

CWI Tracts

Managing Editors

K.R. Apt (CWI, Amsterdam)

M. Hazewinkel (CWI, Amsterdam)

N.M. Temme (CWI, Amsterdam)

J.M. Schumacher (CWI, Amsterdam)

Executive Editor

M. Bakker (CWI Amsterdam, e-mail: Miente.Bakker@cwi.nl)

Editorial Board

W. Albers (Enschede)

P.C. Baayen (Amsterdam)

R.C. Backhouse (Eindhoven)

E.M. de Jager (Amsterdam)

M.A. Kaashoek (Amsterdam)

M.S. Keane (Amsterdam)

H. Kwakernaak (Enschede)

J. van Leeuwen (Utrecht)

P.W.H. Lemmens (Utrecht)

M. van der Put (Groningen)

M. Rem (Eindhoven)

H.J. Sips (Delft)

M.N. Spijker (Leiden)

H.C. Tijms (Amsterdam)

CWI

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

Telephone +31 - 20 592 9333

Telefax +31 - 20 592 4199

CWI is the nationally funded Dutch institute for research in Mathematics and Computer Science.

Comparative concurrency semantics and refinement of actions

R.J. van Glabbeek

1991 Mathematics Subject Classification: 68Q55 (Semantics).
ISBN 90 6196 454 7
NUGI-code: 811

Copyright ©1996, Stichting Mathematisch Centrum, Amsterdam
Printed in the Netherlands

Contents

Introduction	5
1. The linear time - branching time spectrum	17
1. Semantic equivalences on labelled transition systems	22
2. The semantic lattice	38
3. Complete axiomatizations	50
2. Modular specifications in process algebra - with curious queues	55
1. Module logic	61
2. Process algebra	66
3. Applications of the module approach in process algebra	80
4. Queues	89
5. A protocol verification	102
Appendix: Logics	111
3. Branching time and abstraction in bisimulation semantics	117
1. Branching and abstraction	120
2. Axioms	130
3. Branches and traces	143
4. Completeness proofs	149
5. Correspondence	158
6. Refinement	159
7. Divergence	163
8. Modal characterizations	164
4. Refinement of actions in causality based models	173
1. Refinement of actions in prime event structures	181
2. Refinement of actions in flow event structures	186
3. Configuration structures and refinement of actions	193
4. Refinement of transitions in Petri nets	197
5. Partial order semantics for refinement of actions - neither necessary nor always sufficient but appropriate when used with care -	213
6. Equivalence notions for concurrent systems and refinement of actions	223
1. Interleaving semantics	225
2. Step semantics	227
3. 'Linear time' partial order semantics	230
4. 'Branching time' partial order semantics	231
7. The refinement theorem for ST-bisimulation semantics	243
1. Concurrent systems and refinement of actions	247
2. The behaviour of concurrent systems I	249
3. Equivalence notions for concurrent systems I	250
4. The behaviour of concurrent systems II	255
5. Equivalence notions for concurrent systems II	257
6. The refinement theorems	265
References	273

Acknowledgements

Jan Willem Klop has been an ideal supervisor for me. He was always ready to discuss my work, but let me free to pursue my own goals in my own way. It has been a great pleasure to have him as a supervisor.

I also like to express my gratitude to Jan Bergstra for initiating my employment at CWI and organizing a very interesting scientific environment. Moreover I appreciated the many discussions with him, in which he was always full of ideas.

I benefited immensely from the collaboration with Frits Vaandrager. It is in my opinion a rare privilege to discuss ones ideas at the time they are being formed with someone able of grasping them in full and providing feedback. With Frits I could discuss all my work in depth and he contributed numerous useful ideas.

I thank Ursula Goltz for our fruitful cooperation since 1987. Our joint work on action refinement forms an important constituent of this book. This work was initiated by a discussion with Albert Meyer and Ernst-Rüdiger Olderog at ICALP 87 in Karlsruhe, to whom I also express my gratitude.

I am grateful to my colleagues Jos Baeten and Jan Friso Groote for our pleasant collaboration and for their useful comments on earlier versions of parts of this book.

Furthermore I gratefully acknowledge the discussions and correspondence with Alex Rabinovich, Walter Vogler, Wolfgang Reisig, Vaughan Pratt, Tony Hoare, Ilaria Castellani, Rocco De Nicola, Pierpaolo Degano, Ugo Montanari and Gérard Boudol that contributed to the work reported in this book.

I thank the participants of the weekly Process Algebra Meetings (as far as not mentioned already): Wiet Bouma, Jeroen Bruijning, Nicolien Drost, Henk Goeman, Karst Koymans, Sjouke Mauw, Kees Middelburg, Hans Mulder, Alban Ponse, Gerard Renardel de Lavalette, Piet Rodenburg, Gert Veltink, Jos Vrancken, Fer-Jan de Vries, Peter Weijland, Freek Wiedijk and Han Zuidweg, for numerous lively discussions and comments.

More in general, my years at the CWI Department of Software Technology, headed by Jaco de Bakker, will remain a pleasant memory for its excellent working conditions and friendly atmosphere.

It is a pleasure to thank Glynn Winskel for his kind willingness to referee this book.

Finally, special thanks goes to Gertrud Jacobs for her careful preparation of Chapters IV, V and VI of the original manuscript, and for being available whenever she was needed.

Affiliations

Most of the work reported in this book was done when I was employed at the Centre for Mathematics and Computer Science (CWI) in Amsterdam. The manuscript was finalized during my employment at the Technical University of Munich, Postfach 202420, D-8000 München 2. Chapter II is joint work with Frits Vaandrager, and Chapter III with Peter Weijland, both affiliated with the Centre for Mathematics and Computer Science. Chapters IV, V and VI are joint work with Ursula Goltz, Gesellschaft für Mathematik und Datenverarbeitung, Postfach 1240, D-5205 Sankt Augustin 1. My current affiliation is: Computer Science Department, Stanford University, CA 94305, USA.

Support

The work reported in this book has been supported by ESPRIT project no. 432, An Integrated Formal Approach to Industrial Software Development (METEOR) (Chapter II, III, V, VI), RACE project no. 1046, Specification and Programming Environment for Communication Software (SPECS) (Chapter II), ESPRIT Basic Research Action 3006, Theories of Concurrency: Unification and Extension (CONCUR) (Chapter III), ESPRIT Basic Research Action 3148, Design Methods Based on Nets (DEMON) (Introduction, Chapter I, III, IV and VI) and Sonderforschungsbereich 342 of the TU München (Introduction, Chapter I, III and VI). The work on this second edition has been supported by ONR under grant number N00014-92-J-1974. Chapter I was partly written in the preparation of a course Comparative Concurrency Semantics, given at the University of Amsterdam, spring 1988.

Prior Publications

The material presented in the seven chapters of this book can also be found in [79, 92, 96, 86, 85, 87, 78] respectively. Furthermore Chapter II without Sections 4 and 5 appeared in the Ph.D. Thesis of Frits Vaandrager [171] and the first two sections of Chapter III partly appeared in the Ph.D. Thesis of Peter Weijland [182]. Finally, Section 6 of Chapter III appeared as [95].

Introduction

1. Comparative concurrency semantics. This book is about comparative concurrency semantics.

Concurrency is the study of concurrent systems. Often concurrency as area of scientific research is located in computer science. In that case the systems which are the subject of study are taken to be computers or computer programs. However, much theory in the field of concurrency applies equally well to other systems, like machines, elementary particles, protocols, networks of falling dominoes or human beings. *Concurrent* or *parallel* systems - as opposed to *sequential* systems - are systems capable of performing different activities at the same time.

Semantics is the study of the meaning of words. In concurrency, one often employs formal languages for the description of concurrent systems. These I call *system description languages*. Like all formal languages, system description languages are usually introduced to avoid the ambiguities of natural languages and to gain accuracy of expression. Therefore their semantics tends to be easier than the semantics of natural languages. Moreover the meaning of the words in a formal language should to some extent be given by the one who defines the language, rather than to be discovered by linguists.

Since system description languages tend to describe abstractions of systems rather than concrete systems, the meaning of an expression in a system description language is in general given by an *equivalence class* of systems (i.e. a class of systems which are considered to be equivalent on a chosen level of abstraction). Thus the meaning of the entire language is determined by a partition of a set of systems into equivalence classes and an allocation of one such equivalence class to each expression. For this reason it is convenient to divide the semantics of system description languages into two subfields, namely the

study of equivalence relations on sets of concurrent systems, and the study of allocating equivalence classes to expressions in particular languages. The first field deals with the establishment of criteria, determining when two systems are sufficiently alike to be collected in the same equivalence class. It can be studied independently of a particular system description language. Therefore it can be simply referred to as *semantics of concurrent systems* or *concurrency semantics* for short.

