ has not had enough time to switch on its Ok flag.

16.2.6 Specification

This leaves us with the following specification of the process $P(l)$, ($0 \leq l \leq 1$):

$P(l : D_{(0,1)})$

glo	ioeo	$Turn$	:	$D_{(0,1)},$
	ioeo	$Gl b$	:	$T,$
	ioeh	$Ok(l)$	:	$\mathbb{B},$
	iheo	$Ok(l + 1)$	:	$\mathbb{B}$

aux	ioeh	$Crit(l)$	:	$\mathbb{B},$
	iheo	$Crit(l + 1)$	:	$\mathbb{B}$

pre	$\neg Crit(l) \wedge Ok(l) \wedge Inv$
rely	$\overline{(Turn = l \oplus 1 \Rightarrow Turn = l \oplus 1)} \wedge$ $\overline{(\neg Crit(l \oplus 1) \wedge Crit(l \oplus 1) \Rightarrow Ok(l))} \wedge$ $\overline{(Inv \Rightarrow Inv)}$
wait	$(Ok(l) \wedge Turn = l) \vee$ $(\neg Ok(l) \wedge Turn = l \oplus 1 \wedge \neg Ok(l \oplus 1))$
guar	$\overline{(Turn = l \Rightarrow Turn = l)} \wedge$ $\overline{(\neg Crit(l) \wedge Crit(l) \Rightarrow Ok(l \oplus 1))} \wedge$ $\overline{(Inv \Rightarrow Inv)}$
eff	false.

16.2.7 Composition Proof

The obvious question at this stage is: May the processes $P(0)$ and $P(1)$ be composed in parallel? The answer is — Yes! Since

$$\begin{aligned} \vdash \neg(&((Ok(l) \wedge Turn = l) \vee \\ &(\neg Ok(l) \wedge Turn = l \oplus 1 \wedge \neg Ok(l \oplus 1))) \wedge \\ &((Ok(l \oplus 1) \wedge Turn = l \oplus 1) \vee \\ &(\neg Ok(l \oplus 1) \wedge Turn = l \wedge \neg Ok(l))))), \end{aligned}$$

it follows by the consequence- and parallel-rules that $\{P(0) \parallel P(1)\}$ satisfies:

glo	$\{Turn, Ok, Glb\}$
aux	$\{Crit\}$
pre	$\neg Crit(l) \wedge \neg Crit(l \oplus 1) \wedge Ok(l) \wedge Ok(l \oplus 1) \wedge Inv$
rely	I
wait	false
guar	true
eff	false.

16.2.8 Get Access

The next step is to implement $GetAcc(l)$. From the earlier discussion it follows that an implementation of $GetAcc(l)$ must satisfy:

```

glo    {Turn, Ok, Glb}
aux    {Crit}
pre     $\overline{Ok(l)} \wedge Inv$ 
rely     $(\overline{Turn = l \oplus 1} \Rightarrow Turn = l \oplus 1) \wedge$ 
         $(\overline{\neg Crit(l \oplus 1)} \wedge Crit(l \oplus 1) \Rightarrow Ok(l)) \wedge$ 
         $(\overline{Inv} \Rightarrow Inv)$ 
wait     $(Ok(l) \wedge Turn = l) \vee$ 
         $(\neg Ok(l) \wedge Turn = l \oplus 1 \wedge \neg Ok(l \oplus 1))$ 
guar     $(\overline{Turn = l} \Rightarrow Turn = l) \wedge$ 
         $(\overline{\neg Crit(l)} \wedge Crit(l) \Rightarrow Ok(l \oplus 1)) \wedge$ 
         $Glb = \overline{Glb} \wedge (\overline{Inv} \Rightarrow Inv)$ 
eff      $Crit(l) \wedge Inv.$ 

```

Thus, the following is a correct implementation:

```

begin
  loc  $V_l$ ;

   $\{Ok(l) \wedge \neg Crit(l) \wedge Inv\}$ 

   $Ok(l) := \text{false};$ 
   $V_l := Ok(l \oplus 1);$ 

   $\{\neg Ok(l) \wedge (Crit(l) \Leftrightarrow V_l) \wedge Inv\}$ 

  if  $\neg V_l$  then

     $\{\neg Ok(l) \wedge \neg Crit(l) \wedge Inv\}$ 

    if  $Turn = l$  then
       $Ok(l) := \text{true};$ 

       $\{Ok(l) \wedge \neg Crit(l) \wedge Inv\}$ 

      await  $Turn = l \oplus 1$  do skip od;

       $\{Ok(l) \wedge \neg Crit(l) \wedge Turn = l \oplus 1 \wedge Inv\}$ 

       $Ok(l) := \text{false};$ 
    fi;

     $\{\neg Ok(l) \wedge \neg Crit(l) \wedge Turn = l \oplus 1 \wedge Inv\}$ 

    await  $Ok(l \oplus 1)$  do skip od

     $\{\neg Ok(l) \wedge Crit(l) \wedge \neg Crit(l \oplus 1) \wedge Turn = l \oplus 1 \wedge Inv\}$ 

  fi

   $\{\neg Ok(l) \wedge Crit(l) \wedge \neg Crit(l \oplus 1) \wedge Inv\}$ 
end.

```

16.2.9 Release Access

What remains is to implement *RlsAcc*. In other words, to find a program that satisfies:

```

glo    {Turn, Ok, Glb}
aux    {Crit}
pre    Crit(l) ∧ Inv
rely    (Turn = l ⊕ 1 ⇒ Turn = l ⊕ 1) ∧
        (¬Crit(l ⊕ 1) ∧ Crit(l ⊕ 1) ⇒ Ok(l)) ∧
        (Inv ⇒ Inv)
wait    (Ok(l) ∧ Turn = l) ∨
        (¬Ok(l) ∧ Turn = l ⊕ 1 ∧ ¬Ok(l ⊕ 1))
guar    (Turn = l ⇒ Turn = l) ∧
        (¬Crit(l) ∧ Crit(l) ⇒ Ok(l ⊕ 1)) ∧
        Glb = Glb ∧ (Inv ⇒ Inv)
eff     Ok(l) ∧ ¬Crit(l) ∧ Inv.

```

This is not very difficult:

$$\{\neg Ok(l) \wedge Crit(l) \wedge \neg Crit(l \oplus 1) \wedge Inv\}$$

$$Ok(l) := \text{true}$$

$$\{Ok(l) \wedge \neg Crit(l) \wedge Inv\}.$$

Chapter 17

Soundness

17.1 Parallel-Decomposition Proposition

The object of this chapter is prove that LSP is sound. This will be shown by induction on the depth (see page 10.1.4) of a LSP proof. The following proposition, which characterises a decomposition condition, will be useful:

Proposition 7 *Given that*

$$\models_{\pi} \neg(W_1 \wedge E_2) \wedge \neg(W_2 \wedge E_1) \wedge \neg(W_1 \wedge W_2), \quad (17.1)$$

$$ext_{\pi}[P, R_1] \cap cp_{\pi}[z_1] \subseteq int_{\pi}[W \vee W_1, G \wedge R_2, E_1], \quad (17.2)$$

$$ext_{\pi}[P, R_2] \cap cp_{\pi}[z_2] \subseteq int_{\pi}[W \vee W_2, G \wedge R_1, E_2], \quad (17.3)$$

$$\{z_1 \parallel z_2\} \text{ \underline{sat} } (\vartheta, \{\}) :: (P, R_1 \wedge R_2, W, G, E_1 \wedge E_2) \in SP,$$

then

$$ext_{\pi}[P, R_1 \wedge R_2] \cap cp_{\pi}[\{z_1 \parallel z_2\}] \subseteq int_{\pi}[W, G, E_1 \wedge E_2]. \quad (17.4)$$

Proof: Let

$$\sigma \in ext_{\pi}[P, R_1 \wedge R_2] \cap cp_{\pi}[\{z_1 \parallel z_2\}], \quad (17.5)$$

it follows from proposition 2 on page 51 that there are two computations $\sigma' \in cp_{\pi}[z_1]$ and $\sigma'' \in cp_{\pi}[z_2]$, such that $\sigma' \bullet \sigma'' \leftarrow \sigma$. We will first show that $\sigma' \in ext_{\pi}[P, R_1]$ and $\sigma'' \in ext_{\pi}[P, R_2]$; in other words, that both computations satisfy their respective instances of the three conditions in definition 11 on page 68.

Since (17.5) implies that $\delta(\sigma_1) \models_\pi P$, and since by definition $\delta(\sigma_1) = \delta(\sigma'_1) = \delta(\sigma''_1)$, it follows that $\delta(\sigma'_1) \models_\pi P$ and $\delta(\sigma''_1) \models_\pi P$. Hence, they both satisfy the first condition.

To prove that they also satisfy the second condition, let:

- $\max(\sigma') = \text{len}(\sigma')$, if for all $1 \leq k \leq \text{len}(\sigma')$, there is a $\sigma''' \in \text{ext}_\pi[P, R_1] \cap \text{cp}_\pi[z_1]$, which satisfies

$$\sigma'(1, \dots, k) = \sigma'''(1, \dots, k).$$

Otherwise, $\max(\sigma') = m$, where m is the maximum natural number such that there is a $\sigma''' \in \text{ext}_\pi[P, R_1] \cap \text{cp}_\pi[z_1]$, which satisfies

$$\sigma'(1, \dots, m) = \sigma'''(1, \dots, m).$$

- $\max(\sigma'') = \text{len}(\sigma'')$, if for all $1 \leq k \leq \text{len}(\sigma'')$, there is a $\sigma''' \in \text{ext}_\pi[P, R_2] \cap \text{cp}_\pi[z_2]$, which satisfies

$$\sigma''(1, \dots, k) = \sigma'''(1, \dots, k).$$

Otherwise, $\max(\sigma'') = m$, where m is the maximum natural number such that there is a $\sigma''' \in \text{ext}_\pi[P, R_2] \cap \text{cp}_\pi[z_2]$, which satisfies

$$\sigma''(1, \dots, m) = \sigma'''(1, \dots, m).$$

If $\max(\sigma') = \text{len}(\sigma')$, then it is obvious that σ' satisfies the second condition. The same is of course true for σ'' if $\max(\sigma'') = \text{len}(\sigma'')$.

Assume that $\max(\sigma') \neq \text{len}(\sigma')$ or $\max(\sigma'') \neq \text{len}(\sigma'')$. We will show that this leads to a contradiction. There are two cases:

- $m = \max(\sigma') = \max(\sigma'')$: Since the constraints imposed on the environment have no effect on the number of possible internal transitions in a given configuration, and since $\sigma' \bullet \sigma'' \leftarrow \sigma$, it is clear that

$$\lambda(\sigma'_m) = \lambda(\sigma''_m) = \lambda(\sigma_m) = e.$$

But then, $(\delta(\sigma_m), \delta(\sigma_{m+1})) \models_\pi \neg(R_1 \wedge R_2)$ which contradicts (17.5).

- $\max(\sigma') \neq \max(\sigma'')$: Without loss of generality, it may be assumed that

$$m = \max(\sigma') < \max(\sigma'').$$

Moreover, since the constraints imposed on the environment have no effect on the number of possible internal transitions in a given configuration, there are two possibilities:

- $\lambda(\sigma'_m) = \lambda(\sigma''_m) = e$: This leads to a contradiction by an argument similar to the one above.
- $\lambda(\sigma'_m) = e$ and $\lambda(\sigma''_m) = i$: Then there is a $\sigma''' \in \text{ext}_\pi[P, R_2] \cap \text{cp}_\pi[z_2]$ such that $\sigma''(1, \dots, m+1) = \sigma'''(1, \dots, m+1)$. From (17.3) it follows that $\sigma''' \in \text{int}_\pi[W \vee W_2, G \wedge R_1, E_2]$, which implies that

$$(\delta(\sigma''_m), \delta(\sigma''_{m+1})) \models_\pi G \wedge R_1.$$

Then, since $\sigma' \bullet \sigma'' \leftarrow \sigma$, it is also true that

$$(\delta(\sigma'_m), \delta(\sigma'_{m+1})) \models_\pi G \wedge R_1,$$

which again implies that

$$(\delta(\sigma'_m), \delta(\sigma'_{m+1})) \models_\pi R_1.$$

But, this contradicts that $\max(\sigma') = m$.

Thus, both σ' and σ'' satisfy the second condition in definition 11 on page 68. To see that they also satisfy the final constraint, assume that σ diverges. This means that for any $j \geq 1$, there is a $k \geq j$, such that $\lambda(\sigma_k) = i$. But then $\sigma' \bullet \sigma'' \leftarrow \sigma$ implies that for any $j \geq 1$, there is a $k \geq j$, such that $\lambda(\sigma'_k) = i$ or $\lambda(\sigma''_k) = i$. Thus, $\sigma' \in \text{ext}_\pi[P, R_1]$ or $\sigma'' \in \text{ext}_\pi[P, R_2]$, which contradicts (17.2) or (17.3). This means that both σ' and σ'' are finite, and it is clear that

$$\begin{aligned} \sigma' &\in \text{ext}_\pi[P, R_1] \cap \text{cp}_\pi[z_1], \\ \sigma'' &\in \text{ext}_\pi[P, R_2] \cap \text{cp}_\pi[z_2]. \end{aligned}$$

But then

$$\begin{aligned} \sigma' &\in \text{int}_\pi[W \vee W_1, G \vee R_2, E_1], \\ \sigma'' &\in \text{int}_\pi[W \vee W_2, G \vee R_1, E_2] \end{aligned}$$

follows from (17.2) and (17.3). Hence, (17.1) implies that

$$\sigma \in \text{int}_\pi[W, G, E_1 \wedge E_2].$$

This proves (17.4).

end of proof

17.2 Soundness Proposition

Proposition 8 *For any structure π and specified program ψ , if*

$$\text{Tr}_\pi \vdash \psi$$

then

$$\models_\pi \psi.$$

Proof: We will show the proposition by induction on the proof's depth (see page 88). If the proof's depth is one, there are only two possibilities:

- Skip-Rule:

Assume that we have a proof of depth 1 one whose root

$$\text{Tr}_\pi \vdash \text{skip } \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, R)$$

is deduced by the skip-rule. Since

$$\models \text{skip} \xrightarrow{(\vartheta, \alpha)} \text{skip},$$

it is enough to show that

$$\text{ext}_\pi[P, R] \cap \text{cp}_\pi[\text{skip}] \subseteq \text{int}_\pi[W, G, R]. \quad (17.6)$$

Let

$$\sigma \in \text{ext}_\pi[P, R] \cap \text{cp}_\pi[\text{skip}],$$

then σ is of the form

$$\langle \text{skip}, s_1 \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle \text{skip}, s_j \rangle \xrightarrow{i} \langle \epsilon, s_j \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle \epsilon, s_n \rangle,$$

in which case the reflexivity and transitivity of R , and the reflexivity of G imply that $\sigma \in \text{int}_\pi[W, G, R]$. This proves (17.6).

- Assignment-Rule:

Assume z is of the form

$$v := r,$$

and that we have a proof of depth 1 whose root

$$Tr_\pi \vdash z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, R|E|R)$$

is deduced by the assignment-rule. This means that there are expressions $u_1, \dots, u_{card(\alpha)}$ such that

$$\models_\pi \overleftarrow{P^R} \wedge v = \overleftarrow{r} \wedge I_{\{v\} \cup \alpha} \wedge \bigwedge_{j=1}^{card(\alpha)} a_j = \overleftarrow{u_j} \Rightarrow G \wedge E, \quad (17.7)$$

$$\begin{aligned} & \text{for all } 1 \leq j, k \leq card(\alpha), \\ & \quad var[u_j] \in \vartheta \cup \{a_j\}, \\ & \quad a_j \in \alpha, \text{ and } j \neq k \text{ implies } a_j \neq a_k. \end{aligned} \quad (17.8)$$

Let z' denote the program

```

await true do
   $a_1 := u_1;$ 
   $\vdots$ 
   $a_{card(\alpha)} := u_{card(\alpha)};$ 
   $v := r$ 
od,
```

then it follows from (17.8) that

$$\models z' \xrightarrow{(\vartheta, \alpha)} z,$$

and it is enough to show that

$$ext_\pi[P, R] \cap cp_\pi[z'] \subseteq int_\pi[W, G, R|E|R]. \quad (17.9)$$

Let

$$\sigma \in \text{ext}_\pi[P, R] \cap \text{cp}_\pi[z'],$$

then σ is of the form

$$\langle z', s_1 \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle z', s_{j-1} \rangle \xrightarrow{i} \langle \epsilon, s_j \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle \epsilon, s_{j+m} \rangle,$$

where the internal transition is characterised by

$$\overleftarrow{P^R} \wedge v = \overleftarrow{r} \wedge I_{\{v\} \cup \alpha} \wedge \bigwedge_{j=1}^{\text{card}(\alpha)} a_j = \overleftarrow{u_j}.$$

Hence, (17.7) and the reflexivity and transitivity of R imply that

$$\sigma \in \text{int}_\pi[W, G, R|E|R],$$

which again implies (17.9).

This means that the proposition is true for proofs of depth one. Assume the proposition is true for proofs of depth less than or equal to $n - 1$. We will show that the proposition is true for proofs of depth n . There are fourteen cases:

- Consequence-Rule:

Assume we have a proof of depth n , whose root

$$\text{Tr}_\pi \vdash z \text{ sat } (\vartheta, \alpha) :: (P_2, R_2, W_2, G_2, E_2)$$

is deduced from

$$\text{Tr}_\pi \vdash z \text{ sat } (\vartheta, \alpha) :: (P_1, R_1, W_1, G_1, E_1),$$

by the consequence-rule. This means that

$$\models_\pi P_2 \Rightarrow P_1, \tag{17.10}$$

$$\models_\pi R_2 \Rightarrow R_1, \tag{17.11}$$

$$\models_\pi W_1 \Rightarrow W_2, \tag{17.12}$$

$$\models_\pi G_1 \Rightarrow G_2, \tag{17.13}$$

$$\models_\pi E_1 \Rightarrow E_2. \tag{17.14}$$

Moreover, the induction hypothesis implies that

$$\models_{\pi} z \text{ \underline{sat} } (\vartheta, \alpha) :: (P_1, R_1, W_1, G_1, E_1),$$

which means that there is a program z' such that

$$\begin{aligned} & \models z' \xrightarrow{(\vartheta, \alpha)} z, \\ & \text{ext}_{\pi}[P_1, R_1] \cap \text{cp}_{\pi}[z'] \subseteq \text{int}_{\pi}[W_1, G_1, E_1]. \end{aligned} \quad (17.15)$$

Thus, it is enough to show that

$$\text{ext}_{\pi}[P_2, R_2] \cap \text{cp}_{\pi}[z'] \subseteq \text{int}_{\pi}[W_2, G_2, E_2],$$

which follows from (17.15), since (17.10) and (17.11) imply that

$$\text{ext}_{\pi}[P_2, R_2] \subseteq \text{ext}_{\pi}[P_1, R_1],$$

and (17.12), (17.13) and (17.14) imply that

$$\text{int}_{\pi}[W_1, G_1, E_1] \subseteq \text{int}_{\pi}[W_2, G_2, E_2].$$

- Pre-Rule:

Assume we have a proof of depth n , whose root

$$\text{Tr}_{\pi} \vdash z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, \overleftarrow{P} \wedge E)$$

is deduced from

$$\text{Tr}_{\pi} \vdash z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E),$$

by the pre-rule. Then the induction hypothesis implies that

$$\models_{\pi} z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E),$$

which means that there is a program z' such that

$$\begin{aligned} & \models z' \xrightarrow{(\vartheta, \alpha)} z, \\ & \text{ext}_{\pi}[P, R] \cap \text{cp}_{\pi}[z'] \subseteq \text{int}_{\pi}[W, G, E]. \end{aligned}$$

Thus, it is enough to show that

$$ext_\pi[P, R] \cap cp_\pi[z'] \subseteq int_\pi[W, G, \overleftarrow{P} \wedge E]. \quad (17.16)$$

Let

$$\sigma \in ext_\pi[P, R] \cap cp_\pi[z'].$$

Since

$$int_\pi[W, G, \overleftarrow{P} \wedge E] = \{\sigma \in int_\pi[W, G, E] \mid \delta(\sigma_1) \models_\pi P\},$$

and since by definition $\delta(\sigma_1) \models_\pi P$, it follows that

$$\sigma \in int_\pi[W, G, \overleftarrow{P} \wedge E].$$

In other words, (17.16) has been shown.

- Access-Rule:

Assume we have a proof of depth n , whose root

$$Tr_\pi \vdash z \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E)$$

is deduced from

$$Tr_\pi \vdash z \underline{\text{sat}} (\vartheta, \alpha) :: (P, R \wedge v = \overleftarrow{v}, W, G, E)$$

by the access-rule. This means that

$$v \in hid[z] \cap \vartheta.$$

Moreover, the induction hypothesis implies that

$$\models_\pi z \underline{\text{sat}} (\vartheta, \alpha) :: (P, R \wedge v = \overleftarrow{v}, W, G, E)$$

This means that there is a program z' such that

$$\begin{aligned} & \models z' \xrightarrow{(\vartheta, \alpha)} z, \\ & ext_\pi[P, R \wedge v = \overleftarrow{v}] \cap cp_\pi[z'] \subseteq int_\pi[W, G, E]. \end{aligned} \quad (17.17)$$

Thus, it is enough to show that

$$ext_\pi[P, R] \cap cp_\pi[z'] \subseteq int_\pi[W, G, E]. \quad (17.18)$$

Let

$$\sigma \in ext_\pi[P, R] \cap cp_\pi[z'].$$

Since $v \in hid[z']$ and the environment is required to respect $hid[z']$, it is clear that

$$\sigma \in ext_\pi[P, R \wedge v = \overleftarrow{v}] \cap cp_\pi[z'],$$

in which case (17.17) implies that

$$\sigma \in int_\pi[W, G, E],$$

which proves (17.18).