In concurrency semantics a criterion, determining when two systems are sufficiently alike to be collected in the same equivalence class, is called a *semantics*, and the induced equivalence relation a *semantic equivalence*. In the literature on concurrency semantics many semantics have been proposed and most likely also a multitude of sensible semantics have never been proposed. The classification of these semantics is called *comparative concurrency semantics* and will be the primary subject of this book.

2. *Design and verification.* Much work in concurrency is motivated by an interest in design problems for concurrent systems. A fruitful method to design concurrent systems is by means of *stepwise refinement*. Here one starts with a description S_0 of the system one has in mind. This initial description is called a *specification* of the desired system. It abstracts from all the details of the desired system that are not essential in its behaviour and leaves open many design decisions that have to be taken later on. Then one starts refining the specification by adding step for step the details one needs to know when the system is going to be built. In this way one obtains a sequence

$$S_0 \rightarrow S_1 \rightarrow \cdots \rightarrow S_n$$

of system descriptions of which the last one says exactly how the system will look like. This final state in the design process describes the *implementation* of the desired system.

Roughly one can distinguish two different kinds of refinement steps in such a sequence of system descriptions. First of all there are steps in which information is added about *what* the system ought to do. These steps concern the goal of the entire exercise and can therefore not be proven correct in terms of this goal. Secondly there are steps that add information about *how* the system is going to do it. It is one of the tasks of concurrency theory to prove the correctness of such steps.

When considering only one step from a stepwise refinement sequence, the left-hand side of this step is called *specification* and the right-hand side *implementation*. Let $S \rightarrow I$ be a 'how'-step. The question is now to find criteria for determining whether or not this step is correct. Here at least two situations can be distinguished:

1. Although I describes much more activities of the desired system than S , all these extra activities can be considered as internal actions in which the user of the system is not interested. After abstraction from all these details, I and S are equivalent according to some suitable semantic equivalence.

2. Some choices about how the final system should behave, that were left open in S , are resolved in I . Therefore I and S cannot be equivalent. Here one needs a partial order between equivalence classes of concurrent systems, specifying when one system is a correct implementation of the other.

In order to tackle both cases one needs to define a suitable semantic equivalence and a partial order on the equivalence classes. Together these ingredients can be coded as a *preorder*, a reflexive and transitive relation, on system descriptions.

In this book, for reasons of convenience, attention is restricted to equivalences rather than arbitrary preorders. However, there exists a close correspondence between semantic equivalences and preorders. Most semantic equivalences are defined, or can be characterized, in terms of the properties that are shared by equivalent systems. For each system p , a set of properties $O(p)$ is defined, such that two systems p and q are equivalent iff $O(p) = O(q)$. Often $O(p)$ describes the observable behaviour of p according to some testing scenario. Now a corresponding preorder $\leq$ can be defined by $p \leq q$ iff $O(p) \subseteq O(q)$. Most preorders encountered in the literature on concurrency are of that form. I expect that using this insight, much work on classifying semantic equivalences can be generalized to preorders.


Above I argued that semantic equivalences (and preorders) can be relevant for the design of concurrent systems. However, in fact they are more often employed for *verification* purposes. In this case one is offered a specification and an implementation of a certain system and is asked to determine if the implementation is correct. In such applications the distance between the specification and the implementation tends to be larger than in one step in a design process. Therefore it is even more important to have solid criteria for deciding on the correctness of the implementation.

When semantic equivalences are used in the design of concurrent systems, or for verification purposes, they should be chosen in such a way that two system descriptions are considered equivalent only if the described behaviours share the properties that are essential in the context in which the system will be embedded. It depends on this context and on the interests of a particular user which properties are essential. Therefore it is not a task of concurrency semantics to find the 'true' semantic equivalence, but rather to determine which equivalence is suitable for which applications. It is the intention of this book to carry out a bit of this task. In particular it addresses the question which semantic equivalences are suitable for dealing with *action refinement*.

3. *Refinement of actions.* In this book concurrent systems are represented by expressions in a system description language or by elements of some mathematical model. The basic building blocks in the languages and models that occur in this book are the *actions* which may be performed by a system. By an *action* here any activity is understood which is considered as a conceptual entity on a chosen level of abstraction. This allows design steps in which actions are replaced by more complex system descriptions. Such a step in the

design of a system is referred to as *refinement of actions*. This book deals with *uniform concurrency* only, meaning that the actions are represented by uninterpreted symbols $a, b, c, \dots$ that are almost treated like variables.¹ In this context, action refinement should be regarded as a design step that adds information about what the system ought to do (a ‘what’-step), at least if the refined actions are not considered to be internal. Therefore the ‘correctness’ of such refinement steps cannot be proven.² However, the possibility of doing such steps puts some restrictions on the kind of equivalences that can be used for proving the correctness of ‘how’-steps occurring in the same design process.

EXAMPLE: Consider the following specification of a concurrent system: ‘The actions a and b should in principle be performed independently on different processors, but if one of the processors happens to be ready with a before the other starts with b , b may also be executed on this processor instead of the other one’. This system description is represented by the Petri net K below.



An introduction to Petri nets and the way they model concurrent systems can be found in REISIG [85] or PETERSON [148]


Suppose that someone comes up with an implementation in which first it is

1. In particular, the convention applies that different occurrences of the same action a are to be refined in the same way. In other words, when the possibility exists that two activities in a specification are going to be implemented differently, they should be given different names.
2. By no means I want to exclude the employment of action refinement as a ‘how’-step, admitting correctness proofs. However, this requires the classification of the refined action as ‘internal’ and falls outside the scope of this book. Doing so may also be a reason to give up the convention adopted above. This would give rise to a completely different theory.

In RENSINK [158] a mode of action refinement is proposed that I would classify as a combined ‘what’- and ‘how’-step. Correctness proofs of the ‘how’-component of such a refinement step are possible, but relative to a preorder that is explicitly parametrized with the ‘what’-component of the refinement step.

determined whether the actions a and b will happen sequentially or independently, and subsequently one of these alternatives will take place, as represented by the Petri net L . Although this implementation does not seem very convincing, it will be considered 'correct' by many equivalences occurring in the literature.

Let the next step in the design process consist of refining the action a in the sequential composition of two actions a_1 and a_2 . From L one thereby obtains the net L' on the right.



If L' is going to be placed in an environment where a_2 becomes causally dependent on b - it may be the case that b is an output action, a_2 is an input action, and the environment needs data from b in order to compute the data that are requested by a_2 - then deadlock can occur. However, if the refinement step splitting a in a_1 and a_2 is carried out on K already, the resulting system K' is deadlock free in the environment sketched above.

Thus the possibility of refining a somehow invalidates the correctness of the design step from K to L . $\square$

A semantic equivalence is said to be *preserved under refinement of actions* if two equivalent processes remain equivalent after replacing all occurrences of an action a by a more complicated process $r(a)$. The example above indicates that for certain applications it may be fruitful to employ equivalences that are preserved under refinement of actions. It is one of the topics of this book to find out which equivalences have this property.

4. *About the contents of this book.* This book consists of seven chapters which are all based on separate papers and have their own introduction. This general introduction is only meant to give an indication of their contents and their role in the book.

In the first chapter several semantic equivalences for concrete sequential systems are presented, and motivated in terms of the observable behaviour of systems, according to some testing scenario. Here *concrete* means that no internal actions or internal choice are considered. These semantics are partially ordered by the relation ‘makes strictly more identifications than’, thus constituting a complete lattice. For ten of these semantic equivalences complete axiomatizations are provided. As in the rest of my book, stochastic and real-time aspects of concurrent systems are completely neglected. Chapter I serves partly to give an overview of the literature on semantic equivalences for concrete sequential processes. The various notions that can be found elsewhere can easily be compared, since they are all presented in the same style, and using the same formalism. In order for the semantics of this chapter to be applicable for design and verification purposes, they have to be generalized to a setting with internal moves, and with parallelism. This can be done in many ways. In the last two chapters the two extreme points on the semantic lattice, trace semantics and bisimulation semantics, are generalized to a setting with parallelism and in Chapter III, bisimulation semantics is generalized to a setting with internal moves.

In the second chapter it is shown how semantic notions can be used in protocol verification and other applications. This chapter is entirely algebraic in style and employs axiom systems of which only classes of models are considered, rather than a particular model. It is based on the Algebra of Communicating Processes of BERGSTRÅ & KLOP [25,28]. In order to combine axiom systems representing semantic notions that are difficult to combine a new notion of ‘proof’ is developed.

The third chapter is devoted to the generalization of bisimulation equivalence to a setting with silent moves. It is argued that the solution of MILNER [134] (observation equivalence) does not respect the branching structure of processes and hence lacks an important feature of bisimulation semantics without internal moves. A finer equivalence is proposed which indeed respects branching structure. This new *branching bisimulation equivalence* turns out to have some practical advantages as well. In particular, we show that in a setting without parallelism it is preserved under refinement of actions, whereas observation equivalence is not.

In the fourth chapter an operator for refinement of actions is defined on four causality based models for concurrent systems, namely on three kinds of event structures and on Petri nets, and in the remaining three chapters it is investigated which of the ‘linear time’ and ‘branching time’ semantic equivalences proposed in the literature are preserved under refinement of actions and which are not. Chapter V can be regarded as an informal summary of the Chapters VI and VII. It uses Petri nets rather than event structures and contains no technicalities like definitions and proofs. Instead more

attention has been paid to the examples.

All chapters in this book can be read independently, although for motivation it may be helpful to read the introduction to Chapter IV before Chapters V-VII, and depending on the taste of the reader it may be fruitful to consult Chapter V before or simultaneously with the last two chapters. Furthermore Chapter VI depends on Section 1 *or* 2 of Chapter IV. Conceptually Chapter VII follows Chapter VI, and it recalls its results.

5. Results. The main new results contributed by this book and its constituent papers are the following:

1. A short non-categorical proof showing that two processes are bisimulation equivalent if and only if they have the same representation in Aczel's universe of non-well-founded sets (Proposition 1.9 in Chapter I). This result follows immediately from the categorical considerations in ACZEL [7], but it is difficult to point out a precise spot in Aczel's text where this result is obtained.
2. Basically all semantics for concrete sequential processes found in the literature that can be defined uniformly in terms of action relations are presented in a uniform framework and motivated in terms of testing scenarios (Chapter I, Section 1).
3. These semantics are partially ordered with respect to their distinguishing power. (Theorems 2.1 and 2.9 for finitely branching processes; Theorems 2.2 and 2.10 for infinitely branching processes.)
4. A complete axiomatization of completed trace equivalence on finite closed process expressions (Theorem 3.1 in Chapter I).
5. Complete axiomatizations of simulation equivalence and ready simulation equivalence on finite process expressions (Theorems 3.1 and 3.2).
6. The introduction of an algebraic specification technique with homomorphism and subalgebra operators (Chapter II, Section 1).
7. A new logic with a new concept of proof for statements involving these specifications (Chapter II, Section 1.8).
8. A concise overview of equational, conditional and first order logic (Appendix to Chapter II).
9. A new formulation of the *Recursive Specification Principle* in ACP (Theorem 2.8.2 in Chapter II) and more in general a concise presentation of the module $ACP^\#$ (Chapter II, Section 2).
10. The inconsistency of the law $T4$ ($\tau(\tau x + y) = \tau x + y$) in combination with ACP (Proposition 3.1.2.2). Equivalences in which this law is valid fail to be congruences for the left-merge operator FOO of ACP.
11. The application of our module logic to combine the left-merge and $T4$ in ACP verifications (Section 3.1 in Chapter II).
12. The application of our specification technique to express an associative chaining operator in $ACP^\#$ (Section 3.2 in Chapter II).
13. The specification of several types of (faulty) queues and bags using the chaining operator, and the verification of their basic properties (Chapter II, Section 4). Thanks to the parametrized nature of our specifications,

- many properties that we verified for the normal queue, can be inferred to hold for the various faulty queues as well.
14. An example of a plausible identity concerning faulty queues that does not hold in weak bisimulation semantics (Theorem 4.5.1). This identity *can* be proved when the axiom T4 above is added to the axioms of weak bisimulation semantics (Theorem 4.5.2).
 15. A verification of the Concurrent Alternating Bit Protocol in no more than 5 pages (Chapter II, Section 5).
 16. The observation that observation equivalence (weak bisimulation) does not respect branching time (Introduction to Chapter III).
 17. The introduction of branching bisimulation equivalence as an alternative (Chapter III, Section 1).
 18. The stuttering lemma (1.1) saying that all intermediate states on a τ -path between two equivalent states are also equivalent with those states. This lemma allows a simpler presentation of a branching bisimulation.
 19. The classification of η -bisimulation and delay bisimulation as two incomparable notions between weak and branching bisimulation (Section 1).
 20. The characterization of rooted weak bisimulation equivalence as the coarsest congruence for $+$ finer than weak bisimulation equivalence, provided that there is at least one action different from τ (Theorem 2.4 in Chapter III). The proof does not use the Fresh Atom Principle and is valid independent of possible cardinality restrictions imposed on the space of processes to which it applies. The same proof also applies to (rooted) branching bisimulation, η -bisimulation and delay bisimulation.
 21. The introduction of coloured traces and consistent colourings, which show how branching bisimulation equivalence preserves the branching structure of processes (Section 3); the characterization of (rooted) branching bisimulation equivalence as (rooted) coloured trace equivalence (Th. 3.2).
 22. The characterization of branching bisimulation equivalence as hypertrace equivalence (Theorem 3.3).
 23. The definition of a unique normal form in every branching bisimulation congruence class, which constitutes a minimal representation of branching congruent processes (Section 3, end).
 24. An explanation of the way in which deadlock and termination are treated in Basic Process Algebra (BPA), the Calculus of Communicating Systems (CCS) and the Algebra of Communicating Processes (ACP) (Section 2).
 25. A complete axiomatization of branching bisimulation congruence on finite closed process expressions. This is done for (essential parts of) BPA, CCS as well as ACP (Sections 2 and 4).
 26. Alternative completeness proofs for the known axiomatizations of weak, η - and delay bisimulation congruence, based on a saturation technique that reduces these congruences to branching congruence (Section 4, end).
 27. A simple proof-rule
$$\frac{\tau x \leftrightarrow_{rb} \tau y}{x \leftrightarrow_b y}$$
 that converts a complete axiomatization of rooted branching bisimulation into one of unrooted branching

bisimulation equivalence (Theorem 2.15). Such a rule was known already for weak bisimulation and also holds for delay and η -bisimulation.

28. The discovery that for a large class of processes weak and branching bisimulation coincide, and that, as far as we know, all protocols that have been verified in weak bisimulation semantics are also valid in branching bisimulation semantics (Section 5).
29. The observation that (rooted) weak bisimulation equivalence is not preserved under refinement of actions, even if only sequential processes without interleaving operators are considered (Section 6).
30. The result that for sequential processes branching bisimulation *is* preserved under action refinement (Theorem 6.1).
31. A method for obtaining a refinement theorem (like the one mentioned above) directly from the complete axiomatization of an equivalence.
32. The definition of branching bisimulation preorders and equivalences taking divergence into account, analogous to the existing preorders and equivalences for weak bisimulation; and similarly for η - and delay bisimulation (Section 7).
33. A sketch of a testing scenario for branching bisimulation equivalence (Section 8, were also an overview of the work on modal characterizations for branching bisimulation can be found). This result was added in this second edition of my book.
34. A list of 11 arguments for choosing branching bisimulation semantics in verifications and other applications of concurrency theory:
 1. Verifications in branching bisimulation semantics are sound independent of your notion of observable behaviour.
 2. No coarser semantics (like weak bisimulation) has this property.
 3. In abstract interleaving semantics no finer notion of bisimulation is suitable.
 4. There is a reasonable operator for which branching bisimulation is a congruence and weak bisimulation or coarser notions are not. On the other hand no examples testifying for the opposite are known.
 5. Branching bisimulation equivalence is the only known equivalence in the linear time - branching time spectrum that supports an 'eventually' operator as part of a temporal logic on transition systems. It even supports all the operators of CTL* and corresponds with stuttering equivalence of kripke structures.
 6. There are practical applications in which weak bisimulation poses a problem that can be solved by moving to branching bisimulation. No applications have been found in which the reverse holds.
 7. Branching bisimulation equivalence has a lower complexity than any other abstract semantic equivalence used in concurrency theory.
 8. Branching bisimulation is preserved under action refinement, whereas weak bisimulation is not.
 9. Branching congruence has a very appealing complete axiomatization
 10. and better term rewriting properties than other (abstract) bisimulations.
 11. And it has a nice characterization as back-and-forth bisimulation.