- Block-Rule:

Given that z_1 is of the form

$$\text{begin loc } v_1, \dots, v_m; z_2 \text{ end},$$

and assume we have a proof of depth n , whose root

$$Tr_\pi \vdash z_1 \underline{\text{sat}} (\vartheta \setminus \bigcup_{j=1}^m \{v_j\}, \alpha) :: (P, R, W, G, E)$$

is deduced from

$$Tr_\pi \vdash z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}, W, G, E),$$

by the block-rule. Then the induction hypothesis implies that

$$\models_{\pi} z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}, W, G, E),$$

which means that there is a program z'_2 such that

$$\models z'_2 \xrightarrow{(\vartheta, \alpha)} z_2, \quad (17.19)$$

$$\text{ext}_{\pi}[P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}] \cap \text{cp}_{\pi}[z'_2] \subseteq \text{int}_{\pi}[W, G, E]. \quad (17.20)$$

Moreover, if z'_1 denotes the program

begin loc $v_1, \dots, v_m; z'_2$ **end**

it follows from (17.19) that

$$\models z'_1 \xrightarrow{(\vartheta, \alpha)} z_1.$$

Thus, it is enough to show that

$$\text{ext}_{\pi}[P, R] \cap \text{cp}_{\pi}[z'_1] \subseteq \text{int}_{\pi}[W, G, E]. \quad (17.21)$$

Observe that the constraints on a specified program imply that $v_1, \dots, v_m$ cannot occur in P, R, W, G or E . Let

$$\sigma \in \text{ext}_{\pi}[P, R] \cap \text{cp}_{\pi}[z'_1].$$

Since any external transition in σ must respect $\{v_1, \dots, v_n\}$, it follows that σ is of the form

$$\langle z'_1, s_1 \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle z'_1, s_j \rangle \xrightarrow{i} \sigma',$$

where

$$\sigma' \in \text{ext}_{\pi}[P^R, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}] \cap \text{cp}_{\pi}[z'_2].$$

Moreover, since $\delta(\sigma'_1) = s_j$, it is also clear that

$$< z'_2, s_1 > \xrightarrow{e} \dots \xrightarrow{e} < z'_2, s_j > \xrightarrow{e} \sigma'$$

is an element of $ext_\pi[P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}] \cap cp_\pi[z'_2]$. Thus, it follows from (17.20) that $\sigma \in int_\pi[W, G, E]$. This proves (17.21).

- Sequential-Rule:

Given that z_1 is of the form

$$z_3; z_2,$$

and assume we have a proof of depth n , whose root

$$Tr_\pi \vdash z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P_1, R, W, G, E_1 | E_2)$$

is deduced from

$$\begin{aligned} Tr_\pi \vdash z_3 \underline{\text{sat}} (\vartheta, \alpha) &:: (P_1, R, W, G, P_2 \wedge E_1), \\ Tr_\pi \vdash z_2 \underline{\text{sat}} (\vartheta, \alpha) &:: (P_2, R, W, G, E_2), \end{aligned}$$

by the sequential-rule. Then the induction hypothesis implies that

$$\begin{aligned} \models_\pi z_3 \underline{\text{sat}} (\vartheta, \alpha) &:: (P_1, R, W, G, P_2 \wedge E_1), \\ \models_\pi z_2 \underline{\text{sat}} (\vartheta, \alpha) &:: (P_2, R, W, G, E_2), \end{aligned}$$

which means that there are two programs z'_3 and z'_2 such that

$$\models z'_3 \xrightarrow{(\vartheta, \alpha)} z_3, \tag{17.22}$$

$$\models z'_2 \xrightarrow{(\vartheta, \alpha)} z_2, \tag{17.23}$$

$$ext_\pi[P_1, R] \cap cp_\pi[z'_3] \subseteq int_\pi[W, G, P_2 \wedge E_1], \tag{17.24}$$

$$ext_\pi[P_2, R] \cap cp_\pi[z'_2] \subseteq int_\pi[W, G, E_2]. \tag{17.25}$$

Moreover, if z'_1 denotes the program

$$z'_3; z'_2,$$

it follows from (17.22) and (17.23) that

$$\models z'_1 \xrightarrow{(\vartheta, \alpha)} z_1.$$

Thus, it is enough to show that

$$ext_\pi[P_1, R] \cap cp_\pi[z'_1] \subseteq int_\pi[W, G, E_1|E_2]. \quad (17.26)$$

Let

$$\sigma \in ext_\pi[P_1, R] \cap cp_\pi[z'_1].$$

It is clear from (17.24) and (17.25) that σ converges. If σ deadlocks before z'_3 has terminated, it follows from (17.24) that $\sigma \in int_\pi[W, G, P_2 \wedge E_1]$, in which case it is also true that $\sigma \in int_\pi[W, G, E_1|E_2]$.

On the other hand, if σ terminates or deadlocks after z'_3 has terminated, it follows that σ is of the form

$$< z'_3; z'_2, s_1 > \xrightarrow{a_1} \dots \xrightarrow{a_{k-2}} < z'_{k-1}; z'_2, s_{k-1} > \xrightarrow{a_{k-1}} < z'_2, s_k > \xrightarrow{a_k} \sigma'$$

where

$$< z_2, s_k > \xrightarrow{a_k} \sigma' \in ext_\pi[P_2, R] \cap cp_\pi[z'_2],$$

and there is a $\sigma'' \in ext_\pi[P_1, R] \cap cp_\pi[z'_3]$ of the form

$$< z'_3, s_1 > \xrightarrow{a_1} \dots \xrightarrow{a_{k-2}} < z'_{k-1}, s_{k-1} > \xrightarrow{a_{k-1}} < \epsilon, s_k > .$$

Thus, (17.24) and (17.25) imply that $\sigma \in int_\pi[W, G, E_1|E_2]$, which proves (17.26).

- If-Rule:

Given that z_1 is of the form

if b then z_2 else z_3 fi,

and assume we have a proof of depth n , whose root

$$Tr_\pi \vdash z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E)$$

is deduced from

$$\begin{aligned} Tr_\pi \vdash z_2 \underline{\text{sat}} (\vartheta, \alpha) &:: (P \wedge b, R, W, G, E), \\ Tr_\pi \vdash z_3 \underline{\text{sat}} (\vartheta, \alpha) &:: (P \wedge \neg b, R, W, G, E), \end{aligned}$$

by the if-rule. Then the induction hypothesis implies that

$$\begin{aligned} \models_\pi z_2 \underline{\text{sat}} (\vartheta, \alpha) &:: (P \wedge b, R, W, G, E), \\ \models_\pi z_3 \underline{\text{sat}} (\vartheta, \alpha) &:: (P \wedge \neg b, R, W, G, E), \end{aligned}$$

which means that there are two programs z'_2 and z'_3 such that

$$\models z'_2 \xrightarrow{(\vartheta, \alpha)} z_2, \quad (17.27)$$

$$\models z'_3 \xrightarrow{(\vartheta, \alpha)} z_3, \quad (17.28)$$

$$\text{ext}_\pi[P \wedge b, R] \cap \text{cp}_\pi[z_2] \subseteq \text{int}_\pi[W, G, E], \quad (17.29)$$

$$\text{ext}_\pi[P \wedge \neg b, R] \cap \text{cp}_\pi[z_3] \subseteq \text{int}_\pi[W, G, E]. \quad (17.30)$$

Moreover, if z'_1 denotes the program

$$\text{if } b \text{ then } z'_2 \text{ else } z'_3 \text{ fi},$$

it follows from (17.27) and (17.28) that

$$\models z'_1 \xrightarrow{(\vartheta, \alpha)} z_1.$$

Thus, it is enough to show that

$$\text{ext}_\pi[P, R] \cap \text{cp}_\pi[z_1] \subseteq \text{int}_\pi[W, G, E]. \quad (17.31)$$

Let

$$\sigma \in \text{ext}_\pi[P, R] \cap \text{cp}_\pi[z'_1].$$

It is clear that σ is of the form

$$\langle z'_1, s_1 \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle z'_1, s_j \rangle \xrightarrow{i} \sigma'$$

where

$$\sigma' \in (\text{ext}_\pi[P^R \wedge b, R] \cap \text{cp}_\pi[z'_2]) \cup (\text{ext}_\pi[P^R \wedge \neg b, R] \cap \text{cp}_\pi[z'_3]).$$

Since any external transition must respect $\text{hid}[z'_1]$, it follows that no external transition can change the truth value of b . This means that even

$$\langle z'_2, s_1 \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle z'_2, s_j \rangle \xrightarrow{e} \sigma'$$

is an element of

$$(\text{ext}_\pi[P \wedge b, R] \cap \text{cp}_\pi[z'_2]) \cup (\text{ext}_\pi[P \wedge \neg b, R] \cap \text{cp}_\pi[z'_3]).$$

Thus, the reflexivity of G together with (17.29) and (17.30) imply that $\sigma \in \text{int}_\pi[W, G, E]$, which again implies (17.31).

- While-Rule:

Given that z_1 is of the form

while b do z_2 od,

and assume we have a proof of depth n , whose root

$$\text{Tr}_\pi \vdash z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, (Z^\dagger \vee R) \wedge \neg b)$$

is deduced from

$$\text{Tr}_\pi \vdash z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P \wedge b, R, W, G, P \wedge Z),$$

by the while-rule. This means that

$$\models_\pi \underline{\text{wf}} \ Z. \tag{17.32}$$

Moreover, the induction hypothesis implies that

$$\models_{\pi} z_2 \text{ sat } (\vartheta, \alpha) :: (P \wedge b, R, W, G, P \wedge Z),$$

which means that there is a program z'_2 such that

$$\models z'_2 \xrightarrow{(\vartheta, \alpha)} z_2, \quad (17.33)$$

$$\text{ext}_{\pi}[P \wedge b, R] \cap \text{cp}_{\pi}[z'_2] \subseteq \text{int}_{\pi}[W, G, P \wedge Z]. \quad (17.34)$$

Moreover, if z'_1 denotes the program

$$\text{while } b \text{ do } z'_2 \text{ od},$$

it follows from (17.33) that

$$\models z'_1 \xrightarrow{(\vartheta, \alpha)} z_1.$$

Thus, it is enough to show that

$$\text{ext}_{\pi}[P, R] \cap \text{cp}_{\pi}[z'_1] \subseteq \text{int}_{\pi}[W, G, (Z^{\dagger} \vee R) \wedge \neg b]. \quad (17.35)$$

Let

$$\sigma \in \text{ext}_{\pi}[P, R] \cap \text{cp}_{\pi}[z'_1].$$

There are three cases to consider:

- Assume that σ diverges:

This means there is an infinite sequence of natural numbers

$$n_1 n_2 n_3 n_4 n_5 n_6 n_7 \quad \dots \quad n_k n_{k+1} \quad \dots$$

such that:

- * for all j , $j < n_1$ implies that $\lambda(\sigma_j) = e$,
- * for all j ,
 - $n_j < n_{j+1}$,
 - $\tau(\sigma_{n_j}) = \tau(\sigma_{n_1})$,

- $\lambda(\sigma_{n_j}) = i$,
- $\tau(\sigma_{n_{j+1}}) = z'_2; \tau(\sigma_{n_1})$,
- $n_j < k < n_{j+1}$ and $\lambda(\sigma_k) = i$ imply $\tau(\sigma_k) \neq \tau(\sigma_{n_1})$.

In other words, $\sigma(n_j, \dots, n_{j+1})$ characterises the j -th iteration with its interference.

Since the environment cannot change the truth value of b , it is clear that for any computation $\sigma' \in \text{ext}_\pi[P, R] \cap \text{cp}_\pi[z'_1]$ of the form

$$< z'_1, s_1 > \xrightarrow{e} \dots \xrightarrow{e} < z'_1, s_m > \xrightarrow{i} \sigma'',$$

where the $m-1$ first transitions are due to the environment, there is a computation $\sigma''' \in \text{ext}_\pi[P, R] \cap \text{cp}_\pi[z'_1]$ of the form

$$< z'_1, s_1 > \xrightarrow{i} < z'_2, s_1 > \xrightarrow{e} \dots \xrightarrow{e} < z'_2, s_m > \xrightarrow{e} \sigma''.$$

Thus, we may assume that $n_1 = 1$, in which case it follows that $\delta(\sigma_{n_1}) \models_\pi P$, which together with (17.34) imply that for all $j \geq 1$, $\delta(\sigma_{n_j}) \models_\pi P$, and it is also clear that for all j ,

$$(\delta(\sigma_{n_j}), \delta(\sigma_{n_{j+1}})) \models_\pi Z.$$

This contradicts (17.32). In other words, the statement converges.

– Assume that σ deadlocks:

This means that there is a finite sequence of natural numbers

$$n_1 n_2 n_3 n_4 n_5 n_6 n_7 \quad \dots \quad n_{m-1} n_m,$$

such that

- * for all j , $j < n_1$ implies $\lambda(\sigma_j) = e$,
- * for all $j \leq m$:
 - $n_j < n_{j+1}$,
 - $\tau(\sigma_{n_j}) = \tau(\sigma_{n_1})$,
 - $\lambda(\sigma_{n_j}) = i$,
 - $\tau(\sigma_{n_{j+1}}) = z'_2; \tau(\sigma_{n_1})$,
 - $n_j < k < n_{j+1}$ and $\lambda(\sigma_k) = i$ imply $\tau(\sigma_k) \neq \tau(\sigma_{n_1})$,
- * $n_m < k < \text{len}(\sigma)$ and $\lambda(\sigma_k) = i$ implies $\tau(\sigma_k) \neq \tau(\sigma_{n_1})$.

In other words, the loop iterates $m-1$ times before it deadlocks. In the same way as above, we may assume that $n_1 = 1$, in which case it follows that $\delta(\sigma_{n_1}) \models_\pi P$, which together with (17.34) imply that for all $1 \leq j \leq m$, $\delta(\sigma_{n_j}) \models_\pi P$. Hence, it follows from (17.34) that $\delta(\text{len}(\sigma)) \models_\pi W$. Thus, $\sigma \in \text{int}_\pi[W, G, (Z^\dagger \vee R) \wedge \neg b]$.

– Assume that σ terminates:

This means that there is a finite sequence of natural numbers

$$n_1 n_2 n_3 n_4 n_5 n_6 n_7 \quad \dots \quad n_{m-1} n_m,$$

such that

- * for all j , $j < n_1$ or $j > n_m$ implies $\lambda(\sigma_j) = e$,
- * for all $j < m$:
 - $n_j < n_{j+1}$,
 - $\tau(\sigma_{n_j}) = \tau(\sigma_{n_1})$,
 - $\lambda(\sigma_{n_j}) = i$,
 - $\tau(\sigma_{n_{j+1}}) = z'_2; \tau(\sigma_{n_1})$,
 - $n_j < k < n_{j+1}$ and $\lambda(\sigma_k) = i$ implies $\tau(\sigma_k) \neq \tau(\sigma_{n_1})$,
- * $\lambda(\sigma_{n_m}) = i$,
- * $\tau(\sigma_{n_m}) = \tau(\sigma_{n_1})$,
- * $\tau(\sigma_{n_m+1}) = \epsilon$.

In other words, the loop iterates $m-1$ times before it terminates. In the same way as above, we may assume that $n_1 = 1$, in which case it follows that $\delta(\sigma_1) \models_\pi P$, which together with (17.34) imply that for all $1 \leq j \leq m$, $\delta(\sigma_{n_j}) \models_\pi P$.

Moreover, (17.34) and the transitivity of $Z^\dagger$, imply that

$$(\delta(\sigma_1), \delta(\sigma_{n_m})) \models_\pi Z^\dagger.$$

Therefore, since Z covers any interference after the last internal transition, and since any external transition must respect $\text{hid}[z_1]$ and hence cannot change the truth value of b , it is clear that:

$$(\delta(\sigma_1), \delta(\sigma_{\text{len}(\sigma)})) \models_\pi Z^\dagger \wedge \neg b.$$

On the other hand, if $m = 1$, then:

$$(\delta(\sigma_1), \delta(\sigma_{\text{len}(\sigma)})) \models_\pi R \wedge \neg b.$$

That each internal transition satisfies G follows from (17.34) and the reflexivity of G . Again, it is clear that $\sigma \in \text{int}_\pi[W, G, (Z^\dagger \vee R) \wedge \neg b]$.

This proves (17.35).

- Parallel-Rule:

Given that z_1 is of the form

$$\{z_2 \parallel z_3\},$$

and assume we have a proof of depth n , whose root

$$\text{Tr}_\pi \vdash z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_1 \wedge R_2, W, G, E_1 \wedge E_2)$$

is deduced from

$$\text{Tr}_\pi \vdash z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_1, W \vee W_1, G \wedge R_2, E_1),$$

$$\text{Tr}_\pi \vdash z_3 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_2, W \vee W_2, G \wedge R_1, E_2),$$

by the parallel-rule. This means that

$$\models_\pi \neg(W_1 \wedge E_2) \wedge \neg(W_2 \wedge E_1) \wedge \neg(W_1 \wedge W_2). \quad (17.36)$$

Moreover, the induction hypothesis implies that

$$\models_\pi z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_1, W \vee W_1, G \wedge R_2, E_1),$$

$$\models_\pi z_3 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R_2, W \vee W_2, G \wedge R_1, E_2),$$

which means that there are two programs z'_2 and z'_3 such that

$$\models z'_2 \xrightarrow{(\vartheta, \alpha)} z_2, \quad (17.37)$$

$$\models z'_3 \xrightarrow{(\vartheta, \alpha)} z_3, \quad (17.38)$$

$$\text{ext}_\pi[P, R_1] \cap \text{cp}_\pi[z'_2] \subseteq \text{int}_\pi[W \vee W_1, G \wedge R_2, E_1], \quad (17.39)$$

$$\text{ext}_\pi[P, R_2] \cap \text{cp}_\pi[z'_3] \subseteq \text{int}_\pi[W \vee W_2, G \wedge R_1, E_2]. \quad (17.40)$$

Moreover, if z'_1 denotes the program

$$\{z'_2 \parallel z'_3\},$$

it follows from (17.37) and (17.38) that

$$\models z'_1 \xrightarrow{(\vartheta, \alpha)} z_1.$$

Thus, it is enough to show that

$$ext_\pi[P, R_1 \wedge R_2] \cap cp_\pi[z'_1] \subseteq int_\pi[W, G, E_1 \wedge E_2],$$

which follows from (17.36), (17.39) and (17.40) by proposition 7 on page 153.

- Await-Rule:

Given that z_1 is of the form

$$\text{await } b \text{ do } z_2 \text{ od},$$

and assume we have a proof of depth n , whose root

$$Tr_\pi \vdash z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, R|E_2|R)$$

is deduced from

$$Tr_\pi \vdash z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P^R \wedge b, I, \text{false}, \text{true}, E_1),$$

by the await-rule. This means that there are expressions $u_1, \dots, u_{card(\alpha)}$ such that

$$\models_\pi P^R \wedge \neg b \Rightarrow W, \quad (17.41)$$

$$\models_\pi E_1 | \left(\bigwedge_{j=1}^{card(\alpha)} a_j = \overleftarrow{u_j} \wedge I_\alpha \right) \Rightarrow G \wedge E_2, \quad (17.42)$$

$$\begin{aligned} & \text{for all } 1 \leq j, k \leq card(\alpha), \\ & \text{var}[u_j] \in \vartheta \cup \{a_j\}, a_j \in \alpha, \text{ and} \\ & j \neq k \text{ implies } a_j \neq a_k. \end{aligned} \quad (17.43)$$

Moreover, the induction hypothesis implies that

$$\models_\pi z_2 \underline{\text{sat}} (\vartheta, \alpha) :: (P^R \wedge b, I, \text{false}, \text{true}, E_1),$$

which means that there is a program z'_2 such that

$$\models z'_2 \xrightarrow{(\vartheta, \alpha)} z_2, \quad (17.44)$$

$$ext_\pi[P^R \wedge b, I] \cap cp_\pi[z'_2] \subseteq int_\pi[\text{false}, \text{true}, E_1]. \quad (17.45)$$

Moreover, if z'_1 denotes the program

```

await  $b$  do
   $z'_2$ ;
   $a_1 := u_1$ ;
   $\vdots$ 
   $a_{card(\alpha)} := u_{card(\alpha)}$ 
od,
```

it follows from (17.43) and (17.44) that

$$\models z'_1 \xrightarrow{(\vartheta, \alpha)} z_1.$$

Thus, it is enough to show that

$$ext_\pi[P, R] \cap cp_\pi[z'_1] \subseteq int_\pi[W, G, R|E_2|R]. \quad (17.46)$$

Let

$$\sigma \in ext_\pi[P, R] \cap cp_\pi[z'_1].$$

It is clear from (17.45) that σ converges. If σ deadlocks, (17.41) implies that σ is of the form

$$\langle z'_1, s_1 \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle z'_1, s_l \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle z'_1, s_j \rangle,$$

where $s_j \models_\pi W$. Otherwise, σ is of the form

$$\langle z'_1, s_1 \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle z'_1, s_j \rangle \xrightarrow{i} \langle \epsilon, s_{j+1} \rangle \xrightarrow{e} \dots \xrightarrow{e} \langle \epsilon, s_k \rangle,$$

in which case (17.42) and the transitivity and reflexivity of R imply that $\sigma \in int_\pi[W, G, R|E_2|R]$. Thus, (17.46) has been proved.

- Elimination-Rule:

Assume we have a proof of depth n , whose root

$$Tr_\pi \vdash z \underline{\text{sat}} (\vartheta, \alpha \setminus \{a\}) :: (\exists a : P, \forall \overline{a} : \exists a : R, W, G, E)$$

is deduced from

$$Tr_\pi \vdash z \underline{\text{sat}} (\vartheta, \alpha \cup \{a\}) :: (P, R, W, G, E)$$

by the elimination-rule. Then the induction hypothesis implies that

$$\models_\pi z \underline{\text{sat}} (\vartheta, \alpha \cup \{a\}) :: (P, R, W, G, E),$$

which means that there is a program z' such that

$$\models z' \xrightarrow{(\vartheta, \alpha \cup \{a\})} z, \quad (17.47)$$

$$ext_\pi[P, R] \cap cp_\pi[z'] \subseteq int_\pi[W, G, E]. \quad (17.48)$$

Moreover, it follows from (17.47) that there is a program z'' , which can be got from z' by removing all occurrences of assignments to a , such that

$$\models z'' \xrightarrow{(\vartheta, \alpha)} z.$$

Thus, it is enough to show that

$$ext_\pi[\exists a : P, \forall \overline{a} : \exists a : R] \cap cp_\pi[z''] \subseteq int_\pi[W, G, E]. \quad (17.49)$$

Let

$$\sigma \in ext_\pi[\exists a : P, \forall \overline{a} : \exists a : R] \cap cp_\pi[z''].$$

But then, because of the constraints on the pre- and rely-conditions, we can easily construct a computation

$$\sigma' \in ext_\pi[P, R] \cap cp_\pi[z'],$$

such that for all $1 \leq j \leq \text{len}(\sigma)$,

$$\delta(\sigma_j) \stackrel{\vartheta \cup (\alpha \setminus \{a\})}{=} \delta(\sigma'_j).$$

It follows from (17.48) that $\sigma' \in \text{int}_\pi[W, G, E]$, in which case it also true that $\sigma \in \text{int}_\pi[W, G, E]$, since W , G and E are only constraining the elements of $\vartheta \cup (\alpha \setminus \{a\})$. Thus, (17.49) has been proved.

- Effect-Rule:

Assume we have a proof of depth n , whose root

$$\text{Tr}_\pi \vdash z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, (R \vee G)^\dagger \wedge E)$$

is deduced from

$$\text{Tr}_\pi \vdash z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E),$$

by the effect-rule. Then the induction hypothesis implies that

$$\models_\pi z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E),$$

which means that there is a program z' such that

$$\begin{aligned} & \models z' \stackrel{(\vartheta, \alpha)}{\hookrightarrow} z, \\ & \text{ext}_\pi[P, R] \cap \text{cp}_\pi[z'] \subseteq \text{int}_\pi[W, G, E]. \end{aligned} \tag{17.50}$$

Thus, it is enough to show that

$$\text{ext}_\pi[P, R] \cap \text{cp}_\pi[z'] \subseteq \text{int}_\pi[W, G, (R \vee G)^\dagger \wedge E]. \tag{17.51}$$

Let

$$\sigma \in \text{ext}_\pi[P, R] \cap \text{cp}_\pi[z'].$$

Since any external transition in σ satisfies R and (17.50) implies that any internal transition satisfies G , it is clear that

$$\sigma \in \text{int}_\pi[W, G, (R \vee G)^\dagger \wedge E].$$

Thus, (17.51) has been shown.

- Global-Rule:

Assume we have a proof of depth n , whose root

$$Tr_\pi \vdash z \underline{\text{sat}} (\vartheta \cup \{v\}, \alpha) :: (P, R, W, G \wedge v = \overleftarrow{v}, E)$$

is deduced from

$$Tr_\pi \vdash z \underline{\text{sat}} (\vartheta \setminus \{v\}, \alpha) :: (P, R, W, G, E)$$

by the global-rule. The induction hypothesis implies that

$$\models_\pi z \underline{\text{sat}} (\vartheta \setminus \{v\}, \alpha) :: (P, R, W, G, E),$$

which means that there is a program z' such that

$$\begin{aligned} & \models z' \xrightarrow{(\vartheta, \alpha)} z, \\ & ext_\pi[P, R] \cap cp_\pi[z'] \subseteq int_\pi[W, G, E]. \end{aligned} \tag{17.52}$$

Thus it is enough to show that

$$ext_\pi[P, R] \cap cp_\pi[z'] \subseteq int_\pi[W, G \wedge v = \overleftarrow{v}, E]. \tag{17.53}$$

Let

$$\sigma \in ext_\pi[P, R] \cap cp_\pi[z'],$$

then it follows from (17.52) that $\sigma \in int_\pi[W, G, E]$. Moreover, since the constraints on specified programs imply that the variable v occur in neither z' nor G , it is also clear that $\sigma \in int_\pi[W, G \wedge v = \overleftarrow{v}, E]$. This proves (17.53).

- Auxiliary-Rule:

Assume we have a proof of depth n , whose root

$$Tr_\pi \vdash z \underline{\text{sat}} (\vartheta, \alpha \cup \{a\}) :: (P, R, W, G \wedge a = \overleftarrow{a}, E)$$

is deduced from

$$Tr_\pi \vdash z \underline{\text{sat}} (\vartheta, \alpha \setminus \{a\}) :: (P, R, W, G, E)$$

by the auxiliary-rule. The induction hypothesis implies that

$$\models_\pi z \underline{\text{sat}} (\vartheta, \alpha \setminus \{a\}) :: (P, R, W, G, E),$$

which means that there is a program z' such that

$$\begin{aligned} \models z' \xrightarrow{(\vartheta, \alpha)} z, \\ ext_\pi[P, R] \cap cp_\pi[z'] \subseteq int_\pi[W, G, E]. \end{aligned} \quad (17.54)$$

Thus it is enough to show that

$$ext_\pi[P, R] \cap cp_\pi[z'] \subseteq int_\pi[W, G \wedge a = \overleftarrow{a}, E]. \quad (17.55)$$

Let

$$\sigma \in ext_\pi[P, R] \cap cp_\pi[z'],$$

then it follows from (17.54) that $\sigma \in int_\pi[W, G, E]$. Moreover, since the constraints on specified programs imply that the variable a occur in neither z' nor G , it is also clear that $\sigma \in int_\pi[W, G \wedge a = \overleftarrow{a}, E]$. This proves (17.55).

- Introduction-Rule:

Assume we have a proof of depth n , whose root

$$Tr_\pi \vdash z_2 \underline{\text{sat}} (\vartheta \setminus \alpha, \alpha) :: (P, R, W, G, E)$$

is deduced from

$$Tr_\pi \vdash z_1 \underline{\text{sat}} (\vartheta \cup \alpha, \{\}) :: (P, R, W, G, E)$$

by the introduction-rule. This means that

$$\models z_1 \xrightarrow{(\vartheta, \alpha)} z_2. \quad (17.56)$$

Moreover, the induction hypothesis implies that

$$ext_{\pi}[P, R] \cap cp_{\pi}[z_1] \subseteq int_{\pi}[W, G, E] \quad (17.57)$$

Clearly, (17.56) and (17.57) imply

$$\models_{\pi} z_2 \underline{\text{sat}} (\vartheta \setminus \alpha, \alpha) :: (P, R, W, G, E).$$

end of proof

Chapter 18

Relative Completeness

18.1 Motivation

18.1.1 Completeness

We have already proved that the decomposition-rules in LSP are sound. This means that for any structure π and specified program ψ , if

$$Tr_{\pi} \vdash \psi$$

then

$$\models_{\pi} \psi.$$

The topic of this chapter is to discuss the converse question: in what way can we characterise the applicability of LSP, and in particular the applicability of LSP_B ?

18.1.2 Incompleteness of LSP

Not surprisingly, because of the dependence upon first-order logic, LSP is incomplete. To see this, it is enough to observe that since L_P , the language of Peano arithmetic, is contained in L_{wf} , and given its standard interpretation in any structure π , the set Tr_{π} is not recursively enumerable [Sho67]. Moreover, since for any unary expression E ,

$$\models_{\pi} E$$

if and only if

$$\models_{\pi} \text{skip } \underline{\text{sat}} (var[E], \{\}) :: (\text{true}, \text{true}, \text{false}, \text{true}, E),$$

it is clear that for any structure π , the number of valid specified programs is not recursively enumerable either. Thus, since for any formal system the set of provable formulas is recursively enumerable, it follows that LSP is incomplete.

This means that the best we can hope for is to prove relative completeness (see [Coo78], [Wan78], [Har79], [Apt81] for a more detailed discussion); namely that for any structure π and specified program ψ , if

$$\models_{\pi} \psi$$

then

$$Tr_{\pi} \vdash \psi.$$

It will be shown below that LSP_B satisfies this criterion. The proof depends upon the assumption that for any structure π and binary assertion A in L , there is an assertion A' in L_{wf} , which is valid in π , if and only if A is well-founded in π (see page 85).

18.1.3 Structure of Relative Completeness Proof

The relative completeness proof is split into four main sections. Firstly, three new semantic concepts are introduced — the strongest eff-, wait- and guar-relations characterising respectively the strongest eff-, wait- and guar-conditions for a specified program of auxiliary form (see page 74).

Secondly, we define a function which transforms a certain type of specified program into another specified program, called its historic form. Moreover, it is shown that a specified program of historic form has some very useful properties.

The historic form function is then employed to prove that we can always express the strongest eff-, wait- and guar-conditions for a specified program

of auxiliary form.

Finally, due to this result and the properties satisfied by a specified program of historic form, relative completeness is shown by structural induction on the program component.

18.2 Construction of New Assertions

18.2.1 Decomposition

In the relative-completeness proof it will often be necessary to construct new assertions. For example, to decompose the specified program

$$z_1; z_2 \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E),$$

we must find three new assertions P_1, E_1, E_2 such that

$$\begin{aligned} z_1 \text{ \underline{sat} } (\vartheta, \alpha) &:: (P, R, W, G, E_1), \\ z_2 \text{ \underline{sat} } (\vartheta, \alpha) &:: (P_1, R, W, G, E_2), \end{aligned}$$

are specified programs, and

$$\models_{\pi} (E_1 \Rightarrow P_1) \wedge (E_1 | E_2 \Rightarrow E).$$

This ensures that

$$Tr_{\pi} \vdash_B z_1; z_2 \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E)$$

is deducible from

$$\begin{aligned} Tr_{\pi} \vdash_B z_1 \text{ \underline{sat} } (\vartheta, \alpha) &:: (P, R, W, G, E_1), \\ Tr_{\pi} \vdash_B z_2 \text{ \underline{sat} } (\vartheta, \alpha) &:: (P_1, R, W, G, E_2) \end{aligned}$$

by the consequence- and sequential-rules.

To simplify the arguments we have found it useful to introduce three new concepts at the semantic level called respectively the *strongest eff-relation*, the *strongest wait-relation* and the *strongest guar-relation*. They are all defined with respect to a specified program of auxiliary form.

18.2.2 Closed Relations

Let z denote the program

```
begin
  loc  $y$ ;
   $y := 1$ ;
  while  $y \leq 10$  do
     $x := x + y$ ;
     $y := y + 1$ 
  od
end,
```

and assume we want to characterise the strongest eff-condition E , such that

$$\models_{\pi} z \text{ sat } (\{x\}, \{\}) :: (\text{true}, x = \overline{x}, \text{false}, \text{true}, E).$$

Clearly, for any computation

$$\sigma \in \text{ext}_{\pi}[\text{true}, x = \overline{x}] \cap \text{cp}_{\pi}[z],$$

it must be true that

$$(\delta(\sigma_1), \delta(\sigma_{\text{len}(\sigma)})) \models_{\pi} E.$$

Since

$$\begin{aligned} \delta(\sigma_{\text{len}(\sigma)})(x) &= \delta(\sigma_1)(x) + 55, \\ \delta(\sigma_{\text{len}(\sigma)})(y) &= \delta(\sigma_1)(y) + 11, \end{aligned}$$

the assertion

$$x = \overline{x} + 55 \wedge y = \overline{y} + 11$$

satisfies this criterion. Nevertheless, if we substitute this formula for E in the specified program above, the result is no longer a specified program. The reason is that y is a local variable and therefore restricted from occurring in the specification. In other words, E is required to be ‘closed’ with respect to $\{x\}$, in which case

$$x = \frac{1}{x} + 55$$

is the strongest eff-condition. This motivates the following two definitions:

Definition 20 *A unary relation L on states is closed with respect to a set of variables V , if and only if for all states s_1 and s_2 ,*

- $s_1 \in L$ and $s_1 \stackrel{V}{=} s_2$ imply $s_2 \in L$.

Definition 21 *A binary relation L on states is closed with respect to a set of variables V , if and only if for all states s_1, s_2, s_3, s_4 ,*

- $(s_1, s_2) \in L$, $s_1 \stackrel{V}{=} s_3$, and $s_2 \stackrel{V}{=} s_4$ imply $(s_3, s_4) \in L$.

18.2.3 Strongest Eff-Relation

The strongest eff-relation is supposed to model the semantic equivalent of the strongest eff-condition. Given a specified program

$$z \text{ \underline{sat} } (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, we will use $ef_\pi[z, \vartheta, P, R]$ to denote its strongest eff-relation.

Definition 22 *Given a specified program $z \text{ \underline{sat} } (\vartheta, \{\}) :: (P, R, W, G, E)$ of auxiliary form, let $ef_\pi[z, \vartheta, P, R]$ be the least binary relation on states, closed with respect to ϑ , such that for any terminating computation $\sigma \in ext_\pi[P, R] \cap cp_\pi[z]$,*

- $(\delta(\sigma_1), \delta(\sigma_{len(\sigma)})) \in ef_\pi[z, \vartheta, P, R]$.

18.2.4 Strongest Wait-Relation

Similarly, the strongest wait-relation is intended to characterise the semantic equivalent of the strongest wait-condition. We will use $wa_\pi[z, \vartheta, P, R]$ to denote the strongest wait-condition with respect to a specified program $z \text{ \underline{sat} } (\vartheta, \{\}) :: (P, R, W, G, E)$ of auxiliary form.

Definition 23 *Given a specified program $z \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$ of auxiliary form, let $wa_\pi[z, \vartheta, P, R]$ be the least unary relation on states, closed with respect to ϑ , such that for any deadlocking computation $\sigma \in ext_\pi[P, R] \cap cp_\pi[z]$,*

- $\delta(\sigma_{len(\sigma)}) \in wa_\pi[z, \vartheta, P, R]$.

18.2.5 Strongest Guar-Relation

The strongest guar-relation is the semantic equivalent of the strongest guar-condition. For any specified program $z \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$ of auxiliary form, we will use $gu_\pi[z, \vartheta, P, R]$ to denote its strongest guar-relation.

Definition 24 *Given a specified program $z \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$ of auxiliary form, let $gu_\pi[z, \vartheta, P, R]$ be the least binary, reflexive relation on states, closed with respect to ϑ , such that for any computation $\sigma \in ext_\pi[P, R] \cap cp_\pi[z]$ and $1 \leq l < len(\sigma)$,*

- if $\lambda(\sigma_l) = i$ then $(\delta(\sigma_l), \delta(\sigma_{l+1})) \in gu_\pi[z, \vartheta, P, R]$.

18.2.6 Conversion Function

We will later prove that for any specified program

$$z \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, the strongest eff-, wait- and guar-relations can always be expressed in L. In other words, we can always find two binary assertions A_1, A_2 , and one unary assertion A_3 , such that for all states s_1, s_2 :

- $(s_1, s_2) \models_\pi A_1$ if and only if $(s_1, s_2) \in ef_\pi[z, \vartheta, P, R]$,
- $(s_1, s_2) \models_\pi A_2$ if and only if $(s_1, s_2) \in gu_\pi[z, \vartheta, P, R]$,
- $s_1 \models_\pi A_3$ if and only if $s_1 \in wa_\pi[z, \vartheta, P, R]$.

To make it easy to convert a strongest eff-, wait- or guar-relation into an assertion, we have found it helpful to define a specific conversion function $\wp$, which, when applied to a strongest eff-, wait- or guar-relation, returns a corresponding assertion in L.

18.3 Historic Form

18.3.1 Enrichment of the Global State

At any stage during the execution of a program, the set of possible internal transitions is a function of the current configuration. Since specifications are restricted to constrain only the global state transitions, the provability of many programs depends upon the possibility to enrich the global state with auxiliary variables, and use them to encode information about the local state and the program counter into the specification. Let for example z denote the program

$$x := x + 1$$

then

$$\begin{aligned} \wp(gu_\pi[z, \{x\}, \text{true}, x \geq \overline{x}]) &= x = \overline{x} \vee x = \overline{x} + 1, \\ \wp(wa_\pi[z, \{x\}, \text{true}, x \geq \overline{x}]) &= \text{false}, \\ \wp(ef_\pi[z, \{x\}, \text{true}, x \geq \overline{x}]) &= x > \overline{x}. \end{aligned}$$

Thus, if

$$\models_\pi z \text{ sat } (\{x\}, \{\}) :: (\text{true}, x \geq \overline{x}, W, G, E),$$

it follows that

$$\begin{aligned} \models_\pi x = \overline{x} \vee x = \overline{x} + 1 &\Rightarrow G, \\ \models_\pi \text{false} &\Rightarrow W, \\ \models_\pi x > \overline{x} &\Rightarrow E. \end{aligned}$$

Furthermore, it is clear that

$$\models_\pi \{z \parallel z\} \text{ sat } (\{x\}, \{\}) :: (\text{true}, x = \overline{x}, \text{false}, x = \overline{x} \vee x = \overline{x} + 1, x = \overline{x} + 2).$$

Unfortunately, without taking advantage of auxiliary variables, this is not provable in LSP. The closest we can get is

$$Tr_\pi \vdash \{z \parallel z\} \underline{\text{sat}} (\{x\}, \{\}) :: (\text{true}, x = \underline{x}, \text{false}, x = \underline{x} \vee x = \underline{x} + 1, x > \underline{x}).$$

Thus, LSP without the auxiliary-variable rules is incomplete not only because the first-order logic is incomplete. The problem is not that the strongest guar-, wait- and eff-conditions cannot be expressed in L , because, as we have seen in the example above, they can, but that the global state is not rich enough.

18.3.2 History Variables

To make it easy to deal with auxiliary structure in the completeness proof, we have found it useful to introduce a function which transforms a specified program of auxiliary form

$$\{z_1 \parallel z_2\} \underline{\text{sat}} (\vartheta, \{\}) :: (P, R, W, G, E),$$

whose program component has a parallel-statement as main construct, into another specified program, called its *historic form*, which is sufficiently expressive for our purposes.

The basic idea is to introduce auxiliary history variables; one for each variable in ϑ , to record the use of the global state, one for each local variable to record changes to the local state, and one history variable to record if an update was due to the overall environment, due to z_1 or due to z_2 . For example, if z denotes the program

$$\text{await true do } x := x + 1 \text{ od},$$

then the specified program

$$\{z \parallel z\} \underline{\text{sat}} (\{x\}, \{\}) :: (\text{true}, x = \underline{x}, \text{false}, \text{true}, x = \underline{x} + 2)$$

is transformed into a specified program of the form

$$\{z' \parallel z''\} \underline{\text{sat}} (\vartheta', \{\}) :: (P', R', \text{false}, \text{true}, x = \underline{x} + 2),$$

where

- ϑ' is a set $\{x, h_x, h\}$ such that h is of sort seq of $\mathbb{N}$, and if x is of sort Σ , then h_x is of sort seq of Σ ,
- z' and z'' denote respectively

```

await true do
   $h_x := [h_x(1) + 1] \curvearrowright h_x;$ 
   $h := [1] \curvearrowright h;$ 
   $x := x + 1$ 
od,
```

and

```

await true do
   $h_x := [h_x(1) + 1] \curvearrowright h_x;$ 
   $h := [2] \curvearrowright h;$ 
   $x := x + 1$ 
od,
```

- P' represents

$$h_x = [x] \wedge h = [3],$$

- while R' stands for

$$(x = \overleftarrow{x} \wedge h = [3] \curvearrowright \overleftarrow{h} \wedge h_x = [x] \curvearrowright \overleftarrow{h_x})^*.$$

Clearly, we can now use h to determine if the n 'th state change was due to z' , in which case $h(\text{len } h - n) = 1$, to z'' , in which case $h(\text{len } h - n) = 2$, or to the environment, in which case $h(\text{len } h - n) = 3$.

Moreover, if $m \circ h$ stands for the number of occurrences of m in h , and

$$\begin{aligned}
G_1 &= 1 \circ \overleftarrow{h} = 0 \wedge x = \overleftarrow{x} + 1 \wedge h_x = [x] \curvearrowright \overleftarrow{h_x} \wedge h = [1] \curvearrowright \overleftarrow{h}, \\
G_2 &= 2 \circ \overleftarrow{h} = 0 \wedge x = \overleftarrow{x} + 1 \wedge h_x = [x] \curvearrowright \overleftarrow{h_x} \wedge h = [2] \curvearrowright \overleftarrow{h}, \\
R_1 &= (R' \vee G_2)^\dagger, \\
R_2 &= (R' \vee G_1)^\dagger, \\
E_1 &= 1 \circ h = 1 \wedge ((2 \circ h = 0 \wedge x = \overleftarrow{x} + 1) \vee \\
&\quad (2 \circ h = 1 \wedge x = \overleftarrow{x} + 2)), \\
E_2 &= 2 \circ h = 1 \wedge ((1 \circ h = 0 \wedge x = \overleftarrow{x} + 1) \vee \\
&\quad (1 \circ h = 1 \wedge x = \overleftarrow{x} + 2)),
\end{aligned}$$

it follows by the assignment-rule that

$$\begin{aligned}
Tr_\pi \vdash z' \text{ \underline{sat} } (\vartheta', \{\}) &:: (P', R_1, \text{false}, \text{true} \wedge R_2, E_1), \\
Tr_\pi \vdash z'' \text{ \underline{sat} } (\vartheta', \{\}) &:: (P', R_2, \text{false}, \text{true} \wedge R_1, E_2),
\end{aligned}$$

which by the parallel-rule give

$$Tr_\pi \vdash \{z' \parallel z''\} \text{ \underline{sat} } (\vartheta', \{\}) :: (P', R_1 \wedge R_2, \text{false}, \text{true}, E_1 \wedge E_2),$$

in which case

$$Tr_\pi \vdash \{z \parallel z\} \text{ \underline{sat} } (\{x\}, \{\}) :: (\text{true}, x = \overleftarrow{x}, \text{false}, \text{true}, x = \overleftarrow{x} + 2)$$

follows by the elimination-, introduction- and consequence-rules.

18.3.3 Component Functions

To simplify the definition of the historic-form function, we will first define four component functions: $\bar{h}$ which generates the set of history variables, $\mathfrak{S}$ which extends a program with auxiliary structure, Pre which initialises the auxiliary structure, and Rel which constrains the external updates of the auxiliary structure.

- If V is an ordered set of m variables $\{v_1, \dots, v_m\}$, then $\bar{h}[V]$ denotes an ordered set of $m + 1$ variables

$$\{h_{v_1}, \dots, h_{v_m}, h\},$$

disjoint from V , such that h is of sort seq of $\mathbb{N}$, and for all $1 \leq j \leq m$, if v_j is of sort Σ_{v_j} , then h_{v_j} is of sort seq of Σ_{v_j} .