(Conclusion to Chapter III; arguments 2, 3, 4 and 6b were added in the second edition; arguments 5b, 7, 11 and part of 4 depend on work by others).

35. An illustration of how action refinement can be used in the top-down design of concurrent systems (Introduction to Chapter IV).
36. An argument against the use of forgetful refinements (refining an action by the empty process, implemented by erasure) for this purpose (ibidem).
37. The definition of an operator on prime event structures for refinement of actions by non-empty, finite, conflict-free processes (Section 1). It is established that this operator is well-defined up to isomorphism (Prop. 1.5) and compositional w.r.t. configuration semantics (Proposition 1.7).
38. The observation that prime event structures are less suitable for refinement by conflicting or infinite behaviours (Section 1).
39. The modelling of deadlock behaviour and sequential composition of flow event structures (Section 2).
40. The definition of an operator for general action refinement on flow event structures, generalizing the one we had on prime event structures (Section 2). It is established that this operator is well-defined up to isomorphism (Prop. 2.7) and compositional w.r.t. configuration semantics (Prop. 2.8).
41. The introduction of *configuration structures* as a model for concurrency (Section 3). They are essentially just Winskel's families of configurations of general event structures. The new name merely reflects our opinion that they form a pleasant model of concurrency even without considering event structures.
42. The observation that the infinite configurations in configuration structures (or families of configurations) are redundant, and can better be left out.
43. The definition of an operator for action refinement on configuration structures (Section 3). It is established that this operator is well-defined up to isomorphism (Proposition 3.7) and for deadlock-free refinements agrees with the one on flow event structures (Theorem 3.8).
44. A construction for refinement of transitions in Petri nets, generalizing similar constructions proposed earlier (Section 4). For occurrence nets we establish that the construction is consistent with our operator for action refinement on prime event structures (Theorem 4.18).
45. We show that several prospective generalizations of our approach do not work.
46. The comparison of our notion of transition refinement with the notions of (vicinity respecting) net morphism and quotient (Section 4).
47. An example showing that preservation of semantic equivalence of processes under action refinement can be a useful property in applications (this introduction).
48. Counterexamples showing that, besides the interleaving equivalences, also none of the known step equivalences is preserved under action refinement (Chapter V and Section 2 of Chapter VI).
49. A formal proof that linear time partial order semantics is preserved under action refinement (Section 3 of Chapter VI).

50. Counterexamples showing that pomset bisimulation equivalence, generalized pomset bisimulation equivalence, NMS partial ordering equivalence, and a combination of those, are not preserved under refinement of actions (Chapter V and Section 4 of Chapter VI).
51. The discovery that NMS partial ordering equivalence is incomparable with pomset bisimulation equivalence, and the proposal of a modification of this equivalence, which we call *history preserving bisimulation equivalence*, that is finer than both, and still satisfies the absorption laws (Section 4.2 of Chapter VI). This notion was earlier proposed in a different model as behavior structure equivalence. We argue that it is the coarsest equivalence that models the interplay of causality and branching in full detail (Conclusion to Chapter VI).
52. The result that history preserving bisimulation is preserved under action refinement (Section 4.2 of Chapter VI).
53. The result that for systems without autoconcurrency history preserving bisimulation and NMS partial ordering equivalence coincide (ibidem).
54. An example strongly suggesting that history preserving bisimulation is not the coarsest equivalence refining pomset bisimulation and being preserved under action refinement (Conclusion to Chapter VI).
55. An overview of equivalence notions for concurrent systems (Chapters V, VI and especially VII).
56. The introduction of ST-configurations to model the global state of an event structure, just as ST-markings model the global state of a Petri net. I argue that when actions are allowed to have duration, they do so better than ordinary markings or configurations (Chapter VII, Section 4).
57. Several characterizations of ST-trace and ST-bisimulation equivalence on event structures (Section 5). ST-bisimulation had already been defined on Petri nets, and ST-trace equivalence is obtained as its linear time variant.
58. The results that ST-trace and ST-bisimulation equivalence are preserved under refinement of actions (Section 6). It follows that partial order semantics is not needed for dealing with action refinement (Chapters V and VII).
59. An example showing that ST-bisimulation is not preserved under forgetful refinement (Conclusion to Chapter VII).
60. The conjecture that ST-bisimulation is the coarsest equivalence contained in bisimulation equivalence, and similarly ST-trace the coarsest in trace equivalence, that is preserved under action refinement. Also the main idea for the proof is provided (ibidem).

6. *Introduction to the second edition.* This second edition contains the same material as the first edition. In addition a large amount of follow-up work is discussed and referenced.

Chapter I

The Linear Time - Branching Time Spectrum

Rob van Glabbeek

In this chapter various semantics in the linear time - branching time spectrum are presented in a uniform, model-independent way. Restricted to the domain of finitely branching, concrete, sequential processes, only twelve of them turn out to be different, and most semantics found in the literature that can be defined uniformly in terms of action relations coincide with one of these twelve. Several testing scenarios, motivating these semantics, are presented, phrased in terms of 'button pushing experiments' on generative and reactive machines. Finally ten of these semantics are applied to a simple language for finite, concrete, sequential, nondeterministic processes, and for each of them a complete axiomatization is provided.

TABLE OF CONTENTS

Introduction	17
1. Semantic equivalences on labelled transition systems	22
2. The semantic lattice	38
3. Complete axiomatizations	50
Concluding remarks	52

INTRODUCTION

Process theory. A *process* is the behaviour of a system. The system can be a machine, an elementary particle, a communication protocol, a network of falling dominoes, a chess player, or any other system. Process theory is the study of processes. Two main activities of process theory are *modelling* and *verification*. Modelling is the activity of representing processes, mostly as elements of a mathematical domain or as expressions in a system description language. Verification is the activity of proving statements about processes, for instance that the actual behaviour of a system is equal to its intended behaviour. Of course, this is only possible if a criterion has been defined, determining whether or not two processes are equal, i.e. two systems behave similarly. Such a criterion constitutes the *semantics* of a process theory. (To be precise, it constitutes the semantics of the equality concept employed in a process theory.) Which aspects of the behaviour of a system are of importance

to a certain user depends on the environment in which the system will be running, and on the interests of the particular user. Therefore it is not a task of process theory to find the 'true' semantics of processes, but rather to determine which process semantics is suitable for which applications.

Comparative concurrency semantics. This book aims at the classification of process semantics.¹ The set of possible process semantics can be partially ordered by the relation 'makes strictly more identifications on processes than', thereby becoming a complete lattice². Now the classification of some useful process semantics can be facilitated by drawing parts of this lattice and locating the positions of some interesting process semantics, found in the literature. Furthermore the ideas involved in the construction of these semantics can be unraveled and combined in new compositions, thereby creating an abundance of new process semantics. These semantics will, by their intermediate positions in the semantic lattice, shed light on the differences and similarities of the established ones. Sometimes they also turn out to be interesting in their own right. Finally the semantic lattice serves as a map on which it can be indicated which semantics satisfy certain desirable properties, and are suited for a particular class of applications.

Most semantic notions encountered in contemporary process theory can be classified along four different lines, corresponding with four different kinds of identifications. First there is the dichotomy of linear time versus branching time: to what extent should one identify processes differing only in the branching structure of their execution paths? Secondly there is the dichotomy of interleaving semantics versus partial order semantics: to what extent should one identify processes differing only in the causal dependencies between their actions (while agreeing on the possible orders of execution)? Thirdly one encounters different treatments of abstraction from internal actions in a process: to what extent should one identify processes differing only in their internal or silent actions? And fourthly there are different approaches to infinity: to what extent should one identify processes differing only in their infinite behaviour? These considerations give rise to a four dimensional representation of the proposed semantic lattice.

However, at least three more dimensions can be distinguished. In this book, stochastic and real-time aspects of processes are completely neglected. Furthermore it deals with *uniform concurrency*³ only. This means that processes are studied, performing actions⁴ $a, b, c, \dots$ which are not subject to further investigations. So it remains unspecified if these actions are in fact assignments to variables or the falling of dominoes or other actions. If also the options are considered of modelling (to a certain degree) the stochastic and real-time aspects of processes and the operational behaviour of the elementary actions, three more parameters in the classification emerge.