In other words, h is a function which returns a new history variable of appropriate sort for each element of the argument, plus a new history variable to record the origin of updates.

- If z is a program, V is an ordered set of m variables $\{v_1, \dots, v_m\}$, $h[V] = \{h_{v_1}, \dots, h_{v_m}, h\}$ and $n \in \{1, 2\}$, then $\mathfrak{F}[z, V, n]$ denotes the program that can be obtained from z by replacing any await-statement of the form

```

await  $b$  do
   $x_1 := r_{x_1};$ 
   $\vdots$ 
   $x_w := r_{x_w}$ 
od,
```

not contained in the body of an await-statement¹, with an await-statement of the form

¹Since the historic-form function is only defined for specified programs of auxiliary form, this means that the Boolean test b is of the form true and for all $1 \leq j < k \leq w$, $x_j \neq x_k$.

```

await true do
   $h_{v_1} := [t_{v_1}] \curvearrowright h_{v_1};$ 
   $\vdots$ 
   $h_{v_m} := [t_{v_m}] \curvearrowright h_{v_m};$ 
   $h := [n] \curvearrowright h;$ 
   $x_1 := r_{x_1};$ 
   $\vdots$ 
   $x_w := r_{x_w}$ 
od,

where for all  $1 \leq j \leq m$ ,
   $t_{v_j} = r_{x_k}(h_{v_j}(1)/x_k)$  if there is a  $1 \leq k \leq w$  such that  $v_j = x_k$ ,
  and  $t_{v_j} = h_{v_j}(1)$  otherwise,

```

and by replacing any other await-statement of the form

```

await  $b$  do
   $z'$ 
od,

```

not contained in the body of an await-statement, with an await-statement of the form

```

await  $b$  do
   $\mathfrak{S}[z', V, n];$ 
   $h_{v_1} := [h_{v_1}(1)] \curvearrowright h_{v_1};$ 
   $\vdots$ 
   $h_{v_m} := [h_{v_m}(1)] \curvearrowright h_{v_m};$ 
   $h := [n] \curvearrowright h$ 
od.

```

This means that the history variables are updated only in connection with await-statements. The reason is first of all that only assignment- and await-statements can access the global state. Secondly, since the historic-form

function is defined only for specified programs of auxiliary form, assignments to local variables can only take place inside await-statements.

Observe, that no variable v_1 in V is allowed to occur on the right-hand side of a ‘new’ assignment-statement unless there is an assignment-statement in z of the form $v_2 := r$, where $v_1 \neq v_2$ and v_1 occurs in r . Thus, the new auxiliary structure does not depend upon the auxiliary structure of z .

- If P is an unary assertion, $V \cup U$ is an ordered set of m variables $\{v_1, \dots, v_m\}$, $V \cap U = \{\}$ and $\hbar[V \cup U] = \{h_{v_1}, \dots, h_{v_m}, h\}$, then $Pre[P, V, U]$ denotes the assertion

$$P \wedge h = [3] \wedge \bigwedge_{v \in V} h_v = [v] \wedge \bigwedge_{u \in U} \text{len } h_u = 1.$$

The object of the function Pre is to initialise the history variables. The set V contains the global variables, while the set of local variables is denoted by U .

- If R is a binary assertion, $V \cup U$ is an ordered set of m variables $\{v_1, \dots, v_m\}$, $V \cap U = \{\}$ and $\hbar[V \cup U] = \{h_{v_1}, \dots, h_{v_m}, h\}$, then $Rel[R, V, U]$, $Rel_1[R, V, U]$ and $Rel_2[R, V, U]$ denote respectively the assertions:

$$(R \wedge h = [3] \curvearrowright \overleftarrow{h} \wedge \bigwedge_{v \in V} h_v = [v] \curvearrowright \overleftarrow{h_v} \wedge \bigwedge_{u \in U} h_u = [\overleftarrow{h_u}(1)] \curvearrowright \overleftarrow{h_u})^*,$$

$$(R \wedge (h = 3 \curvearrowright \overleftarrow{h} \vee h = 2 \curvearrowright \overleftarrow{h}) \wedge \bigwedge_{v \in V} h_v = [v] \curvearrowright \overleftarrow{h_v} \wedge \bigwedge_{u \in U} h_u = [\overleftarrow{h_u}(1)] \curvearrowright \overleftarrow{h_u})^*,$$

$$(R \wedge (h = 3 \curvearrowright \overleftarrow{h} \vee h = 1 \curvearrowright \overleftarrow{h}) \wedge \bigwedge_{v \in V} h_v = [v] \curvearrowright \overleftarrow{h_v} \wedge \bigwedge_{u \in U} h_u = [\overleftarrow{h_u}(1)] \curvearrowright \overleftarrow{h_u})^*.$$

Again, the set V contains the global variables, while the set of local variables is denoted by U . Clearly, the Rel functions record any update due

to the relevant environment. (Remember that the environment is restricted from updating local variables.) Observe, that $\mathfrak{S}$, Pre and the Rel functions together ensure that the value of any global variable and any ‘active’ local variable always equals the first element of its history variable.

18.3.4 Historic-Form Function

The historic form function can then be defined:

Definition 25 *Given a specified program of auxiliary form*

$$\{z_1 \parallel z_2\} \underline{\text{sat}} (\vartheta, \{\}) :: (P, R, W, G, E),$$

whose program component has a parallel-statement as main construct, then

$$\{\mathfrak{S}[z_1, V, 1] \parallel \mathfrak{S}[z_2, V, 2]\} \underline{\text{sat}} (\vartheta \cup h[V], \{\}) :: \\ (Pre[P, \vartheta, U], Rel[R, \vartheta, U], W, G, E),$$

where $U = \text{var}[\{z_1 \parallel z_2\}] \setminus \vartheta$ and $V = \vartheta \cup U$, is its historic form.

18.3.5 Historic-Form Proposition

To show that the historic-form function has the desired properties, we will prove four propositions. The first one is rather trivial. Basically, it asserts that the result of applying the historic-form function is a specified program, that the historic-form function preserves validity, and that the history variables introduced by the historic-form function are used in such a way that they do not depend upon auxiliary structure already introduced.

Proposition 9 *Given a specified program*

$$z \underline{\text{sat}} (\vartheta, \{\}) :: \psi$$

of auxiliary form such that z ’s main construct is a parallel-statement, and assume that

$$z' \underline{\text{sat}} (\vartheta \cup \vartheta', \{\}) :: \psi',$$

where $\vartheta \cap \vartheta' = \{\}$, is its historic form, then

- $z' \underline{\text{sat}} (\vartheta \cup \vartheta', \{\}) :: \psi' \text{ is a specified program,}$
- $\models_{\pi} z \underline{\text{sat}} (\vartheta, \{\}) :: \psi \text{ implies } \models_{\pi} z' \underline{\text{sat}} (\vartheta \cup \vartheta', \{\}) :: \psi',$
- $\models z \xrightarrow{(\vartheta \setminus \alpha, \alpha)} z'' \text{ implies } \models z' \xrightarrow{(\vartheta \setminus \alpha, \alpha \cup \vartheta')} z''.$

Proof: The first result follows from the fact that $Rel[R, \vartheta, U]$ by definition is both reflexive and transitive, and that any new variable introduced by $\mathfrak{S}$, Pre or Rel is an element of ϑ' and is distinct from any local variable in z .

Moreover, since

$$\begin{aligned} \models_{\pi} Pre[P, \vartheta, U] &\Rightarrow P, \\ \models_{\pi} Rel[R, \vartheta, U] &\Rightarrow R \end{aligned}$$

the elements of ϑ' do not occur in W , G and E , and the elements of ϑ' has no influence on the elements of ϑ , the second result is also correct.

The third result follows from the fact that no variable v_1 in ϑ is allowed to occur on the right-hand side of a ‘new’ assignment-statement unless there is an assignment-statement in z of the form $v_2 := r$, where $v_1 \neq v_2$ and v_1 occurs in r .

end of proof

18.3.6 Decomposition Propositions

The object of the three remaining propositions is to show that the historic-form function gives us the necessary expressive power to decompose any valid specified program whose program component has a parallel-statement as main construct. Let

$$\{z_1 \parallel z_2\} \underline{\text{sat}} (\vartheta, \{\}) :: (P, R, W, G, E)$$

be a specified program of auxiliary form, such that R respects

$$hid[\{z_1 \parallel z_2\}] \cap \vartheta$$

and

$$U = var[\{z_1 \parallel z_2\}] \setminus \vartheta.$$

Moreover, let

$$\{z'_1 \parallel z'_2\} \underline{\text{sat}} (\vartheta', \{\}) :: (P', R', W, G, E)$$

be its historic form, and assume there are assertions G_1 and G_2 such that

$$\begin{aligned} G_1 &= \wp(gu_\pi[z'_1, \vartheta', Pre[\mathbf{true}, \vartheta, U], Rel_2[\mathbf{true}, \vartheta, U]]), \\ G_2 &= \wp(gu_\pi[z'_2, \vartheta', Pre[\mathbf{true}, \vartheta, U], Rel_1[\mathbf{true}, \vartheta, U]]), \\ R_1 &= (R' \vee G_2)^\dagger, \\ R_2 &= (R' \vee G_1)^\dagger. \end{aligned}$$

Clearly, G_1 defines z'_1 's set of possible state changes as a function of the history, while G_2 characterises a similar relationship with respect to z'_2 .

Proposition 10 *Given the context above, and two computations*

$$\begin{aligned} \sigma &\in ext_\pi[P', R_1] \cap cp_\pi[z'_1], \\ \sigma' &\in ext_\pi[P', R_2] \cap cp_\pi[z'_2], \end{aligned}$$

$1 \leq j \leq len(\sigma)$ and $1 \leq k \leq len(\sigma')$, such that $\delta(\sigma_j) \stackrel{\vartheta'}{=} \delta(\sigma'_k)$, then there are computations

$$\begin{aligned} \sigma'' &\in ext_\pi[P', R_1] \cap cp_\pi[z'_1], \\ \sigma''' &\in ext_\pi[P', R_2] \cap cp_\pi[z'_2] \end{aligned}$$

and $1 \leq l \leq \{len(\sigma''), len(\sigma''')\}$, such that $\sigma''(1, \dots, l) \bullet \sigma'''(1, \dots, l)$, $\sigma_j \stackrel{\vartheta'}{=} \sigma''_l$ and $\sigma'_k \stackrel{\vartheta'}{=} \sigma'''_l$.

Proof: Assume that any external transition in σ satisfies R' or G_2 . σ can easily be transformed into such a computation by splitting external transitions, if this is not the case. Due to the information stored in the history variables no external transition of σ can satisfy both R' and G_2 unless it leaves ϑ' unchanged. Similarly, we may assume that any external transition in σ' satisfies R' or G_1 .

The proof is by induction on j . The base case $j = 1$ is trivial. Assume the proposition is correct for $j \leq n$. We will prove that the proposition is correct for $j = n + 1$. There are four cases to consider:

- If there is an $1 \leq l < j$ such that $\lambda(\sigma_l) = e$ and $\delta(\sigma_l) \stackrel{\vartheta'}{=} \delta(\sigma_{l+1})$:
 - This means that there is a computation

$$\sigma''' \in \text{ext}_\pi[P', R_1] \cap \text{cp}_\pi[z'_1],$$

such that $\sigma'''_{j-1} \stackrel{\vartheta'}{=} \sigma_j$. Thus, the proposition follows by the induction hypothesis.

- Else if there is an $1 \leq l < j$ such that $\lambda(\sigma_l) = i$, $\delta(\sigma_l) = \delta(\sigma_{l+1})$ and for all $l < m < j$, $\lambda(\sigma_l) = e$:
 - Since the internal transition is independent of the global state, and no external transition can change the local state, it follows that

$$\begin{aligned} \sigma(1, \dots, l) &\xrightarrow{e} \langle \tau(\sigma_l), \delta(\sigma_{l+2}) \rangle \xrightarrow{e} \dots \\ &\xrightarrow{e} \langle \tau(\sigma_l), \delta(\sigma_j) \rangle \xrightarrow{i} \sigma_j \end{aligned}$$

is a computation in $\text{ext}_\pi[P', R_1] \cap \text{cp}_\pi[z'_1]$, in which case the proposition follows easily by the induction hypothesis.

- Else if there is an $1 \leq l < j$, such that for all $l \leq m < j$, $\lambda(\sigma_m) = i$, and $l = 1$ or $\lambda(\sigma_{l-1}) = e$:
 - Due to the information stored in the history variables it follows that there is a $1 \leq p < k$, such that $\delta(\sigma_l) \stackrel{\vartheta'}{=} \delta(\sigma'_p)$. Since the value of any active local variable is determined by the first value of its history variable, the proposition follows easily by the induction hypothesis.
- Else:
 - Due to the information stored in the history variables it follows that there are $1 \leq l < j$ and $1 \leq m < k$ such that $\delta(\sigma_l) \stackrel{\vartheta'}{=} \delta(\sigma'_m)$ and $(\delta(\sigma_l), \delta(\sigma_j)) \models_\pi R'$ and $(\delta(\sigma'_m), \delta(\sigma'_k)) \models_\pi R'$, in which case the proposition follows by the induction hypothesis.

end of proof

Proposition 11 *Given the context above, a computation*

$$\sigma \in \text{ext}_\pi[P', R_1] \cap \text{cp}[z'_1]$$

and $1 \leq j \leq \text{len}(\sigma)$, then there are computations

$$\begin{aligned} \sigma' &\in \text{ext}_\pi[P', R_1] \cap \text{cp}[z'_1], \\ \sigma'' &\in \text{ext}_\pi[P', R_2] \cap \text{cp}[z'_2], \end{aligned}$$

and $1 \leq k \leq \min\{\text{len}(\sigma'), \text{len}(\sigma'')\}$, such that $\sigma'(1, \dots, k) \bullet \sigma''(1, \dots, k)$ and $\sigma_j \stackrel{\vartheta'}{=} \sigma'_k$.

Proof: In the same way as above, we may assume that any external transition in σ satisfies R' or G_2 .

The proof is by induction on j . The base case $j = 1$ is trivial. Assume the proposition is correct for all $j \leq n$. We will prove that the proposition is correct for $j = n + 1$. There are three cases to consider:

- If $\lambda(\sigma_n) = i$:
 - The proposition follows by the induction hypotheses, since the value of any active local variable is determined by the first value of its history variable.
- Else if $(\delta(\sigma_n), \delta(\sigma_{n+1})) \models_\pi R'$:
 - Follows easily by the induction hypothesis, since R' respects

$$\text{hid}[\{z'_1 \parallel z'_2\}] \cap \vartheta'.$$

- Else:
 - This means that $\lambda(\sigma_n) = e$ and $(\delta(\sigma_n), \delta(\sigma_{n+1})) \models_\pi G_2$. By definition, there is a computation

$$\sigma' \in \text{ext}_\pi[\text{Pre}[\text{true}, \vartheta', U], \text{Rel}_2[\text{true}, \vartheta', U]] \cap \text{cp}_\pi[z'_2]$$

and $1 \leq l \leq \text{len}(\sigma')$, such that $\delta(\sigma'_l) \stackrel{\vartheta'}{=} \delta(\sigma_{n+1})$. But then, due to the information stored in the history variables, there is a computation

$$\sigma'' \in \text{ext}_\pi[P', R_2] \cap \text{cp}_\pi[z'_2],$$

and $1 \leq l \leq \text{len}(\sigma'')$, such that $\delta(\sigma''_l) \stackrel{\vartheta'}{=} \delta(\sigma_{n+1})$, in which case the desired result follows from proposition 10 on page 193.

end of proof

Proposition 12 *Given the context above, if there is a diverging computation*

$$\sigma \in \text{ext}_\pi[P', R_1] \cap \text{cp}[z'_1],$$

then there is a diverging computation

$$\sigma' \in \text{ext}_\pi[P', R'] \cap \text{cp}[\{z'_1 \parallel z'_2\}].$$

Proof: There are three cases to consider:

- There is a $j \geq 1$, such that for all $k \geq j$, $\lambda(\sigma_k) = e$ implies

$$(\delta(\sigma_k), \delta(\sigma_{k+1})) \models_\pi R' :$$

- Proposition 11 on page 195 implies that there are

$$\begin{aligned} \sigma' &\in \text{ext}_\pi[P', R_1] \cap \text{cp}_\pi[z'_1], \\ \sigma'' &\in \text{ext}_\pi[P', R_2] \cap \text{cp}_\pi[z'_2] \end{aligned}$$

and $1 \leq k \leq \min\{\text{len}(\sigma'), \text{len}(\sigma'')\}$, such that

$$\sigma'(1, \dots, k) \bullet \sigma''(1, \dots, k)$$

and $\sigma'_k \stackrel{\vartheta'}{=} \sigma_j$. Since the value of any active local variable is determined by the first value of its history variable, and since R'

respects $hid[\{z'_1 \parallel z'_2\}] \cap \vartheta'$, it is straightforward to construct an infinite computation

$$\sigma' \in ext_\pi[P', R'] \cap cp_\pi[\{z'_1 \parallel z'_2\}].$$

- There is a $j \geq 1$, such that for all $k \geq j$, $\lambda(\sigma_k) = i$ implies $\delta(\sigma_k) = \delta(\sigma_{k+1})$:

Since after the first $j - 1$ transitions no internal transition depends upon the global state, we can easily construct an infinite computation

$$\sigma' \in ext_\pi[P', R'] \cap cp_\pi[\{z'_1 \parallel z'_2\}]$$

by an argument similar to the one above.

- Else:
 - Since any internal transition which leaves the state unchanged can be moved, and any external transition which leaves ϑ' unchanged can be removed, we may assume that there is an infinite sequence of natural numbers

$$n_1 < m_1 < n_2 < m_2 < \dots < n_k < m_k < \dots$$

such that for all j, k :

$$\begin{aligned} \delta(\sigma_{n_j}) &\stackrel{\vartheta'}{\neq} \delta(\sigma_{m_j}), \\ \delta(\sigma_{m_j}) &\stackrel{\vartheta'}{\neq} \delta(\sigma_{n_{j+1}}), \\ n_j \leq k < m_j &\Rightarrow (\delta(\sigma_k), \delta(\sigma_{k+1})) \models_\pi G_1 \vee R', \\ m_j \leq k < n_{j+1} &\Rightarrow (\delta(\sigma_k), \delta(\sigma_{k+1})) \models_\pi G_2. \end{aligned}$$

For any $j \geq 1$, let T_j be the set of all tuples of the form (σ', σ'', k) such that there are

$$\begin{aligned} \sigma' &\in ext_\pi[P', R_1] \cap cp_\pi[z'_1], \\ \sigma'' &\in ext_\pi[P', R_2] \cap cp_\pi[z'_2] \end{aligned}$$

and $1 \leq k \leq \min\{\text{len}(\sigma'), \text{len}(\sigma'')\}$ which satisfy $\sigma'(1, \dots, k) \bullet \sigma''(1, \dots, k)$ and $\delta(\sigma_{m_j}) \stackrel{\vartheta'}{=} \delta(\sigma'_k)$.

Clearly, for all $1 \leq p < j$, if $(\sigma', \sigma'', k) \in T_j$, then there is an l such that $(\sigma', \sigma'', l) \in T_p$.

Moreover, since the value of any active local variable is determined by the first value of its history variable, if $(\sigma', \sigma'', k), (\sigma''', \sigma''', p) \in T_j$, $\tau(\sigma'_k) = \tau(\sigma'''_p)$, $\tau(\sigma''_k) = \tau(\sigma'''_p)$, and there are $o > j$, extensions σ'_e of $\sigma'(1, \dots, k)$, σ''_e of $\sigma''(1, \dots, k)$ and k_e , such that $(\sigma'_e, \sigma''_e, k_e) \in T_o$, then there are extensions σ'''_e of $\sigma'''(1, \dots, k)$ and σ''''_e of $\sigma''''(1, \dots, k)$ and p_e such that $(\sigma'''_e, \sigma''''_e, p_e) \in T_o$.

Let $Z_j = \{(\tau(\sigma'_k), \tau(\sigma''_k)) \mid (\sigma', \sigma'', k) \in T_j\}$. It follows from proposition 11 on page 195 that Z_j is nonempty, and from the definition of a computation that Z_j is finite.

But then, since a computation is not required to satisfy any fairness constraint, we can easily construct an infinite computation

$$\sigma' \in \text{ext}_\pi[P', R'] \cap \text{cp}_\pi[\{z'_1 \parallel z'_2\}].$$

end of proof

18.4 Expressiveness Proposition

The next proposition shows that we can always express the strongest wait-, guar- and eff-relations.

Proposition 13 *Given a specified program*

$$z \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, then $\text{ef}_\pi[z, \vartheta, P, R]$, $\text{wa}_\pi[z, \vartheta, P, R]$ and $\text{gu}_\pi[z, \vartheta, P, R]$ are expressible in L .

Proof: For any specified program of auxiliary form

$$z \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E),$$

there is a unique program z' such that $\models z \hookrightarrow z'$. We will prove the above proposition by structural induction on this program. Since the environment respects $hid[z] \cap \vartheta$, it follows that for all $v \in hid[z] \cap \vartheta$:

$$\begin{aligned} ef_\pi[z, \vartheta, P, R] &= ef_\pi[z, \vartheta, P, R \wedge v = \overleftarrow{v}], \\ wa_\pi[z, \vartheta, P, R] &= wa_\pi[z, \vartheta, P, R \wedge v = \overleftarrow{v}], \\ gu_\pi[z, \vartheta, P, R] &= gu_\pi[z, \vartheta, P, R \wedge v = \overleftarrow{v}]. \end{aligned}$$

In each case below it is therefore assumed that R respects $hid[z] \cap \vartheta$. The base-cases are the skip- and assignment-statements.