1. This field of research is called *comparative concurrency semantics*, a terminology first used by MEYER in [132].

2. The supremum of a set of process semantics is the semantics identifying two processes whenever

Process domains. In order to be able to reason about processes in a mathematical way, it is common practice to represent processes as elements of a mathematical domain. Such a domain is called a *process domain*. The relation between the domain and the world of real processes is mostly stated informally. The semantics of a process theory can be modelled as an equivalence on a process domain, called a *semantic equivalence*. In the literature one finds among others:

- *graph domains*, in which a process is represented as a *process graph*, or *state transition diagram*,
- *net domains*, in which a process is represented as a (labelled) *Petri net*,
- *event structure domains*, in which a process is represented as a (labelled) *event structure*,
- *explicit domains*, in which a process is represented as a mathematically coded set of its properties,
- *projective limit domains*, which are obtained as projective limits of series of finite term domains,
- and *term domains*, in which a process is represented as a term in a system description language.

Action relations. Write $p \xrightarrow{a} q$ if the process p can evolve into the process q , while performing the action a . The binary predicates $\xrightarrow{a}$ are called *action relations*. The semantic equivalences which are treated in this chapter will be defined entirely in terms of action relations. Hence these definitions apply to any process domain on which action relations are defined. Furthermore they will be defined *uniformly* in terms of action relations, meaning that all actions are treated in the same way. For reasons of convenience, even the usual distinction between internal and external actions is dropped in this chapter.

Finitely branching, concrete, sequential processes. Being a first step, this chapter limits itself to a very simple class of processes. First of all only *sequential* processes are investigated: processes capable of performing at most one action at a time. Moreover the main interest is in *finitely branching* processes: processes having in each state only finitely many possible ways to proceed. Finally, instead of dropping the usual distinction between internal and external actions, one can equivalently maintain to study *concrete* processes in which no internal actions occur (and also no internal choices as in CSP [116]). For this simple class of processes, when considering only semantic equivalences that can be defined uniformly in terms of action relations, the announced semantic lattice collapses in six out of seven dimensions and covers only the *linear time - branching time* spectrum.

Literature. In the literature on uniform concurrency 11 semantics can be found, which are uniformly definable in terms of action relations and different on the domain of finitely branching, sequential processes (see Figure 1). The coarsest one (i.e. the semantics making the most identifications) is *trace semantics*, as presented in HOARE [115]. In trace semantics only *partial traces* are

er they are identified by every semantics in this set.

3. The term uniform concurrency is employed by DE BAKKER et al [19].




FIGURE 1. The linear time - branching time spectrum

employed. The finest one (making less identifications than any of the others) is *bisimulation semantics*, as presented in MILNER [136]. Bisimulation semantics is the standard semantics for the system description language CCS (MILNER [134]). The notion of bisimulation was introduced in PARK [145]. Bisimulation equivalence is a refinement of *observational equivalence*, as

4. Strictly speaking processes do not perform actions, but systems do. However, for reasons of convenience, this book sometimes uses the word process, when actually referring to a system of

introduced by HENNESSY & MILNER in [111]. On the domain of finitely branching, concrete, sequential processes, both equivalences coincide. Also the semantics of DE BAKKER & ZUCKER, presented in [20], coincides with bisimulation semantics on this domain. Then there are nine semantics in between. First of all a variant of trace semantics can be obtained by using *complete traces* besides (or instead of) partial ones. In this chapter it is called *completed trace semantics*. *Failure semantics* is introduced in BROOKES, HOARE & ROSCOE [42], and used in the construction of a model for the system description language CSP (HOARE [114, 116]). It is finer than completed trace semantics. The semantics based on *testing equivalences*, as developed in DE NICOLA & HENNESSY [64], coincides with failure semantics on the domain of finitely branching, concrete, sequential processes, as do the semantics of KENNAWAY [121] and DARONDEAU [52]. This has been established in DE NICOLA [63]. In OLDEROG & HOARE [144] *readiness semantics* is presented, which is slightly finer than failure semantics. Between readiness and bisimulation semantics one finds *ready trace semantics*, as introduced independently in PNUELI [152] (there called *barbed semantics*), BAETEN, BERGSTRA & KLOP [14] and POMELLO [153] (under the name *exhibited behaviour semantics*). The natural completion of the square, suggested by failure, readiness and ready trace semantics yields *failure trace semantics*. For finitely branching processes this is the same as *refusal semantics*, introduced in PHILLIPS [150]. *Simulation equivalence*, based on the classical notion of *simulation* (see e.g. PARK [145]), is independent of the last five semantics. *Ready simulation semantics* was introduced in BLOOM, ISTRAIL & MEYER [35] under the name *GSOS trace congruence*. It is finer than ready trace as well as simulation equivalence. In LARSEN & SKOU [128] a more operational characterization of this equivalence was given under the name $\frac{2}{3}$ -*bisimulation equivalence*. This characterization resembles the one used in this chapter. Finally *2-nested simulation equivalence*, introduced in GROOTE & VAANDRAGER [107], is located between ready simulation and bisimulation equivalence, and *possible-futures semantics*, as proposed in ROUNDS & BROOKES [159], can be positioned between 2-nested simulation and readiness semantics. Among the semantics which are not definable in terms of action relations and thus fall outside the scope of this chapter, one finds semantics that take stochastic properties of processes into account, as in VAN GLABBEEK, SMOLKA, STEFFEN & TOFTS [90] and semantics that make almost no identifications and are hardly used for system verification.

About the contents. In the first section of this chapter all semantics are defined, and motivated by several testing scenarios, which are phrased in terms of button pushing experiments. In Section 2 the semantics are partially ordered by the relation ‘makes at least as many identifications as’. This yields the infinitary linear time - branching time spectrum. Counterexamples are provided, showing that on a graph domain this ordering cannot be further expanded. However, for deterministic processes the spectrum collapses, as was first observed by PARK [145]. Finally, in Section 3, ten of these semantics are applied to a simple language for finite, concrete, sequential, nondeterministic

which the process is the behaviour.

processes, and for each of them a complete axiomatization is provided.

1. SEMANTIC EQUIVALENCES ON LABELLED TRANSITION SYSTEMS

1.1. Labelled transition systems. In this book processes will be investigated, that are capable of performing actions from a given set Act . By an action any activity is understood that is considered as a conceptual entity on a chosen level of abstraction. Actions may be instantaneous or durational and are not required to terminate, but in a finite time only finitely many actions can be carried out. Any activity of an investigated process should be part of some action $a \in Act$ performed by the process. Different activities that are indistinguishable on the chosen level of abstraction are interpreted as occurrences of the same action $a \in Act$.

A process is *sequential* if it can perform at most one action at the same time. In this chapter only sequential processes will be considered. A domain of sequential processes can often be conveniently represented as a labelled transition system. This is a domain $\mathbb{A}$ on which infix written binary predicates $\xrightarrow{a}$ are defined for each action $a \in Act$. The elements of $\mathbb{A}$ represent processes, and $p \xrightarrow{a} q$ means that p can start performing the action a and after completion of this action reach a state where q is its remaining behaviour.

In a labelled transition system it may happen that $p \xrightarrow{a} q$ and $p \xrightarrow{b} r$ for different actions a and b or different processes p and q . This phenomena is called *branching*. It need not be specified how the choice between the alternatives is made, or whether a probability distribution can be attached to it.

NOTATION: For any alphabet Σ , let Σ^* be the set of *strings* over Σ . Write ϵ for the empty string, $\sigma\rho$ for the concatenation of σ and $\rho \in \Sigma^*$, and a for the string, consisting of the single symbol $a \in \Sigma$.

DEFINITION: A *labelled transition system* is a pair $(\mathbb{A}, \rightarrow)$ with $\mathbb{A}$ a class and $\rightarrow \subseteq \mathbb{A} \times Act \times \mathbb{A}$, such that for $p \in \mathbb{A}$ and $a \in Act$ the class $\{q \in \mathbb{A} \mid p \xrightarrow{a} q\}$ is a set.

Let for the remainder of this section $(\mathbb{A}, \rightarrow)$ be a labelled transition system, ranged over by $p, q, r, \dots$. Write $p \xrightarrow{a} q$ for $(p, a, q) \in \rightarrow$. The binary predicates $\xrightarrow{a}$ are called *action relations*.

DEFINITIONS (Remark that the following concepts are defined in terms of action relations only):

- The *generalized action relations* $\xrightarrow{\sigma}$ for $\sigma \in Act^*$ are defined inductively by:

1. $p \xrightarrow{\epsilon} p$, for any process p .
2. $(p, a, q) \in \rightarrow$ with $a \in \text{Act}$ implies $p \xrightarrow{a} q$ with $a \in \text{Act}^*$.
3. $p \xrightarrow{\sigma} q \xrightarrow{\rho} r$ implies $p \xrightarrow{\sigma\rho} r$.