- Skip: Given a specified program

$$z_1 \text{ \underline{sat} } (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, and a program z'_1 of the form

$$\text{skip},$$

such that $\models z_1 \hookrightarrow z'_1$. Clearly, $ef_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\overleftarrow{P} \wedge R,$$

$wa_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\text{false},$$

while $gu_\pi[z_1, \vartheta, P, R]$ is characterised by

$$I.$$

- Assignment: Given a specified program

$$z_1 \text{ \underline{sat} } (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, a program z'_1 of the form

$$v := r,$$

such that $\models z_1 \hookrightarrow z'_1$, and assume α characterises the set of auxiliary variables in z_1 . This means that z_1 is of the form

```

await true do
   $a_1 := u_1;$ 
   $\vdots$ 
   $a_m := u_m;$ 
   $v := r$ 
od,
```

where for all $1 \leq j, k \leq m$, $\text{var}[u_j] \subseteq (\vartheta \setminus \alpha) \cup \{a_j\}$, $j \neq k \Rightarrow a_j \neq a_k$ and $a_j \in \alpha$. Clearly, $ef_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\overleftarrow{P} \wedge R | (v = \overleftarrow{r} \wedge I_{\{v\} \cup \alpha} \wedge \bigwedge_{j=1}^m a_j = \overleftarrow{u_j}) | R,$$

$wa_\pi[z_1, \vartheta, P, R]$ is characterised by

false,

while $gu_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\overleftarrow{P^R} \wedge v = \overleftarrow{r} \wedge I_{\{v\} \cup \alpha} \wedge \bigwedge_{j=1}^m a_j = \overleftarrow{u_j}.$$

- Block: Given a specified program

$$z_1 \underline{\text{sat}} (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, and a program z'_1 of the form

```

begin loc  $v_1, \dots, v_m; z'_2$  end,
```

such that $\models z_1 \hookrightarrow z'_1$. This means that z_1 is of the form

```

begin loc  $v_1, \dots, v_m; z_2$  end,
```

where $\models z_2 \hookrightarrow z'_2$. Thus, by the induction hypothesis it follows that $ef_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\begin{aligned} & \exists \overleftarrow{v_1}, \dots, \overleftarrow{v_m}, v_1, \dots, v_m : \\ & \wp(ef_\pi[z_2, \vartheta \cup \bigcup_{j=1}^m \{v_j\}, P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}]), \end{aligned}$$

$wa_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\exists v_1, \dots, v_m : \wp(wa_\pi[z_2, \vartheta \cup \bigcup_{j=1}^m \{v_j\}, P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}]),$$

while $gu_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\begin{aligned} & \exists \overleftarrow{v_1}, \dots, \overleftarrow{v_m}, v_1, \dots, v_m : \\ & \wp(gu_\pi[z_2, \vartheta \cup \bigcup_{j=1}^m \{v_j\}, P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}]) \end{aligned}$$

- Composition: Given a specified program

$$z_1 \text{ \underline{sat} } (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, and a program z'_1 of the form

$$z'_2; z'_3,$$

such that $\models z_1 \hookrightarrow z'_1$. This means that z_1 is of the form

$$z_2; z_3,$$

where $\models z_2 \hookrightarrow z'_2$ and $\models z_3 \hookrightarrow z'_3$. Thus, by the induction hypothesis it follows that $ef_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\begin{aligned} & \wp(ef_\pi[z_2, \vartheta, P, R]) | \\ & \wp(ef_\pi[z_3, \vartheta, \exists \overleftarrow{\vartheta} : \wp(ef_\pi[z_2, \vartheta, P, R]), R]), \end{aligned}$$

$wa_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\wp(wa_\pi[z_2, \vartheta, P, R]) \vee \wp(wa_\pi[z_3, \vartheta, \exists \overleftarrow{\vartheta} : \wp(ef_\pi[z_2, \vartheta, P, R]), R]),$$

while $gu_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\wp(gu_\pi[z_2, \vartheta, P, R]) \vee \wp(gu_\pi[z_3, \vartheta, \exists \overleftarrow{\vartheta} : \wp(ef_\pi[z_2, \vartheta, P, R]), R).$$

- If: Given a specified program

$$z_1 \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, and a program z'_1 of the form

$$\text{if } b \text{ then } z'_2 \text{ else } z'_3 \text{ fi,}$$

such that $\models z_1 \hookrightarrow z'_1$. This means that z_1 is of the form

$$\text{if } b \text{ then } z_2 \text{ else } z_3 \text{ fi,}$$

where $\models z_2 \hookrightarrow z'_2$ and $\models z_3 \hookrightarrow z'_3$. Thus, by the induction hypothesis it follows that $ef_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\wp(ef_\pi[z_2, \vartheta, P \wedge \neg b, R]) \vee \wp(ef_\pi[z_3, \vartheta, P \wedge b, R]),$$

$wa_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\wp(wa_\pi[z_2, \vartheta, P \wedge \neg b, R]) \vee \wp(wa_\pi[z_3, \vartheta, P \wedge b, R]),$$

while $gu_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\wp(gu_\pi[z_2, \vartheta, P \wedge \neg b, R]) \vee \wp(gu_\pi[z_3, \vartheta, P \wedge b, R]).$$

- While: Given a specified program

$$z_1 \underline{\text{sat}} (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, and a program z'_1 of the form

$$\text{while } b \text{ do } z'_2 \text{ od,}$$

such that $\models z_1 \hookrightarrow z'_1$. This means that z_1 is of the form

$$\text{while } b \text{ do } z_2 \text{ od,}$$

where $\models z_2 \hookrightarrow z'_2$. Thus, by the induction hypothesis, if Z denotes

$$\wp(\text{ef}_\pi[z_2, \vartheta, b, R]),$$

it follows that $\text{ef}_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\overline{P} \wedge (Z^\dagger \vee R) \wedge \neg b,$$

$\text{wa}_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\wp(\text{wa}_\pi[z_2, \vartheta, P^Z \wedge b, R]),$$

while $\text{gu}_\pi[z_1, P, R]$ is characterised by

$$\wp(\text{gu}_\pi[z_2, \vartheta, P^Z \wedge b, R]).$$

- Parallel: Given a specified program

$$z_1 \underline{\text{sat}} (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, and a program z'_1 of the form

$$\{z'_2 \parallel z'_3\},$$

such that $\models z_1 \hookrightarrow z'_1$. This means that z_1 is of the form

$$\{z_2 \parallel z_3\},$$

where $\models z_2 \hookrightarrow z'_2$ and $\models z_3 \hookrightarrow z'_3$.

Let

$$z''_1 \text{ sat } (\vartheta'', \{\}) :: (P'', R'', W, G, E)$$

be the historic form. This means that z''_1 is of the form

$$\{z''_2 \parallel z''_3\},$$

where $\models z''_2 \hookrightarrow z'_2$ and $\models z''_3 \hookrightarrow z'_3$. Let $U = \text{var}[z_1] \setminus \vartheta$, then by the induction hypothesis it follows that there are assertions such that

$$\begin{aligned} G_1 &= \wp(\text{gu}_\pi[z''_2, \vartheta'', \text{Pre}[\text{true}, \vartheta, U], \text{Rel}_1[\text{true}, \vartheta, U]]), \\ G_2 &= \wp(\text{gu}_\pi[z''_3, \vartheta'', \text{Pre}[\text{true}, \vartheta, U], \text{Rel}_2[\text{true}, \vartheta, U]]), \\ R_1 &= (G_2 \vee R'')^\dagger, \\ R_2 &= (G_1 \vee R'')^\dagger, \end{aligned}$$

Hence, the induction hypothesis, proposition 10 on page 193 and proposition 11 on page 195 imply that $\text{ef}_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\begin{aligned} &\overleftarrow{\exists h[\vartheta]}, h[\vartheta] : \\ &\wp(\text{ef}_\pi[z''_2, \vartheta'', P'', R_1]) \wedge \wp(\text{ef}_\pi[z''_3, \vartheta'', P'', R_2]), \end{aligned}$$

$\text{wa}_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\begin{aligned} &\exists h[\vartheta] : \\ &(\wp(\text{wa}_\pi[z''_2, \vartheta'', P'', R_1]) \wedge \wp(\text{wa}_\pi[z''_3, \vartheta'', P'', R_2])) \vee \\ &(\wp(\text{wa}_\pi[z''_2, \vartheta'', P'', R_1]) \wedge \overleftarrow{\exists \vartheta''} : \wp(\text{ef}_\pi[z''_3, \vartheta'', P'', R_2])) \vee \\ &(\wp(\text{wa}_\pi[z''_3, \vartheta'', P'', R_2]) \wedge \overleftarrow{\exists \vartheta''} : \wp(\text{ef}_\pi[z''_2, \vartheta'', P'', R_1])) \end{aligned}$$

while $\text{gu}_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\begin{aligned} &\overleftarrow{\exists h[\vartheta]}, h[\vartheta] : \\ &\wp(\text{gu}_\pi[z''_2, \vartheta'', P'', R_1]) \vee \wp(\text{gu}_\pi[z''_3, \vartheta'', P'', R_2]). \end{aligned}$$

- Await: Given a specified program

$$z_1 \text{ sat } (\vartheta, \{\}) :: (P, R, W, G, E)$$

of auxiliary form, a program z'_1 of the form

$$\text{await } b \text{ do } z'_2 \text{ od,}$$

such that $\models z_1 \hookrightarrow z'_1$, and assume that α is the set of auxiliary variables in z_1 . This means that z_1 is of the form

$$\begin{aligned} &\text{await } b \text{ true} \\ &\quad z_2; \\ &\quad a_1 := u_1; \\ &\quad \vdots \\ &\quad a_m := u_m \\ &\text{od,} \end{aligned}$$

where for all $1 \leq j, k \leq m$, $\text{var}[u_j] \subseteq (\vartheta \setminus \alpha) \cup \{a_j\}$, $a_j \in \alpha$, $j \neq k \Rightarrow a_j \neq a_k$ and $\models z_2 \hookrightarrow z'_2$.

By the induction hypothesis it follows that $ef_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\overleftarrow{P} \wedge R | \wp(ef_\pi[z_2, \vartheta, P^R \wedge b, I]) | (I_\alpha \wedge \bigwedge_{j=1}^m a_j = \overleftarrow{u_j}) | R,$$

$wa_\pi[z_1, \vartheta, P, R]$ is characterised by

$$P^R \wedge \neg b,$$

while $gu_\pi[z_1, \vartheta, P, R]$ is characterised by

$$\wp(ef_\pi[z_2, \vartheta, P^R \wedge b, I]) | (I_\alpha \wedge \bigwedge_{j=1}^m a_j = \overleftarrow{u_j}).$$

end of proof

18.5 Relative-Completeness Propositions

Proposition 14 *Given that*

- $\models_{\pi} z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E),$
- $R \text{ respects } \vartheta \cap \text{hid}[z],$

then

- $Tr_{\pi} \vdash_B z \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E).$

Proof: We will prove the proposition by structural induction on z . The base-cases are the skip- and assignment-statements.

- Skip: Assume that

$$\models_{\pi} \text{skip sat } (\vartheta, \alpha) :: (P, R, W, G, E).$$

The skip- and pre-rules imply that

$$Tr_{\pi} \vdash_B \text{skip sat } (\vartheta, \alpha) :: (P, R, W, G, \overleftarrow{P} \wedge R). \quad (18.1)$$

Moreover, since

$$ef_{\pi}[\text{skip}, \vartheta \cup \alpha, P, R] = \overleftarrow{P} \wedge R,$$

it follows that

$$\models_{\pi} \overleftarrow{P} \wedge R \Rightarrow E,$$

in which case, 18.1 and the consequence-rule give

$$Tr_{\pi} \vdash_B \text{skip sat } (\vartheta, \alpha) :: (P, R, W, G, E).$$

- Assignment: Assume that z_1 is of the form

$$v := r,$$

and that

$$\models_{\pi} z_1 \underline{\text{sat}} (\vartheta, \bigcup_{j=1}^m \{a_j\}) :: (P, R, W, G, E).$$

This means that there is a program z'_1 of the form

await true do

$$a_1 := u_1;$$

$\vdots$

$$a_m := u_m;$$

$$v := r$$

od

such that

$$\begin{aligned} & \models z'_1 \xrightarrow{(\vartheta, \bigcup_{j=1}^m \{a_j\})} z_1, \\ & \models_{\pi} z'_1 \underline{\text{sat}} (\vartheta \cup \bigcup_{j=1}^m \{a_j\}, \{\}) :: (P, R, W, G, E), \end{aligned} \quad (18.2)$$

where for all $1 \leq j, k \leq m$, $\text{var}[u_j] \subseteq \vartheta \cup \{a_j\}$ and $j \neq k \Rightarrow a_j \neq a_k$.

Moreover, since

$$\begin{aligned} ef_{\pi}[z'_1, \vartheta \cup \bigcup_{j=1}^m \{a_j\}, P, R] &= \\ & \overleftarrow{P} \wedge R | (v = \overleftarrow{r} \wedge I_{\{v\} \cup \bigcup_{j=1}^m \{a_j\}} \wedge \bigwedge_{j=1}^m a_j = \overleftarrow{u_j}) | R, \\ gu_{\pi}[z'_1, \vartheta \cup \bigcup_{j=1}^m \{a_j\}, P, R] &= \\ & \overleftarrow{P}^R \wedge v = \overleftarrow{r} \wedge I_{\{v\} \cup \bigcup_{j=1}^m \{a_j\}} \wedge \bigwedge_{j=1}^m a_j = \overleftarrow{u_j}, \end{aligned}$$

it follows from 18.2 that there is an assertion E_1 such that

$$\begin{aligned} \models_{\pi} \overleftarrow{P^R} \wedge v &= \overleftarrow{r} \wedge I_{\{v\} \cup \bigcup_{j=1}^m \{a_j\}} \wedge \bigwedge_{j=1}^m a_j = \overleftarrow{u_j} \Rightarrow G \wedge E_1, \\ \models_{\pi} \overleftarrow{P} \wedge R|E_1|R &\Rightarrow E, \end{aligned}$$

in which case

$$Tr_{\pi} \vdash_B z_1 \underline{\text{sat}} (\vartheta, \bigcup_{j=1}^m \{a_j\}) :: (P, R, W, G, E).$$

follows by the assignment-, consequence- and pre-rules.

- Block: Assume that z_1 is of the form

begin loc $v_1, \dots, v_m; z_2$ end,

and that

$$\models_{\pi} z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E).$$

Since $\bigcup_{j=1}^m \{v_j\} \subseteq \text{hid}[z_1]$, it is clear that

$$\models_{\pi} z_2 \underline{\text{sat}} (\vartheta \cup \bigcup_{j=1}^m \{v_j\}, \alpha) :: (P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}, W, G, E).$$

Moreover, since $R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}$ respects $\vartheta \cap \text{hid}[z_2]$, the induction hypothesis implies that

$$Tr_{\pi} \vdash_B z_2 \underline{\text{sat}} (\vartheta \cup \bigcup_{j=1}^m \{v_j\}, \alpha) :: (P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}, W, G, E),$$

in which case

$$Tr_{\pi} \vdash_B z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E),$$

follows by the block-rule.

- Sequential: Assume that z_1 is of the form

$$z_2; z_3,$$

and that

$$\models_{\pi} z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E).$$

This means that there is a program z'_1 of the form

$$z'_2; z'_3,$$

such that

$$\begin{aligned} & \models z'_1 \xrightarrow{(\vartheta, \alpha)} z_1, \\ & \models_{\pi} z'_1 \underline{\text{sat}} (\vartheta \cup \alpha, \{\}) :: (P, R, W, G, E). \end{aligned} \tag{18.3}$$

It follows from proposition 13 on page 198 that there are assertions such that

$$\begin{aligned} E_1 &= \wp(\text{ef}_{\pi}[z'_2, \vartheta \cup \alpha, P, R]), \\ P_2 &= \exists \overline{\vartheta \cup \alpha} : \overline{P} \wedge E_1, \\ E_2 &= \wp(\text{ef}_{\pi}[z'_3, \vartheta \cup \alpha, P_2, R]), \end{aligned}$$

in which case 18.3 and the induction hypothesis imply

$$\begin{aligned} Tr_{\pi} \vdash_B z_2 \underline{\text{sat}} (\vartheta, \alpha) &:: (P, R, W, G, E_1 \wedge P_2), \\ Tr_{\pi} \vdash_B z_3 \underline{\text{sat}} (\vartheta, \alpha) &:: (P_2, R, W, G, E_2). \end{aligned}$$

It is also clear that

$$\models_{\pi} E_1 | E_2 \Rightarrow E,$$

in which case

$$Tr_{\pi} \vdash_B z_1 \underline{\text{sat}} (\vartheta, \{\}) :: (P, R, W, G, E),$$

follows by the sequential- and consequence-rules.

- If: Assume that z_1 is of the form

if b then z_2 else z_3 fi,

and that

$$\models_{\pi} z_1 \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E).$$

Then it is clear that

$$\begin{aligned} \models_{\pi} z_2 \text{ sat } (\vartheta, \alpha) &:: (P \wedge b, R, W, G, E), \\ \models_{\pi} z_3 \text{ sat } (\vartheta, \alpha) &:: (P \wedge \neg b, R, W, G, E), \end{aligned}$$

and it follows from the induction hypothesis that

$$\begin{aligned} Tr_{\pi} \vdash_B z_2 \text{ sat } (\vartheta, \alpha) &:: (P \wedge b, R, W, G, E), \\ Tr_{\pi} \vdash_B z_3 \text{ sat } (\vartheta, \alpha) &:: (P \wedge \neg b, R, W, G, E), \end{aligned}$$

which means that

$$Tr_{\pi} \vdash_B z_1 \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E)$$

can be deduced by the if-rule.

- While: Assume that z_1 is of the form

while b do z_2 od,

and that

$$\models_{\pi} z_1 \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, E).$$

This means that there is a program z'_1 of the form

while b do z'_2 od,

such that

$$\begin{aligned} & \models z'_1 \xrightarrow{(\vartheta, \alpha)} z_1, \\ & \models_{\pi} z'_1 \text{ \underline{sat} } (\vartheta \cup \alpha, \{\}) :: (P, R, W, G, E). \end{aligned} \quad (18.4)$$

It follows from proposition 13 on page 198 that there are assertions such that

$$\begin{aligned} P' &= \overleftarrow{\exists \vartheta \cup \alpha} : \overleftarrow{P} \wedge \wp(\text{ef}_{\pi}[z'_2, \vartheta \cup \alpha, b, R])^*, \\ Z &= \wp(\text{ef}_{\pi}[z'_2, \vartheta \cup \alpha, P' \wedge b, R]). \end{aligned}$$

Assume Z is not well-founded. But then, there is a diverging computation

$$\sigma \in \text{ext}_{\pi}[P, R] \cap \text{cp}_{\pi}[z'_1].$$

This contradicts 18.4. Thus,

$$\models_{\pi} \text{ \underline{wf} } Z. \quad (18.5)$$

Moreover, 18.4 and the induction hypothesis give

$$\text{Tr}_{\pi} \vdash_B z_2 \text{ \underline{sat} } (\vartheta, \alpha) :: (P', R, W, G, P' \wedge Z).$$

But then,

$$\text{Tr}_{\pi} \vdash_B z_1 \text{ \underline{sat} } (\vartheta, \alpha) :: (P', R, W, G, (Z^{\dagger} \vee R) \wedge \neg b)$$

follows from 18.5 by the while-rule. The consequence- and pre-rules give

$$\text{Tr}_{\pi} \vdash_B z_1 \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, \overleftarrow{P} \wedge (Z^{\dagger} \vee R) \wedge \neg b).$$

Furthermore, it follows easily that

$$\models_{\pi} \overleftarrow{P} \wedge (Z^{\dagger} \vee R) \wedge \neg b \Rightarrow E,$$

which by the consequence-rule gives

$$\text{Tr}_{\pi} \vdash_B z_1 \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E).$$

- Parallel: Assume that z_1 is of the form

$$\{z_2 \parallel z_3\}$$

and that

$$\models_{\pi} z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E).$$

This means that there is a program z'_1 of the form

$$\{z'_2 \parallel z'_3\}$$

such that

$$\begin{aligned} \models z'_1 &\xrightarrow{(\vartheta, \alpha)} z_1, \\ \models_{\pi} z'_1 \underline{\text{sat}} (\vartheta \cup \alpha, \{\}) &:: (P, R, W, G, E). \end{aligned}$$

This specified program is of auxiliary form, which means that it has a historic form

$$\models_{\pi} z''_1 \underline{\text{sat}} (\vartheta \cup \alpha'', \{\}) :: (P'', R'', W, G, E) \quad (18.6)$$

where z''_1 is of the form

$$\{z''_2 \parallel z''_3\}$$

and

$$\models z''_1 \xrightarrow{(\vartheta, \alpha'')} z_1.$$

Let $U = \text{var}[z_1] \setminus \vartheta$, it follows from proposition 13 on page 198 that there are assertions such that

$$\begin{aligned} G_1 &= \wp(gu_{\pi}[z''_2, \vartheta \cup \alpha'', \text{Pre}[\mathbf{true}, \vartheta \cup \alpha, U], \text{Rel}_1[\mathbf{true}, \vartheta \cup \alpha, U]]), \\ G_2 &= \wp(gu_{\pi}[z''_3, \vartheta \cup \alpha'', \text{Pre}[\mathbf{true}, \vartheta \cup \alpha, U], \text{Rel}_2[\mathbf{true}, \vartheta \cup \alpha, U]]), \\ R_1 &= (G_2 \vee R'')^{\dagger}, \\ R_2 &= (G_1 \vee R'')^{\dagger}, \\ E_1 &= \wp(ef_{\pi}[z''_2, \vartheta \cup \alpha'', P'', R_1]), \\ E_2 &= \wp(ef_{\pi}[z''_3, \vartheta \cup \alpha'', P'', R_2]), \\ W_1 &= \wp(wa_{\pi}[z''_2, \vartheta \cup \alpha'', P'', R_1]) \wedge \neg W, \\ W_2 &= \wp(wa_{\pi}[z''_3, \vartheta \cup \alpha'', P'', R_2]) \wedge \neg W. \end{aligned}$$

It follows easily from 18.6 and proposition 10 on page 193 that

$$\models_{\pi} \neg(W_1 \wedge W_2) \wedge \neg(W_1 \wedge E_2) \wedge \neg(W_2 \wedge E_1). \quad (18.7)$$

Moreover, 18.6 and proposition 12 on page 196 imply that any computation

$$\sigma \in \text{ext}_{\pi}[P'', R_1] \cap \text{cp}_{\pi}[z_2'']$$

either deadlocks or terminates. The same is of course true for any computation

$$\sigma' \in \text{ext}_{\pi}[P'', R_2] \cap \text{cp}_{\pi}[z_3''].$$

By proposition 11 on page 195 and the induction hypothesis we get

$$\begin{aligned} Tr_{\pi} \vdash_B z_2 \underline{\text{sat}} (\vartheta, \alpha'') &:: (P'', R_1, W \vee W_1, G \wedge R_2, E_1), \\ Tr_{\pi} \vdash_B z_3 \underline{\text{sat}} (\vartheta, \alpha'') &:: (P'', R_2, W \vee W_2, G \wedge R_1, E_2), \end{aligned}$$

which together with 18.7 and the parallel-rule give

$$Tr_{\pi} \vdash_B z_1 \underline{\text{sat}} (\vartheta, \alpha'') :: (P'', R_1 \wedge R_2, W, G, E_1 \wedge E_2).$$

Moreover, since

$$\begin{aligned} \models_{\pi} R'' &\Rightarrow R_1 \wedge R_2, \\ \models_{\pi} E_1 \wedge E_2 &\Rightarrow E, \end{aligned}$$

it follows by the consequence-rule that

$$Tr_{\pi} \vdash_B z_1 \underline{\text{sat}} (\vartheta, \alpha'') :: (P'', R'', W, G, E),$$

in which case

$$Tr_{\pi} \vdash_B z_1 \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E),$$

can be deduced by the elimination- and consequence-rules.

- Await:

Assume that z_1 is of the form

await b do z_2 od,

and that

$$\models_{\pi} z_1 \underline{\text{sat}} (\vartheta, \bigcup_{j=1}^m \{a_j\}) :: (P, R, W, G, E).$$

This means that there is a program z'_1 of the form

await b do
 z'_2 ;
 $a_1 := u_1$;
 $\vdots$
 $a_m := u_m$
od

such that

$$\begin{aligned} \models z'_1 & \xrightarrow{(\vartheta, \bigcup_{j=1}^m \{a_j\})} z_1, \\ \models_{\pi} z'_1 \underline{\text{sat}} (\vartheta \cup \bigcup_{j=1}^m \{a_j\}, \{\}) & :: (P, R, W, G, E), \end{aligned} \quad (18.8)$$

where for all $1 \leq j, k \leq m$, $\text{var}[u_j] \subseteq \vartheta \cup \{a_j\}$, $j \neq k \Rightarrow a_j \neq a_k$ and $\models z'_2 \hookrightarrow z_2$.

It follows from proposition 13 on page 198 that there is an assertion such that

$$E_1 = \wp(\text{ef}_{\pi}[z'_2, \vartheta \cup \bigcup_{j=1}^m \{a_j\}, P^R \wedge b, I]),$$

Moreover, 18.8 implies that

$$\models_{\pi} P^R \wedge \neg b \Rightarrow W, \quad (18.9)$$

and together with the induction hypothesis also that

$$Tr_\pi \vdash_B z_2 \underline{\text{sat}} (\vartheta, \bigcup_{j=1}^m \{a_j\}) :: (P^R \wedge b, I, \text{false}, \text{true}, E_1). \quad (18.10)$$

Moreover, it is clear from 18.8 that there is an assertion E_2 such that

$$\begin{aligned} \models_\pi E_1 | (\bigwedge_{j=1}^m a_j = \overleftarrow{u_j} \wedge I_{\bigcup_{j=1}^m \{a_j\}}) &\Rightarrow G \wedge E_2, \\ \models_\pi \overleftarrow{P} \wedge R | E_2 | R &\Rightarrow E, \end{aligned} \quad (18.11)$$

in which case

$$Tr_\pi \vdash_B z_1 \underline{\text{sat}} (\vartheta, \bigcup_{j=1}^m \{a_j\}) :: (P, R, W, G, E)$$

follows from 18.9, 18.10, 18.11 by the await-, pre- and consequence-rules.

end of proof

Proposition 15 *Given that*

$$\models_\pi z \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E)$$

then

$$Tr_\pi \vdash_B z \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E).$$

Proof: Assume that

$$\models_\pi z \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E).$$

It follows easily that

$$\models_\pi z \underline{\text{sat}} (\vartheta, \alpha) :: (P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}, W, G, E),$$

where $\vartheta \cap \text{hid}[z] = \bigcup_{j=1}^m \{v_j\}$. But then, it is clear from proposition 14 on page 206 that

$$Tr_\pi \vdash_B z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R \wedge \bigwedge_{j=1}^m v_j = \overleftarrow{v_j}, W, G, E),$$

in which case it follows by the access-rule that

$$Tr_\pi \vdash_B z \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R, W, G, E).$$

end of proof

Chapter 19

Discussion

19.1 Motivation

The object of this chapter is first of all to motivate some of the design decisions taken at different points above. It will be explained how things could have been done differently, and alternative approaches will be discussed and compared with the one chosen.

The reason why this has been postponed until now is that it is first at this stage the reader has an overview of the whole system. Thus it is only now that the reader can fully evaluate and understand the consequences of the alternative approaches.

It is also the object of this chapter to indicate weaknesses and areas for further research, and furthermore to compare LSP with some of the most closely related methods known from the literature.

19.2 Without Hooked Variables

19.2.1 Two Representations

In LSP the *rely*-, *guar*- and *eff*-conditions are binary assertions. This means that they may have free occurrences of hooked variables.

However, our approach does not depend upon the use of binary assertions. LSP can be transformed into a logic where all assertions are unary. The choice between the two representations can be seen as a matter of taste, although we believe that the binary assertions in many cases result in shorter and more

readable specifications.

19.2.2 Unary Eff-Conditions

Since the pre-condition characterises the initial state, and since the use of auxiliary variables is allowed, the eff-condition could just as well be unary. The post-condition in Hoare-logic is for example unary.

One advantage of such a modification is a simpler sequential-rule. Unfortunately, this change has the opposite effect on some of the other rules; for example the while-rule. A unary eff-condition would also have a complicating effect on certain specifications. Consider for example

$$(\{S\}, \{\}) :: (\text{true}, I, \text{false}, \text{true}, \#S = \overleftarrow{\#S} + 1),$$

which restricts an implementation to increase the number of elements in the set S by one. If the eff-condition is constrained to be unary, it is necessary to introduce an auxiliary variable to write an equivalent specification:

$$(\{S\}, \{m\}) :: (\#S = m, I, \text{false}, \text{true}, \#S = m + 1).$$

19.2.3 Rely and Guar as Assertion Sets

In [Sti88] rely- and guar-conditions are represented as sets of assertions. Given a structure π , the basic idea is that for any set of assertions Δ , there is a set of state changes that are invariant with respect to it, namely the set of all pairs of states (s_1, s_2) , such that for all assertions $A \in \Delta$:

- $s_1 \models_{\pi} A$ implies $s_2 \models_{\pi} A$.

To compare this approach with ours, consider

$$x = \overleftarrow{x} \wedge y \geq \overleftarrow{y},$$

and assume the variables x and y are respectively of the sorts $\mathbb{N}$ and $\mathbb{R}$. To express the first conjunct in the notation of [Sti88], it is enough to constrain the environment to maintain any formula of the form $x = n$, while the set of all assertions of the form $y \geq r$ characterises the second. Thus, the binary assertion above can be transformed into the following set of unary assertions:

$$\{x = n | n \in \mathbb{N}\} \cup \{y \geq r | r \in \mathbb{R}\}.$$

19.3 Separate Pre-Conditions

It may be argued that because the eff-condition is binary, a separate pre-condition is not really needed (see [HM87]). The reason is of course that the eff-condition can be redefined to constrain both the set of initial states and the overall effect. Given such a formalism we may write

$$(\vartheta, \alpha) :: (R, W, G, \overleftarrow{P} \wedge E)$$

instead of

$$(\vartheta, \alpha) :: (P, R, W, G, E).$$

Nevertheless, we have decided to use five assertions. There are two reasons for that:

- A specification with a separate pre-condition is more readable.
- The pre-condition is a useful tool when formulating the program decomposition rules. Alternatively, we could have introduced a new syntactic operator on assertions to characterise the domain of the eff-condition. This is for example the view taken in Z [Spi88].

19.4 Scope of Eff-Conditions

19.4.1 Three Alternatives

In LSP the eff-condition not only characterises the overall effect of the specified program, but also any change due to the environment both before the first internal transition and after the last. This means that interference both before the implementation starts up and after it has terminated is included in the eff-condition. A similar view was taken in [Jon81]. However, there are some obvious alternatives to this:

- The eff-condition characterises the overall state transition from the initial state to immediately after the last internal transition.
- The eff-condition characterises the overall state transition from immediately before the first internal transition to the final state. (Remember that a terminating program has only finite computations with respect to a convergent environment.)
- The eff-condition characterises the overall state transition from immediately before the first internal transition to immediately after the last internal transition.

In the first case the eff-condition is only required to be true immediately after the implementation terminates and does not have to be preserved by the environment.

In the second case no interference before the first internal transition is included in the eff-condition, while interference after the last is.

In the third only interference which occurs after the first internal transition and before the last influences the eff-condition.

19.4.2 Stirling

The first position is taken by Stirling in [Sti88]. However, his approach is not completely consistent, because the rely-condition of a program whose main construct is the parallel-statement is required to preserve the post-condition. For example, it is possible to prove that the program $b := \text{false}$ satisfies:

```

pre    $a \wedge b$ 
rely    $\{b, \neg b, b \vee c \Rightarrow a\}$ 
guar    $\{c, \neg c, b \vee c \Rightarrow a\}$ 
post    $a \wedge \neg b,$ 

```

and that the program $c := \text{false}$ satisfies:

```

pre    $a \wedge c$ 
rely    $\{c, \neg c, b \vee c \Rightarrow a\}$ 
guar    $\{b, \neg b, b \vee c \Rightarrow a\}$ 
post    $a \wedge \neg c.$ 

```

However, because Stirling's parallel-rule insists that the two post-conditions are preserved by their respective rely-conditions, it is as far as we can see only possible to prove that $\{b := \text{false} \parallel c := \text{false}\}$ satisfies

```

pre    $a \wedge b \wedge c$ 
rely   $\{c, \neg c, b, \neg b, b \vee c \Rightarrow a\}$ 
guar   $\{\}$ 
post   $\neg b \wedge \neg c,$ 

```

and not that the same program satisfies

```

pre    $a \wedge b \wedge c$ 
rely   $\{c, \neg c, b, \neg b, b \vee c \Rightarrow a\}$ 
guar   $\{\}$ 
post   $\neg b \wedge \neg c \wedge a,$ 

```

which is what one would have expected. To achieve this a more sophisticated parallel-rule is needed.

19.4.3 First Alternative Approach

If we change the semantics of LSP in such a way that interference after the last internal transition is not included in the eff-condition, the assignment-rule can be slightly simplified. The reason is that the second occurrence of R in the conclusions eff-condition can be removed. Moreover, the await-rule can be changed in a similar way. Unfortunately, the parallel-rule becomes more complicated:

$$\frac{
\begin{array}{l}
\neg(W_1 \wedge E_2 | R_2) \wedge \neg(W_2 \wedge E_1 | R_1) \wedge \neg(W_1 \wedge W_2) \\
z_1 \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R_1, W \vee W_1, G \wedge R_2, E_1) \\
z_2 \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R_2, W \vee W_2, G \wedge R_1, E_2)
\end{array}
}{
\{z_1 \parallel z_2\} \text{ \underline{sat} } (\vartheta, \alpha) :: (P, R_1 \wedge R_2, W, G, (E_1 | R_1 \wedge E_2) \vee (E_1 \wedge E_2 | R_2))
}$$

Redefining the eff-condition in this way can therefore not be said to result in any overall simplification of the decomposition rules.

It may be argued that our approach, where the eff-condition covers interference both before the first and after the last internal transition, may lead to extra proof-work when two statements are composed in sequence. The

reason is that the interference after the last internal transition of the first statement and before the first internal transition of the second statement must be included in the eff-conditions of both statements.

But then, what about pre-condition preservation? If the eff-condition of the first statement has been determined, and it is known that this eff-condition is preserved by the rely-condition, it is enough to prove that the first statement's eff-condition implies the second statement's pre-condition to make sure that even the latter assertion is preserved by the rely-condition.

Otherwise it would have been necessary firstly to prove that the first statement's eff-condition implies the second statement's pre-condition, and then show that this pre-condition is preserved by the rely-condition.

Moreover, the first alternative approach certainly results in more proof-work when it comes to parallel composition.

19.4.4 Second Alternative Approach

The arguments for and against this approach are similar to the previous case.

19.4.5 Third Alternative Approach

In this case, the sequential-rule must be changed to

$$\frac{\begin{array}{l} \overline{P_2} \wedge R \Rightarrow P_2 \\ z_1 \text{ sat } (\vartheta, \alpha) :: (P_1, R, W, G, P_2 \wedge E_1) \\ z_2 \text{ sat } (\vartheta, \alpha) :: (P_2^R, R, W, G, E_2) \end{array}}{z_1; z_2 \text{ sat } (\vartheta, \alpha) :: (P_1, R, W, G, E_1 | R | E_2)}.$$

The occurrence of R in the conclusion's eff-condition is needed to cover interference between the last internal transition of the first statement and the first internal transition of the second.

Unfortunately, some of the other rules like the while-rule and the parallel-rule become much more complicated, and the approach has been rejected for that reason.

19.5 Multi-State Assertions

19.5.1 Accessing the Initial State

As pointed out in [GR89], in some cases it is useful to refer to the initial state in the rely- or guar-conditions in the same way that we can use hooked variables to refer to the initial state in the eff-condition. The same is of course true for the wait-condition.

For example, if primed¹ variables in the wait-, rely- and guar-conditions refer to the initial state, then the specification

glo	$\{S\}$
aux	$\{\}$
pre	true
rely	$S = \overleftarrow{S} \vee (\# \overleftarrow{S} = \#S' - 1 \wedge \overleftarrow{S} \subset S \wedge \#S = \#S')$
wait	$\#S = \#S' - 1$
guar	$S = \overleftarrow{S} \vee (\# \overleftarrow{S} = \#S' \wedge S \subset \overleftarrow{S} \wedge \#S = \#S' - 1)$
eff	$\#S = \# \overleftarrow{S},$

where $\#S$ denotes the number of elements in the set S , describes a program which either deadlocks in a state such that the size of S has been reduced with one, or repeats the following procedure a finite number of times:

- Reduce the size of S by one and wait until an element has been added.

19.5.2 Accessing the Final State

If primed variables are employed to let the rely-, guar- and wait-conditions refer to a computation's initial state, why not introduce a similar convention to permit the same assertions to refer to the final state too? (Remember that a terminating program has only finite computations with respect to a convergent environment.) If for example double-primed variables refer to the final state, then an implementation of

¹The combination of primes and hooks is not good. The introduction of three- or four-state assertions would require a better naming convention.

glo	$\{S\}$
aux	$\{\}$
pre	true
rely	$S = \overleftarrow{S}$
wait	false
guar	$\#S \geq \#S''$
eff	true

can only terminate in a state in which the number of elements in S is less than or equal to the sets size at any previous state.

19.5.3 Where to Stop?

It does not have to stop with assertions accessing four states. On some occasions it could be useful to refer to any state between the current state and the initial state. One might then for example express that a specific action can only take place if a particular flag has not been switched off between the initial state and the current state.

So the question is not only, should we allow the different assertions to refer to the initial state, the final state, etc. , but also — where shall we say ‘stop’?

19.5.4 Conclusion

It may be true that the introduction of three- and four-state assertions has a simplifying effect on some specifications. The examples above seem to confirm this.

Such a restricted use of multi-state assertions will not give us the necessary expressive power to get rid of auxiliary variables altogether, but this is of course no argument against employing them to simplify specifications, since binary assertions are used for the same purpose.

The reason, why we have decided not to introduce assertions of more than two states, is the complicating effect they have on the decomposition-rules. Try for example to formulate the sequential-rule when the rely-, wait- and guar-conditions can refer to both the initial and the final states.

19.6 Atomicity

Both expressions and assignments have so far been required to be atomic. As explained above (see page 11.3.4) this does not mean that the execution of two expressions cannot overlap in time, but only that they behave as if they were executed in sequential order.

Some may argue that since LSP constrains the environment with a rely-condition, we can do better than that; more precisely it should be possible to reason on the level of memory reference without any (other) atomicity constraint.

Unfortunately, this is more difficult than it sounds. First of all, it would lead to a more complicated semantics, and the change would certainly have a similar effect on some of the decomposition-rules. But this is of course only what one could expect.

However, the main argument against such a modification is that it would necessitate a fundamental change in the way we are dealing with auxiliary variables, because it would no longer be possible to place an assignment-statement inside the body of an await-statement without changing the behaviour of the algorithm.

19.7 Boolean Tests of If and While

In the definition of our programming language the Boolean tests of if- and while-statements are prohibited from accessing global variables. As explained earlier (see page 37) this constraint does not reduce the set of possible algorithms, but has the obvious disadvantage of increasing the length of programs. However, this is in our opinion outweighed by simpler decomposition-rules, and more importantly, that it is easier to reason with auxiliary variables.

When reasoning with auxiliary variables it is often of great importance that the auxiliary structure is updated in the same internal transition as the global structure is updated or read.

This is achieved in the case of the assignment-statement, because due to the assignment-rule, the execution of an assignment-statement of the form

$$v := r$$

actually corresponds to the execution of a statement of the form

```

await  $b$  do
   $a_1 := u_1;$ 
   $\vdots$ 
   $a_m := u_m;$ 
   $v := r$ 
od.

```

Similarly, due to the await-rule the execution of an await-statement of the form

```

await  $b$  do  $z$  od

```

corresponds to the execution of a statement of the form

```

await  $b$  do
   $z;$ 
   $a_1 := u_1;$ 
   $\vdots$ 
   $a_m := u_m$ 
od.

```

Moreover, due to the constraint on the Boolean tests of if- and while-statements, this is also achieved in their cases, since the execution of a statement of the form

```

begin
  loc  $v;$ 
   $v := b;$ 
  if  $v$  then  $z_1$  else  $z_2$  fi
end

```

corresponds to the execution of a statement of the form


```

begin
  loc  $v$ ;
  await true do
     $a_1 := v_1$ ;
     $\vdots$ 
     $a_m := v_m$ ;
     $v := b$ 
  od;
  if  $v$  then  $z'_1$  else  $z'_2$  fi
end,

```

where z'_1 and z'_2 are the results of adding auxiliary structure to respectively z_1 and z_2 , while the execution of a statement of the form

```

begin
  loc  $v$ ;
   $v := b$ ;
  while  $v$  do  $z; v := b$  od
end

```

corresponds to the execution of a statement of the form

```

begin
  loc v;
  await true do
    a1 := v1;
    ⋮
    am := vm;
    v := b
  od;
  while v do
    z';
    await true do
      a1 := v1;
      ⋮
      am := vm;
      v := b
    od
  od
end,

```

where z' is the result of adding auxiliary structure to z . The same is of course not true for if- and while-statements if the truth values of their Boolean tests are not maintained by the environment. In that case, one possibility is to allow auxiliary structure of the form

$$\begin{aligned}
\langle ite \rangle &::= \text{if } \langle tst \rangle : \langle prg \rangle \text{ then } \langle prg \rangle \text{ else } \langle prg \rangle \text{ fi,} \\
\langle wd \rangle &::= \text{while } \langle tst \rangle : \langle prg \rangle \text{ do } \langle prg \rangle \text{ od,}
\end{aligned}$$

where $\langle tst \rangle : \langle prg \rangle$ is a pair of a Boolean test and a program which is assumed to be executed in isolation; in other words, as one internal transition. However, this would lead to a more complicated and less intuitive semantics.

19.8 Fairness

19.8.1 Introduction

When discussing fairness, it is usual to distinguish between weak fairness; that an event will not be infinitely postponed provided that it remains con-

tinuously enabled, and strong fairness; that an event will not be infinitely postponed provided that it is enabled infinitely often. See [Fra86] for a more detailed discussion.

The programming language discussed in this thesis is unfair. This means that LSP is not complete with respect to a weakly-fair programming language, and even less so if the language is strongly fair.

For example, as explained earlier, if the language is unfair the parallel composition of the two programs

$b := \text{false},$

and

```
begin
  loc v;
  v := b;
  while v do
    x := x + 1;
    v := b
  od
end
```

is not guaranteed to terminate because the first may be infinitely overtaken by the second. In other words, LSP cannot be used to prove properties of programs whose algorithms depend upon busy waiting.

19.8.2 Two Alternative Systems

Nevertheless, we believe that it is possible to transform the present system into two new systems; one which deals with weak fairness, and another which is specially designed to handle strong fairness. The decomposition-rules for the two systems are closely related to the rules given in LSP.

The basic idea is that although a busy-waiting process never becomes blocked, it is actually waiting for the environment to release it. Thus instead of only using the wait-condition to characterise the set of states in which the implementation may become blocked, we will also use the wait-condition to describe the set of states in which a busy process needs help from its environment to be released.

19.9 Data Reification

19.9.1 Introduction

Program development is often divided into two rather separate subfields:

- program decomposition,
- data reification (also called data refinement).

The first topic has been discussed in detail above, but so far nothing has been said about data reification.

The earliest suggestions for a reification method were given in [Mil71], [Hoa72b] and [Jon72]. These ideas have since then been further developed and discussed by many authors, see for example [Jon80], [HHS86], [Nip86], [Sta88] and [AL88].

19.9.2 Jones' System

Jones did not give any reification-rule in [Jon81], but stated informally that in the same way as the VDM-rule allows the pre-condition to be weakened and the post-condition to be strengthened, it is sound to weaken the pre- and rely-conditions and strengthen the guar- and eff-conditions.

This rule is not complete. The incompleteness has been discussed by several authors (see [GR89], [WD88] and [XH90]), but so far no complete rule has been proposed. A reification-rule for rely-guarantee specifications is suggested in [GQNL90], but as pointed out in that paper; the method is incomplete and deals only with safety properties.