In words: the generalized action relations $\xrightarrow{\sigma}$ are the reflexive and transitive closure of the ordinary action relations $\xrightarrow{a}$. $p \xrightarrow{\sigma} q$ means that p can evolve into q , while performing the sequence σ of actions. Remark that the overloading of the notion $p \xrightarrow{a} q$ is quite harmless.

- The set of *initial actions* of a process p is defined by:
 $I(p) = \{a \in \text{Act} \mid \exists q : p \xrightarrow{a} q\}$.
- A process $p \in \mathbb{A}$ is *finitely branching* if for each $q \in \mathbb{A}$ with $p \xrightarrow{\sigma} q$ for some $\sigma \in \text{Act}^*$, the set $\{(a, r) \mid q \xrightarrow{a} r, a \in \text{Act}, r \in \mathbb{A}\}$ is finite.


In the following, several semantic equivalences on $\mathbb{A}$ will be defined in terms of action relations. Most of these equivalences can be motivated by the observable behaviour of processes, according to some testing scenario. (Two processes are equivalent if they allow the same set of possible observations, possibly in response on certain experiments.) I will try to capture these motivations in terms of *button pushing experiments* (cf. MILNER [134], pp. 10-12).

1.2. Trace semantics. $\sigma \in \text{Act}^*$ is a *trace* of a process p , if there is a process q , such that $p \xrightarrow{\sigma} q$. Let $T(p)$ denote the set of traces of p . Two processes p and q are *trace equivalent* if $T(p) = T(q)$. In trace semantics two processes are identified iff they are trace equivalent.

Trace semantics is based on the idea that two processes are to be identified if they allow the same set of observations, where an observation simply consists of a sequence of actions performed by the process in succession.

1.3. Completed trace semantics. $\sigma \in \text{Act}^*$ is a *complete trace* of a process p , if there is a process q , such that $p \xrightarrow{\sigma} q$ and $I(q) = \emptyset$. Let $CT(p)$ denote the set of complete traces of p . Two processes p and q are *completed trace equivalent* if $T(p) = T(q)$ and $CT(p) = CT(q)$. In completed trace semantics two processes are identified iff they are completed trace equivalent.


Completed trace semantics can be explained with the following (rather trivial) *completed trace machine*. The process is modelled as a black box that contains as its interface to the outside world a display on which the name of the action is shown that is currently carried out by the process. The process autonomously choses an execution path that is consistent with its position in the labelled transition system $(\mathbb{A}, \rightarrow)$. During this execution always an action name is visible on the display. As soon as no further action can be carried out, the process reaches a state of deadlock and the display becomes empty. Now the existence of an observer is assumed that watches the display and

FIGURE 2. *The completed trace machine*

records the sequence of actions displayed during a run of the process, possibly followed by deadlock. It is assumed that an observation takes only a finite amount of time and may be terminated before the process stagnates. Two processes are identified if they allow the same set of observations in this sense.

The *trace machine* can be regarded as a simpler version of the completed trace machine, were the last action name remains visible in the display if deadlock occurs (unless deadlock occurs in the beginning already). On this machine traces can be recorded, but stagnation can not be detected, since in case of deadlock the observer may think that the last action is still continuing.

1.4. Failure semantics. The *failure machine* contains as its interface to the outside world not only the display of the completed trace machine, but also a switch for each action $a \in Act$ (as in Figure 3).

FIGURE 3. *The failure trace machine*

By means of these switches the observer may determine which actions are *free* and which are *blocked*. This situation may be changed any time during a run of the process. As before, the process autonomously choses an execution path that fits with its position in $(\mathbf{A}, \rightarrow)$, but this time the process may only start the execution of free actions. If the process reaches a state where all initial actions of its remaining behaviour are blocked, it can not proceed and the

machine stagnates, which can be recognized from the empty display. In this case the observer may record that after a certain sequence of actions σ , the set X of free actions is refused by the process. X is therefore called a *refusal set* and $\langle \sigma, X \rangle$ a *failure pair*. The set of all failure pairs of a process is called its *failure set*, and constitutes its observable behaviour.

DEFINITION: $\langle \sigma, X \rangle \in Act^* \times \mathcal{P}(Act)$ is a *failure pair* of a process p , if there is a process q , such that $p \xrightarrow{\sigma} q$ and $I(q) \cap X = \emptyset$. Let $F(p)$ denote the set of failure pairs of p . Two processes p and q are *failure equivalent* if $F(p) = F(q)$. In failure semantics two processes are identified iff they are failure equivalent.

This version of failure semantics is taken from HOARE [116]. In BROOKES, HOARE & ROSCOE [42], where failure semantics was introduced, the refusal sets are required to be finite. It is not difficult to see that for finitely branching processes the two versions yield the same failure equivalence. In fact this follows immediately from the following proposition, that says that, for finitely branching processes, the failure pairs with infinite refusal set are completely determined by the ones with finite refusal set.

PROPOSITION 1.1: Let $p \in \mathbb{A}$ and $\sigma \in T(p)$. Put $Cont(\sigma) = \{a \in Act \mid \sigma a \in T(p)\}$.

- i. Then, for $X \subseteq Act$, $\langle \sigma, X \rangle \in F(p) \Leftrightarrow \langle \sigma, X \cap Cont(\sigma) \rangle \in F(p)$.
- ii. If p is finitely branching then $Cont(\sigma)$ is finite.

PROOF: Straightforward. □

In DE NICOLA [63] several equivalences, that were proposed in KENNAWAY [121], DARONDEAU [52] and DE NICOLA & HENNESSY [64], are shown to coincide with failure semantics on the domain of finitely branching transition systems without internal moves. For this purpose he uses the following alternative characterization of failure equivalence.

DEFINITION: Write p after σ *MUST* X if for each $q \in \mathbb{A}$ with $p \xrightarrow{\sigma} q$ there is an $r \in \mathbb{A}$ and $a \in X$ such that $q \xrightarrow{a} r$. Put $p \simeq q$ if for all $\sigma \in Act^*$ and $X \subseteq Act$: p after σ *MUST* $X \Leftrightarrow q$ after σ *MUST* X .

PROPOSITION 1.2: Let $p, q \in \mathbb{A}$. Then $p \simeq q \Leftrightarrow F(p) = F(q)$.

PROOF: p after σ *MUST* $X \Leftrightarrow (\sigma, X) \notin F(p)$ [63]. □

In HENNESSY [108], a model for nondeterministic behaviours is proposed in which a process is represented as an *acceptance tree*. An acceptance tree of a finitely branching process p without internal moves or internal nondeterminism can be represented as the set of all pairs $\langle \sigma, X \rangle \in Act^* \times \mathcal{P}(Act)$ for which there is a process q , such that $p \xrightarrow{\sigma} q$ and $X \subseteq I(q)$. It follows that for such processes *acceptance tree equivalence* coincides with failure equivalence.

1.5. Failure trace semantics. The *failure trace machine* has the same layout as the failure machine, but it does not stagnate permanently if the process cannot proceed due to the circumstance that all actions it is prepared to continue with are blocked by the observer. Instead it idles - recognizable from the empty display - until the observer changes its mind and allows one of the actions the process is ready to perform. What can be observed are traces with idle periods in between, and for each such period the set of actions that are not blocked by the observer. Such observations can be coded as sequences of members and subsets of Act .

EXAMPLE: The sequence $\{a,b\}cdb\{b,c\}\{b,c,d\}a(Act)$ is the account of the following observation: At the beginning of the execution of the process p , only the actions a and b were allowed by the observer. Apparently, these actions were not on the menu of p , for p started with an idle period. Suddenly the observer canceled its veto on c , and this resulted in the execution of c , followed by d and b . Then again an idle period occurred, this time when b and c were the actions not being blocked by the observer. After a while the observer decided to allow d as well, but the process ignored this gesture and remained idle. Only when the observer gave the green light for the action a , it happened immediately. Finally, the process became idle once more, but this time not even one action was blocked. This made the observer realize that a state of eternal stagnation had been reached, and disappointed he terminated the observation.

A set $X \subseteq Act$, occurring in such a sequence, can be regarded as an offer from the environment, that is refused by the process. Therefore such a set is called a *refusal set*. The occurrence of a refusal set may be interpreted as a 'failure' of the environment to create a situation in which the process can proceed without being disturbed. Hence a sequence over $Act \cup \mathcal{P}(Act)$, resulting from an observation of a process p may be called a *failure trace* of p . The observable behaviour of a process, according to this testing scenario, is given by the set of its failure traces, its *failure trace set*. The semantics in which processes are identified iff their failure trace sets coincide, is called *failure trace semantics*.

DEFINITIONS:

- The *refusal relations* $\xrightarrow{X}$ for $X \subseteq Act$ are defined by: $p \xrightarrow{X} q$ iff $p = q$ and $I(p) \cap X = \emptyset$.
 $p \xrightarrow{X} q$ means that p can evolve into q , while being idle during a period in which X is the set of actions allowed by the environment.
- The *failure trace relations* $\xrightarrow{\sigma}$ for $\sigma \in (Act \cup \mathcal{P}(Act))^*$ are defined as the reflexive and transitive closure of both the action and the refusal relations. Again the overloading of notation is harmless.
- $\sigma \in (Act \cup \mathcal{P}(Act))^*$ is a *failure trace* of a process p , if there is a process q , such that $p \xrightarrow{\sigma} q$. Let $FT(p)$ denote the set of failure traces of p . Two processes p and q are *failure trace equivalent* if $FT(p) = FT(q)$.

EXERCISES:

1. Explain why $a\{a,b\}a$ can never be a failure trace of a process $p \in \mathbb{A}$.
2. Can $\{a\}b$ and $\{b\}a$ be two failure traces of such a process? And $a\{a\}b$ and $a\{b\}a$?
3. $\{a,b\}cc$, $\{a\}c\{b\}c$, $\{b\}c\{a\}c$, $c\{a,b\}c$, $c\{a\}\{b\}c$ and c are failure traces of a process $p \in \mathbb{A}$. Which selections from this series provide the same information about p ?

1.6. *Ready trace semantics.* The *Ready trace machine* is a variant of the failure trace machine that is equipped with a lamp for each action $a \in \text{Act}$.




FIGURE 4. *The ready trace machine*

Each time the process idles, the lamps of all actions the process is ready to engage in are lit. Of course all these actions are blocked by the observer, otherwise the process wouldn't idle. Now the observer can see which actions could be released in order to let the process proceed. During the execution of an action no lamps are lit. An observation now consists of a sequence of members and subsets of Act , the actions representing information obtained from the display, and the sets of actions representing information obtained from the lights. Such a sequence is called a *ready trace* of the process, and the subsets occurring in a ready trace are referred to as *menus*. The information about the free and blocked actions is now redundant. The set of all ready traces of a process is called its *ready trace set*, and constitutes its observable behaviour.


DEFINITIONS:

- The *ready trace relations* $\xrightarrow{\sigma}$ for $\sigma \in (\text{Act} \cup \mathcal{P}(\text{Act}))^*$ are defined by:
 1. $p \xrightarrow{\epsilon} p$, for any process p .
 2. $p \xrightarrow{a} q$ implies $p \xrightarrow{a} q$.
 3. $p \xrightarrow{X} q$ with $X \subseteq \text{Act}$ whenever $p = q$ and $I(p) = X$.
 4. $p \xrightarrow{\sigma} q \xrightarrow{\rho} r$ implies $p \xrightarrow{\sigma\rho} r$.

The special arrow $\overset{\sigma}{\ast\rightarrow}$ had to be used, since further overloading of $\overset{\sigma}{\rightarrow}$ would cause confusion with the failure trace relations.

- $\sigma \in (Act \cup \mathcal{P}(Act))^*$ is a *ready trace* of a process p , if there is a process q , such that $p \overset{\sigma}{\ast\rightarrow} q$. Let $RT(p)$ denote the set of ready traces of p . Two processes p and q are *ready trace equivalent* if $RT(p) = RT(q)$. In ready trace semantics two processes are identified iff they are ready trace equivalent.

In BAETEN, BERGSTRA & KLOP [14], PNUELI [152] and POMELLO [153] ready trace semantics was defined slightly differently. By the proposition below, their definition yields the same equivalence as mine.



DEFINITION: $X_0 a_1 X_1 a_2 \cdots a_n X_n \in \mathcal{P}(Act) \times (Act \times \mathcal{P}(Act))^*$ is a *normal ready trace* of a process p , if there are processes $p_1, \dots, p_n$ such that $p \xrightarrow{a_1} p_1 \xrightarrow{a_2} \cdots \xrightarrow{a_n} p_n$ and $I(p_i) = X_i$ for $i = 1, \dots, n$. Let $RT_N(p)$ denote the set of normal ready traces of p . Two processes p and q are ready trace equivalent in the sense of [14, 152, 153] if $RT_N(p) = RT_N(q)$.

PROPOSITION 1.3: Let $p, q \in \mathbf{A}$. Then $RT_N(p) = RT_N(q) \Leftrightarrow RT(p) = RT(q)$.

PROOF: The normal ready traces of a process are just the ready traces which are an alternating sequence of sets and actions, and vice versa the set of all ready traces can be constructed from the set of normal ready traces by means of doubling and leaving out menus. $\square$

1.7. Readiness semantics. The *readiness machine* has the same layout as the ready trace machine, but, like the failure machine, can not recover from an idle period. By means of the lights the menu of initial actions of the remaining behaviour of an idle process can be recorded, but this happens at most once during an observation of a process, namely at the end. An observation either results in a trace of the process, or in a pair of a trace and a menu of actions by which the observation could have been extended if the observer wouldn't have blocked them. Such a pair is called a *ready pair* of the process, and the set of all ready pairs of a process is its *ready set*.

DEFINITION: $\langle \sigma, X \rangle \in Act^* \times \mathcal{P}(Act)$ is a *ready pair* of a process p , if there is a process q , such that $p \xrightarrow{\sigma} q$ and $I(q) = X$. Let $R(p)$ denote the set of ready pairs of p . Two processes p and q are *ready equivalent* if $R(p) = R(q)$. In readiness semantics two processes are identified iff they are ready equivalent.

Two preliminary versions of readiness semantics were proposed in ROUNDS &

BROOKES [159]. In *possible-futures semantics* the menu consists of the entire trace set of remaining behaviour of an idle process, instead of only the set of its initial actions; in *acceptance-refusal semantics* a menu may be any finite subset of initial actions, while also the finite refusal sets of Subsection 1.4 are observable.

DEFINITION: $\langle \sigma, X \rangle \in Act^* \times \mathcal{P}(Act^*)$ is a *possible-future* of a process p , if there is a process q , such that $p \xrightarrow{\sigma} q$ and $T(q) = X$. Let $PF(p)$ denote the set of possible futures of p . Two processes p and q are *possible-futures equivalent* if $PF(p) = PF(q)$.

DEFINITION: $\langle \sigma, X, Y \rangle \in Act^* \times \mathcal{P}(Act) \times \mathcal{P}(Act)$ is a *acceptance-refusal triple* of a process p , if X and Y are finite and there is a process q , such that $p \xrightarrow{\sigma} q$, $X \subseteq I(q)$ and $Y \cap I(q) = \emptyset$. Let $AR(p)$ denote the set of acceptance-refusal triples of p . Two processes p and q are *acceptance-refusal equivalent* if $AR(p) = AR(q)$.

It is not difficult to see that for finitely branching processes acceptance-refusal equivalence coincides with readiness equivalence: $\langle \sigma, X \rangle$ is a ready pair of a process p iff p has an acceptance-refusal triple $\langle \sigma, X, Y \rangle$ with $X \cup Y = Cont(\sigma)$ (as defined in the proof of Proposition 1.1).

1.8. *Infinite observations.* All testing scenarios up till now assumed that an observation takes only a finite amount of time. However, they can be easily adapted in order to take infinite behaviours into account.

DEFINITION:

- For any alphabet Σ , let Σ^ω be the set of infinite sequences over Σ .
- $a_1 a_2 \dots \in Act^\omega$ is an *infinite trace* of a process $p \in \mathbb{A}$, if there are processes $p_1, p_2, \dots$ such that $p \xrightarrow{a_1} p_1 \xrightarrow{a_2} \dots$. Let $T^\omega(p)$ denote the set of infinite traces of p .
- Two processes p and q are *infinitary trace equivalent* if $T(p) = T(q)$ and $T^\omega(p) = T^\omega(q)$.
- p and q are *infinitary completed trace equivalent* if $CT(p) = CT(q)$ and $T^\omega(p) = T^\omega(q)$. Note that in this case also $T(p) = T(q)$.
- p and q are *infinitary failure equivalent* if $F(p) = F(q)$ and $T^\omega(p) = T^\omega(q)$.
- p and q are *infinitary ready equivalent* if $R(p) = R(q)$ and $T^\omega(p) = T^\omega(q)$.
- Infinitary failure traces and infinitary ready traces $\sigma \in (Act \cup \mathcal{P}(Act))^\omega$ and the corresponding sets $FT^\omega(p)$ and $RT^\omega(p)$ are defined in the obvious way. Two processes p and q are *infinitary failure trace equivalent* if $FT^\omega(p) = FT^\omega(q)$, and likewise for infinitary ready trace equivalence.