19.9.3 LSP

The reification-rule indicated by Jones can be generalised to LSP in an obvious way; by insisting that the wait-condition is also strengthened. How to formulate a complete rule is an open question.

19.10 Nondeterminacy

19.10.1 Introduction

So far only deterministic program constructs have been considered (although the parallel-statement gives rise to nondeterministic behaviour). By introducing two new statements, selection ($\langle sl \rangle$) and repetition ($\langle rp \rangle$):

$$\begin{aligned} \langle sl \rangle &::= \{ \langle gl \rangle \}, \\ \langle rp \rangle &::= * \{ \langle gl \rangle \}, \\ \langle gl \rangle &::= \langle tst \rangle \rightarrow \langle stm \rangle \mid \langle tst \rangle \rightarrow \langle stm \rangle \sqcap \langle gl \rangle, \end{aligned}$$

we will show how LSP can be generalised to deal with nondeterminism in the style of Dijkstra [Dij75].

19.10.2 Selection

We will use

$$\{ \sqcap_{j \in \{1, \dots, m\}} b_j \rightarrow z_j \}$$

as short for

$$\{ b_1 \rightarrow z_1 \sqcap \dots \sqcap b_m \rightarrow z_m \}.$$

For each j , $b_j \rightarrow z_j$ is called a guarded command, while the Boolean test b_j is referred to as a guard. Basically the selection-statement selects a guarded command with a true guard and executes its body. If no guard is true when the selection takes place, the statement aborts.

The problem with modeling this in the operational semantics is that the statement may abort. Fortunately, since we are only concerned with converging programs, this may be dealt with in the same way as nonterminating await-bodies were dealt with above; namely by employing computations with infinitely many internal identity steps:

- $\langle \sqcap_{j \in A} b_j \rightarrow z_j \rangle, s \rangle \xrightarrow{i} \langle z_k, s \rangle$ if $s \models_{\pi} b_k$ and $k \in A$,
- $\langle \sqcap_{j \in A} b_j \rightarrow z_j \rangle, s \rangle \xrightarrow{i} \langle \sqcap_{j \in A} b_j \rightarrow z_j \rangle, s \rangle$ if $s \models_{\pi} \bigwedge_{j \in A} \neg b_j$,

To see that this has the desired effect, let z denote the program

$$\{x = 1 \rightarrow x := 10 \sqcap x = 2 \rightarrow x := 10\}.$$

Clearly if $s(x) = 0$, then there is an infinite computation σ of the form

$$\langle z, s \rangle \xrightarrow{i} \langle z, s \rangle \xrightarrow{i} \dots \xrightarrow{i} \langle z, s \rangle \xrightarrow{i} \dots ,$$

Moreover, σ is an element of $ext_\pi[\mathbf{true}, \mathbf{true}]$, but not an element of

$$int_\pi[\mathbf{false}, \mathbf{true}, \mathbf{true}].$$

Thus

$$z \text{ sat } (\{x\}, \{\}) :: (\mathbf{true}, \mathbf{true}, \mathbf{false}, \mathbf{true}, \mathbf{true})$$

is not valid.

19.10.3 Repetition

In the same way as above

$$*\{\sqcap_{j \in \{1, \dots, m\}} b_j \rightarrow z_j\}$$

is short for

$$*\{b_1 \rightarrow z_1 \sqcap \dots \sqcap b_m \rightarrow z_m\}.$$

Basically, this statement can be viewed as a loop which iteratively executes the selection-statement $\{\sqcap_{j \in A} b_j \rightarrow z_j\}$ until the selection of the guarded command takes place in a state where none of the guards is true, in which case the statement terminates normally. This can be described in the operational semantics as follows:

- $\langle *\{\sqcap_{j \in A} b_j \rightarrow z_j\}, s \rangle \xrightarrow{i} \langle z_k, *\{\sqcap_{j \in A} b_j \rightarrow z_j\}, s \rangle$ if $s \models_\pi b_k$ and $k \in A$,
- $\langle *\{\sqcap_{j \in A} b_j \rightarrow z_j\}, s \rangle \xrightarrow{i} \langle \epsilon, s \rangle$ if $s \models_\pi \bigwedge_{j \in A} \neg b_j$.

19.10.4 Decomposition-Rules

As in the case of the if- and while-statements, the environment is required to preserve the truth-values of the different guards. This means the following two rules are sufficient:

$$\frac{P \Rightarrow \bigvee_{j \in A} b_j \quad z_j \underline{\text{sat}} (\vartheta, \alpha) :: (P \wedge b_j, R, W, G, E)_{j \in A}}{\{\bigwedge_{j \in A} b_j \rightarrow z_j\} \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E)}$$

$$\frac{\underline{\text{wf}} Z \quad z_j \underline{\text{sat}} (\vartheta, \alpha) :: (P \wedge b_j, R, W, G, P \wedge Z)_{j \in A}}{* \{\bigwedge_{j \in A} b_j \rightarrow z_j\} \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, (Z^\dagger \vee R) \wedge \bigwedge_{j \in A} \neg b_j)}$$

The only ‘new’ with respect to the if- and while-rules is that the selection-rule needs an extra premise to ensure that at least one guard is true when the selection takes place.

19.11 Partial Functions

19.11.1 Introduction

So far it has been assumed that all functions are total. This constraint is too strong if we want to apply LSP to ‘real’ problems. Division by 0 is for example one thing which can lead to program abortion. It should therefore not be possible to prove that

$$x := x/x \underline{\text{sat}} (\{x\}, \{\}) :: (\text{true}, x = \frac{_}{x}, \text{false}, \text{true}, \text{true}).$$

The reason is of course that the program aborts if the initial value of x is 0. If we use LSP as it is, and ignore that all functions are required to be total, this specified program is actually ‘provable’. Hence, to handle partial functions some modifications are necessary.

Several logics have been suggested to deal with partial functions. We will employ LPF [BCJ84] here, but this does not mean that we cannot give similar rules with respect to other ‘partial’ logics, like for example the ‘weak

logic' described in [Owe85] or LCF [GMW79]. A discussion of LPF with respect to its usability in program development can be found in [CJ90].

19.11.2 LPF

In LPF the propositional operators are given the strongest monotone extensions of their two-valued interpretations, while quantifiers range only over the 'proper' values of their bound sets. The resulting logic employs only the normal collection of operators and all theorems of LPF hold in classical first-order predicate calculus as well. However, the opposite is not true; because the 'law of the excluded middle' no longer is valid, there are many truths of classical first-order logic that do not hold in LPF.

We will use $\delta(n)$ to mean a Boolean expression which is true if and only if n denotes a proper element of its sort. For example, if b is an expression of sort $\mathbb{B}$, then $\delta(b)$ is equivalent to

$$b \vee \neg b.$$

Observe, that $\delta(b)$ is undefined if b is undefined. At semantic level, we will use the nonmonotonic operator Δ to express if an expression is undefined or not. This means that $\Delta(b)$ is true if b is true or false; otherwise false.

19.11.3 Changes to the Operational Semantics

Because the right-hand expression of the assignment-statement, and the Boolean tests of the if-, while- and await-statements may be undefined, it is necessary to change the operational semantics of these statements. Abortion will be modeled in the same way as in the case of the selection-statement:

- $\langle v := r, s \rangle \xrightarrow{i} \langle \epsilon, s(r^v) \rangle$ if $s \models_{\pi} \Delta(r)$, and $s(r^v)$ denotes the state that is obtained from s , by mapping the variable v to the value of the term r , determined by π and s , and leaving all other maplets unchanged,
- $\langle v := r, s \rangle \xrightarrow{i} \langle v := r, s \rangle$ if $s \models_{\pi} \neg \Delta(r)$,
- $\langle \text{if } b \text{ then } z_1 \text{ else } z_2 \text{ fi}, s \rangle \xrightarrow{i} \langle z_1, s \rangle$ if $s \models_{\pi} b \wedge \Delta(b)$,
- $\langle \text{if } b \text{ then } z_1 \text{ else } z_2 \text{ fi}, s \rangle \xrightarrow{i} \langle z_2, s \rangle$ if $s \models_{\pi} \neg b \wedge \Delta(b)$,

- $\langle \text{if } b \text{ then } z_1 \text{ else } z_2 \text{ fi}, s \rangle \xrightarrow{i} \langle \text{if } b \text{ then } z_1 \text{ else } z_2 \text{ fi}, s \rangle$ if $s \models_{\pi} \neg \Delta(b)$,
- $\langle \text{while } b \text{ do } z_1 \text{ od}, s \rangle \xrightarrow{i} \langle z_1; \text{while } b \text{ do } z_1 \text{ od}, s \rangle$ if $s \models_{\pi} b \wedge \Delta(b)$,
- $\langle \text{while } b \text{ do } z_1 \text{ od}, s \rangle \xrightarrow{i} \langle \epsilon, s \rangle$ if $s \models_{\pi} \neg b \wedge \Delta(b)$,
- $\langle \text{while } b \text{ do } z_1 \text{ od}, s \rangle \xrightarrow{i} \langle \text{while } b \text{ do } z_1 \text{ od}, s \rangle$ if $s \models_{\pi} \neg \Delta(b)$,
- $\langle \text{await } b \text{ do } z_1 \text{ od}, s_1 \rangle \xrightarrow{i} \langle z_n, s_n \rangle$ if $s_1 \models_{\pi} b \wedge \Delta(b)$, and
 - there is a list of configurations $\langle z_2, s_2 \rangle, \langle z_3, s_3 \rangle, \dots, \langle z_{n-1}, s_{n-1} \rangle$, such that for all $1 < k \leq n$, $\langle z_{k-1}, s_{k-1} \rangle \xrightarrow{i} \langle z_k, s_k \rangle$ and $z_n = \epsilon$,
- $\langle \text{await } b \text{ do } z_1 \text{ od}, s_1 \rangle \xrightarrow{i} \langle \text{await } b \text{ do } z_1 \text{ od}, s_1 \rangle$ if $s_1 \models_{\pi} \neg \Delta(b)$, or if $s_1 \models_{\pi} b$ and
 - there is an infinite list of configurations $\langle z_2, s_2 \rangle, \langle z_3, s_3 \rangle, \dots, \langle z_n, s_n \rangle, \dots$, such that for all $k > 1$, $\langle z_{k-1}, s_{k-1} \rangle \xrightarrow{i} \langle z_k, s_k \rangle$, or
 - there is a finite list of configurations $\langle z_2, s_2 \rangle, \langle z_3, s_3 \rangle, \dots, \langle z_n, s_n \rangle$, where $z_n \neq \epsilon$, there is no configuration $\langle z_{n+1}, s_{n+1} \rangle$ such that $\langle z_n, s_n \rangle \xrightarrow{i} \langle z_{n+1}, s_{n+1} \rangle$, and for all $1 < k \leq n$, $\langle z_{k-1}, s_{k-1} \rangle \xrightarrow{i} \langle z_k, s_k \rangle$.

19.11.4 Modified Rules

Since the assignment-statement aborts if the assigned values are undefined, it is necessary to add an extra premise to the await-rule to make this impossible:

$$\frac{\frac{P^R \Rightarrow \delta(r) \wedge \bigwedge_{a \in \alpha} \delta(u_a)}{\overline{P^R} \wedge v = \overline{r} \wedge I_{\{v\} \cup \alpha} \wedge \bigwedge_{a \in \alpha} a = \overline{u_a} \Rightarrow G \wedge E}}{v := r \text{ sat } (\vartheta, \alpha) :: (P, R, W, G, R|E|R)}$$

The if-statement aborts if the Boolean test is undefined when it is evaluated. Since the truth-value of the Boolean test cannot be changed by the environment, it is enough to give the conclusion's pre-condition an extra conjunct:

$$\frac{z_1 \text{ sat } (\vartheta, \alpha) :: (P \wedge b, R, W, G, E) \quad z_2 \text{ sat } (\vartheta, \alpha) :: (P \wedge \neg b, R, W, G, E)}{\text{if } b \text{ then } z_1 \text{ else } z_2 \text{ fi sat } (\vartheta, \alpha) :: (P \wedge \delta(b), R, W, G, E)}$$

In the case of the while-statement, the Boolean test may be evaluated any number of times. Thus the second premise must be changed to make sure that if the Boolean test evaluates to true, then it is also defined the next time it is evaluated, in which case it is enough to strengthen the conclusion's pre-condition in the same way as above:

$$\frac{\text{wf } Z \quad z \text{ sat } (\vartheta, \alpha) :: (P \wedge b, R, W, G, P \wedge \delta(b) \wedge Z)}{\text{while } b \text{ do } z \text{ od sat } (\vartheta, \alpha) :: (P \wedge \delta(b), R, W, G, (Z^\dagger \vee R) \wedge \neg b)}$$

It is a bit more difficult to formulate the await-rule:

$$\frac{P^R \Rightarrow (W \wedge \neg b) \vee b \quad E_1 | (\bigwedge_{a \in \alpha} a = \overline{u_a} \wedge I_\alpha) \Rightarrow G \wedge E_2 \quad z \text{ sat } (\vartheta, \alpha) :: (P^R \wedge b, I, \text{false}, \text{true}, E_1 \wedge \bigwedge_{a \in \alpha} \delta(u_a))}{\text{await } b \text{ do } z \text{ od sat } (\vartheta, \alpha) :: (P, R, W, G, R | E_2 | R)}$$

The first premise ensures that the statement can only become blocked in a state which satisfies W , and that the Boolean test is defined when it is evaluated. The third premise guarantees that the expressions assigned to the auxiliary variables are defined.

19.12 CSP

In this thesis we are only concerned with shared-state concurrency. This does not mean that similar strategies cannot be employed to develop communication-based programs.

Levin and Gries have suggested a proof-system [LG81] which is closely related to the method for shared-state concurrency described in [OG76a]. Since LSP can be seen as a compositional version of the latter, it should be fairly obvious that our system can be transformed into a compositional version of the Levin/Gries approach.

19.13 Without Reflexivity and Transitivity Constraints

19.13.1 Modifications

The rely-condition has so far been constrained to be both reflexive and transitive, while the guar-condition has been required to be reflexive. The reason why we introduced these requirements is that they have a slightly simplifying effect on both the decomposition-rules and the meaning of a specification. On the other hand, the obvious disadvantage is that some specifications become more complicated.

However, these constraints can easily be removed by redefining the rely- and guar-conditions to denote respectively any atomic state change by the environment and any atomic state change by the implementation. This means that *ext* and *int* (see page 69) must be reformulated as below:

Definition 26 *Given a glo-set ϑ , a pre-condition P , a rely-condition R , and a structure π , then $\text{ext}_\pi[\vartheta, P, R]$ denotes the set of all computations σ in π , such that:*

- $\delta(\sigma_1) \models_\pi P$,
- for all $1 \leq j < \text{len}(\sigma)$, if $\lambda(\sigma_j) = e$ and $\delta(\sigma_j) \stackrel{\vartheta}{\neq} \delta(\sigma_{j+1})$ then

$$(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_\pi R,$$

- if $\text{len}(\sigma) = \infty$, then for all $j \geq 1$, there is a $k \geq j$, such that $\lambda(\sigma_k) = i$.

Definition 27 *Given a glo-set ϑ , a wait-condition W , a guar-condition G , an eff-condition E , and a structure π , then $\text{int}_\pi[\vartheta, W, G, E]$ denotes the set of all computations σ in π , such that:*

- $\text{len}(\sigma) \neq \infty$,
- for all $1 \leq j < \text{len}(\sigma)$, if $\lambda(\sigma_j) = i$ and $\delta(\sigma_j) \stackrel{\vartheta}{\neq} \delta(\sigma_{j+1})$ then

$$(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_\pi G,$$

- if $\tau(\sigma_{len(\sigma)}) \neq \epsilon$ then $\delta(\sigma_{len(\sigma)}) \models_{\pi} W$,
- if $\tau(\sigma_{len(\sigma)}) = \epsilon$ then $(\delta(\sigma_1), \delta(\sigma_{len(\sigma)})) \models_{\pi} E$.

Satisfaction can then be defined as earlier.

19.13.2 Alternative Rules

It is only necessary to change the skip-, assignment-, while-, await- and effect-rules (observe that P^R is equivalent to P^{R^*}):

$$\text{skip } \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, R^*)$$

$$\frac{\overline{P^R} \wedge v = \overline{r} \wedge I_{\{v\} \cup \alpha} \wedge \bigwedge_{a \in \alpha} a = \overline{u_a} \Rightarrow (G \vee I) \wedge E}{v := r \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, R^* | E | R^*)}$$

$$\text{where for all } a \in \alpha, \text{var}[u_a] \subseteq \vartheta \cup \{a\}$$

$$\frac{\underline{\text{wf}} Z}{z \underline{\text{sat}} (\vartheta, \alpha) :: (P \wedge b, R, W, G, P \wedge Z)} \\ \text{while } b \text{ do } z \text{ od } \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, (Z^\dagger \vee R^*) \wedge \neg b)$$

$$\frac{P^R \wedge \neg b \Rightarrow W}{E_1 | (\bigwedge_{a \in \alpha} a = \overline{u_a} \wedge I_\alpha) \Rightarrow (G \vee I) \wedge E_2} \\ z \underline{\text{sat}} (\vartheta, \alpha) :: (P^R \wedge b, I, \text{false}, \text{true}, E_1) \\ \text{await } b \text{ do } z \text{ od } \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, R^* | E_2 | R^*)$$

$$\text{where for all } a \in \alpha, \text{var}[u_a] \subseteq \vartheta \cup \{a\}$$

$$\frac{z \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E)}{z \underline{\text{sat}} (\vartheta, \alpha) :: (P, R, W, G, E \wedge (R \vee G)^*)}$$

19.14 Allowing Environments to Diverge

As mentioned earlier, the assumption that the environment is convergent can easily be removed. It is enough to change the definitions of *ext* (see page 68) and *int* (see page 69) to:

Definition 28 *Given a pre-condition P , a rely-condition R , and a structure π , then $\text{ext}_\pi[P, R]$ denotes the set of all computations σ in π , such that:*

- $\delta(\sigma_1) \models_\pi P$,
- for all $1 \leq j < \text{len}(\sigma)$, if $\lambda(\sigma_j) = e$ then $(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_\pi R$.

Definition 29 *Given a wait-condition W , a guar-condition G , an eff-condition E , and a structure π , then $\text{int}_\pi[W, G, E]$ denotes the set of all computations σ in π , such that:*

- if $\text{len}(\sigma) = \infty$, then there is a $j \geq 1$, such that for all $k \geq j$, $\lambda(\sigma_k) = e$,
- for all $1 \leq j < \text{len}(\sigma)$, if $\lambda(\sigma_j) = i$ then $(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_\pi G$,
- if $\text{len}(\sigma) \neq \infty$ and $\tau(\sigma_{\text{len}(\sigma)}) \neq \epsilon$ then $\delta(\sigma_{\text{len}(\sigma)}) \models_\pi W$,
- if $\text{len}(\sigma) \neq \infty$ and $\tau(\sigma_{\text{len}(\sigma)}) = \epsilon$ then $(\delta(\sigma_1), \delta(\sigma_{\text{len}(\sigma)})) \models_\pi E^2$.

Observe, that *int* is not constraining the infinite elements of *ext* that have only a finite number of internal transitions. Thus the only difference from earlier is that we can no longer use LSP to prove termination, but only that a program terminates or is infinitely overtaken by the environment. (This is related to the view taken in [XH90].)

19.15 Wait as an Assumption, Not as a Commitment

So far the wait-condition has characterised a commitment. It is also possible to interpret the wait-condition as an assumption on the environment.

²This constraint is sufficient, because if there are an infinite computation σ and $j \geq 1$, such that $\tau(\sigma_j) = \epsilon$, then $\sigma(1, \dots, j)$ is a computation.

It is enough to strengthen the definition of a specification (see page 57) with another constraint³; namely that the environment can always reach a state which does not satisfy the wait-condition, and thereafter redefine *ext* and *int* (see page 69) as below:

Definition 30 *A specification is a tuple of the form*

$$(\vartheta, \alpha) :: (P, R, W, G, E),$$

where R , G and E are binary assertions, P and W are unary assertions, ϑ and α are finite, disjoint sets of variables, such that for all structures π :

- R is reflexive and transitive,
- G is reflexive,
- for any state s_1 , there is a state s_2 , such that $(s_1, s_2) \models_{\pi} R$ and $s_2 \models_{\pi} \neg W$,
- the unhooked version of any free variable occurring in (P, R, W, G, E) is an element of $\vartheta \cup \alpha$.

Definition 31 *Given a pre-condition P , a rely-condition R , a wait-condition W , and a structure π , then $\text{ext}_{\pi}[P, R, W]$ denotes the set of all computations σ in π , such that:*

- $\delta(\sigma_1) \models_{\pi} P$,
- for all $1 \leq j < \text{len}(\sigma)$, if $\lambda(\sigma_j) = e$ then $(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_{\pi} R$,
- if $\text{len}(\sigma) = \infty$, then for all $j \geq 1$, there is a $k \geq j$, such that $\lambda(\sigma_k) = i$,
- if $\text{len}(\sigma) \neq \infty$ and $\tau(\sigma_{\text{len}(\sigma)}) \neq \epsilon$, then $\delta(\sigma_{\text{len}(\sigma)}) \models_{\pi} \neg W$.

Definition 32 *Given a guar-condition G , an eff-condition E , and a structure π , then $\text{int}_{\pi}[G, E]$ denotes the set of all computations σ in π , such that:*

- $\text{len}(\sigma) \neq \infty$,

³This constraint can be dropped. It was introduced to ensure that the fourth assumption in definition 31 is consistent with the second assumption. (This is related to implementability — see page 242.)

- $\tau(\sigma_{len(\sigma)}) = \epsilon$,
- for all $1 \leq j < len(\sigma)$, if $\lambda(\sigma_j) = i$ then $(\delta(\sigma_j), \delta(\sigma_{j+1})) \models_{\pi} G$,
- $(\delta(\sigma_1), \delta(\sigma_{len(\sigma)})) \models_{\pi} E$.

The decomposition-rules are the same as before, but due to the new constraint on a specification, it is no longer possible for a program to become blocked in a state in which it cannot be released by the environment. Moreover, due to the new assumption on the environment, an implementation can be thought of as totally correct even when the wait-condition is not equivalent to false.

This means that, in the same way as the wait-condition characterises the set of states in which it is safe for the implementation to become blocked, it is safe for the environment to become blocked in a state which neither satisfies the wait- nor the eff-condition (with all free hooked variables bound by existential quantifiers).

In other words, the wait- and eff-conditions generate an assumption/commitment relationship similar to the one between the rely- and guar-conditions.

19.16 Proof-Obligations on Specifications

19.16.1 Implementability

In VDM [Jon86] specifications are required to satisfy an obligation called implementability⁴; namely that for any state which satisfies the pre-condition, there is a state transition which satisfies the eff-condition. Since we are only interested in total correctness, it is sensible to formulate a similar implementability constraint for our system.

More formally, given a specification $(\vartheta, \alpha) :: (P, R, W, G, E)$, this means that

- for all structures π and all states s_1 , $s_1 \models P$ implies there is a state s_2 such that $(s_1, s_2) \models_{\pi} E$.

By insisting that specifications satisfy this constraint we may eliminate some inconsistencies, but far from all. Consider for example the following specification:

⁴In [Jon90] the same obligation is called satisfiability.

glo	$\{x\}$
aux	$\{\}$
pre	true
rely	$x = \frac{_}{x}$
wait	false
guar	$x \geq \frac{_}{x}$
eff	$x < \frac{_}{x}.$

Clearly, the eff-condition is defined for any state which satisfies the pre-condition. Nevertheless, this specification does not make sense, since none of the alternatives offered by the eff-condition is ‘reachable’ by the transitive closure of the rely- and guar-conditions.

One possibility is instead to employ the constraint⁵:

- for all structures π and all states s_1 , $s_1 \models P$ implies there is a state s_2 such that $(s_1, s_2) \models_{\pi} E \wedge (G \vee R)^{\dagger}$.

This disqualifies the specification above, but it is not difficult to find cases where a still stronger constraint is desirable. There is for example no program which satisfies the specification

glo	$\{x\}$
aux	$\{\}$
pre	$x = 0$
rely	true
wait	false
guar	$x \geq \frac{_}{x}$
eff	$x = 2,$

because the environment can update x as it likes.

It may also be argued that since LSP is a method for the development of totally correct programs, an implementability constraint should restrict specifications in such a way that an implementation cannot become blocked in a state in which it cannot be released by the environment.

It is still an open question how to formulate a sufficiently strong requirement, and how difficult it would be to prove that a specification satisfies such a constraint. See [GR89] for a further discussion.

⁵This was suggested in [WD88] with respect to Jones’ rely/guar-method.

19.17 Related Work

19.17.1 Owicki/Gries

LSP can be understood as a compositional version of the proof system proposed by Owicki and Gries in [OG76a]⁶. The rely-, guar- and wait-conditions have been introduced to avoid the final non-interference and freedom-from-deadlock proofs.

The handling of auxiliary variables has also been changed. Auxiliary variables do not have to be implemented. Constraints on their use are built into the decomposition-rules. Moreover, in the Owicki/Gries approach auxiliary variables can only be used as a verification tool, while in LSP auxiliary variables can be employed both as verification and specification tools.

19.17.2 Jones

It is correct to consider Jones' system [Jon83a] as a restricted version of our system. There are two main differences. First of all, LSP has a wait-condition which makes it possible to deal with synchronisation. Secondly, because auxiliary variables may be employed both as specification and verification tools, LSP is more expressive.

19.17.3 Soundararajan

One important advantage of LSP with respect to Soundararajan's method [Sou84] is that auxiliary variables can be of any sort. It is shown in [Owe90] how reasoning with this system can be simplified by splitting up the original traces into history variables. However, the use of auxiliary variables is still more restricted than in LSP.

Another obvious disadvantage with the Soundararajan approach is that it only deals with partial correctness.

19.17.4 Barringer/Kuiper/Pnueli

The system proposed by Barringer, Kuiper and Pnueli [BKP84] is much more general than ours. However, we believe that our approach is conceptually

⁶The additional interference-freedom requirement for total correctness proposed by Qwicki and Gries is not correct [AdBO90].

simpler and therefore better suited in the area where it can be employed.

Another advantage with LSP is that it makes a clearer distinction between assumptions (about the environment) and commitments (which the implementation is required to satisfy).

19.17.5 Stirling

We have already discussed some of the differences between Stirling's method [Sti88] and LSP, namely that the rely- and guar-conditions are sets of invariants and not binary assertions, that the eff-condition is unary and not binary, and that interference after the last internal transition is treated differently.

Another important difference is that Stirling is only concerned with partial correctness. Moreover, his system does not allow auxiliary variables to be employed as a specification tool.

Bibliography

- [Abr79] K. Abrahamson. Modal logic of concurrent programs. In *Proc. Semantics of Concurrent Computation*, Lecture Notes in Computer Science 70, pages 21–33. Springer, 1979.
- [Acz83] P. Aczel. On an inference rule for parallel composition. Unpublished Paper, February 1983.
- [AdBO90] K.R. Apt, F.S. de Boer, and E.R. Olderog. Proving termination of parallel programs. In W.H.J. Feijen, A.J.M. van Gasteren, D. Gries, and J. Misra, editors, *Beauty Is Our Business, A Birthday Salute to Edsger W. Dijkstra*. Springer-Verlag, 1990.
- [AFdR80] K. R. Apt, N. Francez, and W. P. de Roever. A proof system for communicating sequential processes. *ACM Transactions on Programming Languages and Systems*, 2(3):359–385, 1980.
- [Ak89] J. A. Ah-kee. *Operation Decomposition Proof Obligations for Blocks and Procedures*. PhD thesis, University of Manchester, 1989.
- [AL88] M. Abadi and L. Lamport. The existence of refinement mappings. Technical Report 29, Digital, Palo Alto, 1988.
- [Ame89] P. America. Issues in the design of a parallel object-oriented language. *Formal Aspects of Computing*, 1(4):366–411, 1989.
- [AP86] K. R. Apt and G. D. Plotkin. Countable nondeterminism and random assignment. *Journal of the ACM*, 33(4):724–767, 1986.
- [Apt81] K. R. Apt. Ten years of Hoare’s logic: A survey – part I. *ACM Transactions on Programming Languages and Systems*, 3(4):431–483, 1981.

- [Apt84] K. R. Apt. Ten years of Hoare's logic: A survey, part II, non-determinism. *Theoretical Computer Science*, 28:83–109, 1984.
- [AS85] B. Alpern and F. B. Schneider. Defining liveness. *Information Processing Letters*, 21:181–185, 1985.
- [AS89] B. Alpern and F. B. Schneider. Verifying temporal properties without temporal logic. *ACM Transactions on Programming Languages and Systems*, 11(1):147–167, 1989.
- [Bac59] J. Backus. The syntax and semantics of the proposed international algebraic language of the Zurich ACM-GAMM conference. In *Proc. Int. Conf. Information Processing*, pages 125–132, 1959.
- [Bac88] R. J. R. Back. Refining atomicity in parallel algorithms. Technical Report 57, Åbo Akademi, 1988.
- [Bar85] H. Barringer. *A Survey of Verification Techniques for Parallel Programs*, volume 191 of *Lecture Notes in Computer Science*. Springer-Verlag, 1985.
- [BCJ84] H. Barringer, J. H. Cheng, and C. B. Jones. A logic covering undefinedness in program proofs. *Acta Informatica*, 21:251–269, 1984.
- [BH73] P. Brinch Hansen. Concurrent programming concepts. *Com. Sur.*, 5(4):223–245, 1973.
- [BK84] H. Barringer and R. Kuiper. Towards the hierarchical temporal logic specification of concurrent systems. In *Proc. The Analysis of Concurrent Systems*, Lecture Notes in Computer Science 207, pages 157–184. Springer, 1984.
- [BKP84] H. Barringer, R. Kuiper, and A. Pnueli. Now you may compose temporal logic specifications. In *Proc. Sixteenth ACM Symposium on Theory of Computing*, pages 51–63, 1984.
- [BKP85] H. Barringer, R. Kuiper, and A. Pnueli. A compositional temporal approach to a CSP-like language. In N.J. Neuhold and G. Chroust, editors, *Formal Models in Programming*. Elsevier Science Publishers B.V., 1985.

- [Bro89] M. Broy. Towards a design methodology for distributed systems. In *Proc. Constructive Methods in Computing Science, Summer-school, Marktoberdorf*, pages 311–364. Springer, 1989.
- [BS90] R. J. R. Back and K. Sere. Stepwise refinement of parallel algorithms. *Science of Computer Programming*, 13:133–180, 1989/90.
- [Bur74] R. M. Burstall. Program proving as hand simulation with a little induction. In *Proc. Information Processing 74*, pages 308–312, 1974.
- [CJ90] J. H. Cheng and C. B. Jones. On the usability of logics which handle partial functions. In C. Morgan and J. Woodcock, editors, *Proc. Third Refinement Workshop*, 1990.
- [Cli73] M. Clint. Program proving: Coroutines. *Acta Informatica*, 2:50–63, 1973.
- [CM88] K. M. Chandy and J. Misra. *Parallel Program Design, A Foundation*. Addison-Wesley, 1988.
- [Coo78] S. A. Cook. Soundness and completeness of an axiom system for program verification. *SIAM Journal on Computing*, 7:70–90, 1978.
- [dB80] J. de Bakker. *Mathematical Theory of Program Correctness*. Prentice-Hall, 1980.
- [DH72] O.J. Dahl and C. A. R. Hoare. Hierarchical program structures. In O.J. Dahl, E. W. Dijkstra, and C. A. R. Hoare, editors, *Structured Programming*. Academic Press, Inc, 1972.
- [Dij65] E. W. Dijkstra. Programming considered as a human activity. In *Proc. 1965 IFIP Congress*, pages 213–217, 1965.
- [Dij68] E. W. Dijkstra. Cooperating sequential processes. In F. Genuys, editor, *Programming Languages*. Academic Press, Inc, 1968.
- [Dij75] E. W. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Communications of the ACM*, 18(8):453–457, 1975.

- [Dij76] E.W. Dijkstra. *A Discipline of Programming*. Prentice-Hall International, 1976.
- [Dij82] E. W. Dijkstra. A correctness proof for communicating processes: A small exercise. In *Selected Writings on Computing: A Personal Perspective*. Springer-Verlag, 1982.
- [dR85] W. P. de Roever. The quest for compositionality, formal models in programming. In F. J. Neuhold and G. Chroust, editors, *Proc. IFIP 85*, pages 181–205, 1985.
- [ELW74] B. Elspas, K. N. Levitt, and R. J. Waldinger. An interactive system for the verification of computer programs. Technical report, Menlo Park, 1974.
- [Flo67] R. W. Floyd. Assigning meaning to programs. In Schwartz J. T., editor, *Proc. Symposium in Applied Mathematics, 19*, pages 19–32, 1967.
- [FP78] N. Francez and A. Pnueli. A proof method for cyclic programs. *Acta Informatica*, 9:133–157, 1978.
- [Fra76] N. Francez. *The Analysis of Cyclic Programs*. PhD thesis, The Weizmann Institute of Science, 1976.
- [Fra86] N. Francez. *Fairness*. Springer-Verlag, 1986.
- [GMW79] M. Gordon, R. Milner, and C. Wadsworth. *Edinburgh LCF*, volume 78 of *Lecture Notes in Computer Science*. Springer-Verlag, 1979.
- [Gol90] D. M. Goldschlag. Mechanizing Unity. In *Proc. IFIP Working Conference on Programming Concepts and Methods*, pages 374–401, 1990.
- [GQNL90] P. Grønning, T. Qvist Nielsen, and H. H. Løvengreen. Refinement and composition of transition-based rely-guarantee specifications with auxiliary variables. Unpublished Paper, May 1990.
- [GR89] D. Grosvenor and A. Robinson. An evaluation of rely-guarantee. Unpublished Paper, March 1989.

- [Har79] D. Harel. *First Order Dynamic Logic*, volume 68 of *Lecture Notes in Computer Science*. Springer-Verlag, 1979.
- [HHS86] J. He, C. A. R. Hoare, and J. W. Sanders. Data refinement refined. In *Proc. 1st ESOP, Lecture Notes in Computer Science 213*, pages 187–196, 1986.
- [HM87] E. C. R. Hehner and A. J. Malton. Termination conventions and comparative semantics. In *Proc. Logic of Programming and Calculi of Discrete Design, International Summer School, Marktoberdorf*, pages 79–97, 1987.
- [HO80] B. T. Hailpern and S. S. Owicki. Verifying network protocols using temporal logic. Technical Report 192, Stanford University, 1980.
- [Hoa69] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–583, 1969.
- [Hoa72a] C. A. R. Hoare. Notes on data structuring. In O.J. Dahl, E. W. Dijkstra, and C. A. R. Hoare, editors, *Structured Programming*, pages 175–220. Academic Press, Inc, 1972.
- [Hoa72b] C. A. R. Hoare. Proof of correctness of data representations. *Acta Informatica*, 1:271–282, 1972.
- [Hoa78] C. A. R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8):666–678, 1978.
- [Hoa85] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall International, 1985.
- [HP73] P. Hitchcock and D. Park. Induction rules and termination proofs. In M. Nivat, editor, *Proc. 1st ICALP Symp.*, pages 225–251, 1973.
- [HP79] M. Hennessy and G. Plotkin. Full abstraction for a simple parallel programming language. In *Proc. Mathematical Foundations of Computer Science, Lecture Notes in Computer Science 74*, pages 108–120, 1979.

- [Jon72] C. B. Jones. Formal development of correct algorithms: An example based on Earley's recogniser. In *Proc. ACM Conference on Proving Assertions about Programs, SIGPLAN Notices 7*, pages 150–169, 1972.
- [Jon80] C. B. Jones. *Software Development: A Rigorous Approach*. Prentice-Hall International, 1980.
- [Jon81] C. B. Jones. *Development Methods for Computer Programs Including a Notion of Interference*. PhD thesis, Oxford University, 1981.
- [Jon83a] C. B. Jones. Specification and design of (parallel) programs. In R.E.A. Mason, editor, *Proc. Information Processing 83*, pages 321–331. North-Holland, 1983.
- [Jon83b] C. B. Jones. Tentative steps towards a development method for interfering programs. *ACM Transactions on Programming Languages and Systems*, 5(4):576–619, 1983.
- [Jon86] C. B. Jones. *Systematic Software Development Using VDM*. Prentice-Hall International, 1986.
- [Jon90] C. B. Jones. *Systematic Software Development Using VDM, Second Edition*. Prentice-Hall International, 1990.
- [KKZ89] R. Koymans, R. Kuiper, and E. Zijlstra. Specification specified. In: R. Kuiper, *Combining Linear Time Temporal Logic Descriptions of Concurrent Computations*, PhD. thesis, University of Eindhoven, 1989.
- [KM75] S. Katz and Z. Manna. A closer look at termination. *Acta Informatica*, 5:333–352, 1975.
- [Knu68] D. E. Knuth. *Fundamental Algorithms*, volume 1 of *The Art of Computer Programming*. Addison Wesley, 1968.
- [Lam80] L. Lamport. The 'Hoare logic' of concurrent programs. *Acta Informatica*, 14:21–37, 1980.

- [Lam83] L. Lamport. Specifying concurrent program modules. *ACM Transactions on Programming Languages and Systems*, 5(2):190–222, 1983.
- [Lam85] L. Lamport. An axiomatic semantics of concurrent programming languages. In K. R. Apt, editor, *Logics and Tools for Concurrent Systems, NATO ASI Series Vol. F13*, 1985.
- [Lam90] L. Lamport. A temporal logic of actions. Technical Report 57, Digital, Palo Alto, 1990.
- [Lau73] H. C. Lauer. *Correctness in Operating Systems*. PhD thesis, Carnegie-Mellon University, 1973.
- [Len82] C. Lengauer. *A Methodology for Programming with Concurrency*. PhD thesis, University of Toronto, 1982.
- [LG81] G. M. Levin and D. Gries. A proof technique for communicating sequential processes. *Acta Informatica*, 15:281–302, 1981.
- [LPS81] D. Lehmann, A. Pnueli, and J. Stavi. Impartiality, justice and fairness: The ethics of concurrent termination. In *Proc. Automata, Languages, and Programming, Lecture Notes in Computer Science 115*, pages 264–277, 1981.
- [LS84] L. Lamport and F. Schneider. The ‘Hoare logic’ of CSP, and all that. *ACM Transactions on Programming Languages and Systems*, 6(2):281–296, 1984.
- [MC81] J. Misra and K. M. Chandy. Proofs of networks of processes. *IEEE Transactions on Software Engineering*, 7(4), 1981.
- [MCS82] J. Misra, K. M. Chandy, and T. Smith. Proving safety and liveness of communicating processes with examples. In *Proc. ACM Symposium on Principles of Distributed Computing*, pages 201–208, 1982.
- [Mid89] K. Middelburg. VVSL: A language for structured VDM specifications. *Formal Aspects of Computing*, 1(1):115–135, 1989.

- [Mil71] R. Milner. An algebraic definition of simulation between programs. In *Proc. 2nd International Joint Conference on Artificial Intelligence*, 1971.
- [Mos74] Y. N. Moschovakis. *Elementary Induction on Abstract Structures*. North-Holland, 1974.
- [Mos86] B. Moszkowski. *Executing Temporal Logic Programs*. North-Holland, 1986.
- [MP83] Z. Manna and A. Pnueli. Verification of concurrent programs: A temporal proof system. In J. W. de Bakker and J. van Leuwen, editors, *Proc. Foundations of Computer Science 4, Distributed Systems: Part 2*, pages 163–255, 1983.
- [MP84] Z. Manna and A. Pnueli. Adequate proof principles for invariance and liveness properties of concurrent programs. *Science of Computer Programming*, 4(3):257–289, 1984.
- [Nau66] P. Naur. Proofs of algorithms by general snapshots. *BIT*, 6:310–316, 1966.
- [Nip86] T Nipkow. Non-deterministic data types: Models and implementations. *Acta Informatica*, 22:629–661, 1986.
- [OG76a] S. Owicki and D. Gries. An axiomatic proof technique for parallel programs. *Acta Informatica*, 6:319–340, 1976.
- [OG76b] S. Owicki and D. Gries. Verifying properties of parallel programs: An axiomatic approach. *Communications of the ACM*, 19(5):279–285, 1976.
- [OL82] S. Owicki and L. Lamport. Proving liveness properties of concurrent programs. *ACM Transactions on Programming Languages and Systems*, 4:455–495, 1982.
- [Owe85] O. Owe. An approach to program reasoning based on a first order logic for partial functions. Technical Report 89, University of Oslo, 1985.
- [Owe90] O. Owe. Axiomatic treatment of processes with shared variables revisited. Unpublished Paper, April 1990.

- [Owi75] S. Owicki. *Axiomatic Proof Techniques for Parallel Programs*. PhD thesis, Cornell University, 1975.
- [Pan90] P. K. Pandya. Some comments on the assumption-commitment framework for compositional verification of distributed programs. In J. W. de Bakker, W. P. de Roever, and G. Rozenberg, editors, *Proc. REX Workshop on Stepwise Refinement of Distributed Systems, Lecture Notes in Computer Science 430*, pages 622–640, 1990.
- [PJ87] P.K. Pandya and M. Joseph. P-a logic – a compositional proof system for distributed programs. Technical report, Univ. of Warwick, 1987.
- [Plo81] G. D. Plotkin. A structural approach to operational semantics. Technical Report DAIMI FN-19, Aarhus University, 1981.
- [SdRG89] F. A. Stomp, W. P. de Roever, and R. T. Gerth. The μ -calculus as an assertion-language for fairness arguments. *Information and Computation*, 82:278–322, 1989.
- [Sho67] J. R. Shoenfield. *Mathematical Logic*. Addison Wesley, 1967.
- [SMSV83] R. L. Schwartz, P. M. Melliar-Smith, and F. H. Vogt. An interval logic for higher-level temporal reasoning. In *Proc. Second Annual ACM Symposium on Principles of Distributed Computing*, pages 173–186, 1983.
- [Sou84] N. Soundararajan. A proof technique for parallel programs. *Theoretical Computer Science*, 31:13–29, 1984.
- [Spi88] J. M. Spivey. *Understanding Z, A Specification Language and its Formal Semantics*, volume 3 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1988.
- [Sta85] E. W. Stark. A proof technique for rely/guarantee properties. In S. N. Maheshwari, editor, *Proc. 5th Conference on the Foundation of Software Technology and Theoretical Computer Science, Lecture Notes in Computer Science 206*, pages 369–391, 1985.

- [Sta88] E. W. Stark. Proving entailment between conceptual state specifications. *Theoretical Computer Science*, 56:135–154, 1988.
- [Sti88] C. Stirling. A generalization of Owicki-Gries’s Hoare logic for a concurrent while language. *Theoretical Computer Science*, 58:347–359, 1988.
- [Stø90] K. Stølen. *Development of Parallel Programs on Shared Data-Structures*. PhD thesis, University of Manchester, 1990.
- [Tur49] A. Turing. On checking a large routine. In *Proc. Conference on high-speed automatic calculating machines, University Mathematical Laboratory, Cambridge*, 1949.
- [Wan78] M. Wand. A new incompleteness result for Hoare’s system. *Journal of the ACM*, 25(1):168–175, 1978.
- [WD88] J. C. P. Woodcock and B. Dickinson. Using VDM with rely and guarantee-conditions. Experiences from a real project. In R. Bloomfield, L. Marshall, and R. Jones, editors, *Proc. VDM’88, Lecture Notes in Computer Science 328*, pages 434–458, 1988.
- [Wir71] N. Wirth. The development of programs by stepwise refinement. *Communications of the ACM*, 14:221–227, 1971.
- [XH90] Q. Xu and J. He. Towards a theory of interfering programming. Unpublished Paper, January 1990.
- [ZH81] C. C. Zhou and C. A. R. Hoare. Partial correctness of communicating processes and protocols. Technical Report PRG-20, Oxford University, 1981.
- [Zwi89] J. Zwiers. *Compositionality, Concurrency and Partial Correctness: Proof Theories for Networks of Processes and Their Relationship*, volume 321 of *Lecture Notes in Computer Science*. Springer-Verlag, 1989.