With Königs lemma one easily proves that for finitely branching processes all infinitary equivalences coincide with the corresponding finitary ones.

1.9. Simulation semantics. The testing scenario for finitary simulation semantics resembles that for trace semantics, but in addition the observer is, at any time during a run of the investigated process, capable of making arbitrary (but finitely) many copies of the process in its present state and observe them independently. Thus an observation yields a tree rather than a sequence of actions. Such a tree can be coded as an expression in a simple modal language.

DEFINITIONS:

- The set $\mathcal{L}_S$ of *simulation formulas* over Act is defined inductively by:
 1. $T \in \mathcal{L}_S$.
 2. If $\phi, \psi \in \mathcal{L}_S$ then $\phi \wedge \psi \in \mathcal{L}_S$.
 3. If $\phi \in \mathcal{L}_S$ and $a \in Act$ then $a\phi \in \mathcal{L}_S$.
- The *satisfaction relation* $\models \subseteq \mathbb{A} \times \mathcal{L}_S$ is defined inductively by:
 1. $p \models T$ for all $p \in \mathbb{A}$.
 2. $p \models \phi \wedge \psi$ if $p \models \phi$ and $p \models \psi$.
 3. $p \models a\phi$ if for some $q \in \mathbb{A}$: $p \xrightarrow{a} q$ and $q \models \phi$.
- Let $S^*(p)$ denote the set of all simulation formula that are satisfied by the process p :
 $S^*(p) = \{\phi \in \mathcal{L}_S \mid p \models \phi\}$. Two processes p and q are *finitary simulation equivalent* if $S^*(p) = S^*(q)$.

The following concept of *simulation*, occurs frequently in the literature (see e.g. PARK [145]). The derived notion of *simulation equivalence* coincides with finitary simulation equivalence for finitely branching processes.

DEFINITION: A *simulation* is a binary relation R on processes, satisfying, for $a \in Act$:

- if pRq and $p \xrightarrow{a} p'$, then $\exists q': q \xrightarrow{a} q'$ and $p'Rq'$.

Process p can be simulated by q , notation $s \subseteq t$, if there is a simulation R with pRq .

p and q are *similar*, notation $p \Leftarrow q$, if $p \subseteq q$ and $q \subseteq p$.

PROPOSITION 1.4: Similarity is an equivalence on the domain of processes.

PROOF: It has to be checked that $p \subseteq p$, and $p \subseteq q \ \& \ q \subseteq r \Rightarrow p \subseteq r$.

- The identity relation is a simulation with pRp .
- If R is a simulation with pRq and S is a simulation with qSr , then the relation $R \circ S$, defined by $x(R \circ S)z$ iff $\exists y: xRy \ \& \ ySz$, is a simulation with $p(R \circ S)r$. □

Hence the relation will be called *simulation equivalence*.

PROPOSITION 1.5: Let $p, q \in \mathbb{A}$ be finitely branching processes. Then $p \Leftarrow q \Leftrightarrow S^*(p) = S^*(q)$.

PROOF: See HENNESSY & MILNER [112]. □

The testing scenario for simulation semantics differs from that for finitary simulation semantics, in that both the duration of observations and the amount of copies that can be made each time are not required to be finite.

1.10. Ready simulation semantics. Of course one can also combine the copying facility with any of the other testing scenarios. The observer can then plan experiments on one of the generative machines from the Subsections 1.3 to 1.7 together with a *duplicator*, an ingenious device by which one can duplicate the machine whenever and as often as one wants. In order to represent observations, the modal language from the previous subsection needs to be slightly extended.

DEFINITIONS:

- The *completed simulation formulas* and the corresponding satisfaction relation are defined by means of the extra clauses:
 - 4. $0 \in \mathcal{L}_{CS}$.
 - 4. $p \models 0$ if $I(p) = \emptyset$.
- For the *failure simulation formulas* one needs:
 - 4. If $X \subseteq Act$ then $X \in \mathcal{L}_{FS}$.
 - 4. $p \models X$ if $I(p) \cap X = \emptyset$.
- For the *ready simulation formulas*:
 - 4. If $X \subseteq Act$ then $X \in \mathcal{L}_{RS}$.
 - 4. $p \models X$ if $I(p) = X$.
- For the *failure trace simulation formulas*:
 - 4. If $\phi \in \mathcal{L}_{FTS}$ and $X \subseteq Act$ then $X\phi \in \mathcal{L}_{FTS}$.
 - 4. $p \models X\phi$ if $I(p) \cap X = \emptyset$ and $p \models \phi$.
- And for the *ready trace simulation formulas*:
 - 4. If $\phi \in \mathcal{L}_{RTS}$ and $X \subseteq Act$ then $X\phi \in \mathcal{L}_{RTS}$.
 - 4. $p \models X\phi$ if $I(p) = X$ and $p \models \phi$.

Note that traces, complete traces, failure pairs, etc. can be obtained as the corresponding kind of simulation formulas without the operator $\wedge$.

By means of the formulas defined above one can define the finitary versions of *completed simulation equivalence*, *ready simulation equivalence*, etc. It is obvious that failure trace simulation equivalence coincides with failure simulation equivalence and ready trace simulation equivalence with ready simulation equivalence ($p \models X\phi \Leftrightarrow p \models X \wedge \phi$). Also it is not difficult to see that for finitely branching processes¹ failure simulation equivalence and ready simulation equivalence coincide. So for finitely branching processes two different equivalences remain, the finitary versions of which coincide with the following infinitary versions.

1. My thanks to PH. SCHNOEBELEN [163] for pointing out that this cannot be said for infinitely branching processes, as mistakenly done in the original version of this book.

DEFINITION: A *complete simulation* is a binary relation R on processes, satisfying, for $a \in Act$:

- if pRq and $p \xrightarrow{a} p'$, then $\exists q': q \xrightarrow{a} q'$ and $p'Rq'$;
- if pRq then $I(p) = \emptyset \Leftrightarrow I(q) = \emptyset$.

Two processes p and q are *completed simulation equivalent* if there exists a complete simulation R with pRq and a complete simulation S with qSp .

DEFINITION: A *ready simulation* is a binary relation R on processes, satisfying, for $a \in Act$:

- if pRq and $p \xrightarrow{a} p'$, then $\exists q': q \xrightarrow{a} q'$ and $p'Rq'$;
- if pRq then $I(p) = I(q)$.

Two processes p and q are *ready simulation equivalent* if there exists a ready simulation R with pRq and a ready simulation S with qSp .

Failure simulation equivalence can be defined analogously, and, unlike its finitary version, is easily seen to coincide with ready simulation equivalence.

An alternative and maybe more natural testing scenario for finitary ready simulation semantics (or simulation semantics) can be obtained by exchanging the duplicator for an *undo*-button on the (ready) trace machine (Figure 5).




FIGURE 5. *The ready simulation machine*

It is assumed that all intermediate states that are past through during a run of a process are stored in a memory inside the black box. Now pressing the *undo*-button causes the machine to shift one state backwards. In case the button is pressed during the execution of an action, this execution will be interrupted and the process assumes the state just before this action began. In the initial state pressing the button has no effect. An observation now consists of a (ready) trace, enriched with *undo*-actions. Such observations can easily be translated in (ready) simulation formulas.

1.11. Refusal (simulation) semantics. In the testing scenarios presented so far, a process is considered to perform actions and make choices autonomously. The investigated behaviours can therefore be classified as *generative processes*. The observer merely restricts the spontaneous behaviour of the generative machine by cutting off some possible courses of action. An alternative view of the investigated processes can be obtained by considering them to react on stimuli from the environment and be passive otherwise. *Reactive machines* can be obtained out of the generative machines presented so far by replacing the switches by buttons and the display by a green light.




FIGURE 6. *The reactive ready simulation machine*

Initially the process waits patiently until the observer tries to press one of the buttons. If the observer tries to press an *a*-button, the machine can react in two different ways: if the process can not start with an *a*-action the button will not go down and the observer may try another one; if the process can start with an *a*-action it will do so and the button goes down. Furthermore the green light switches on. During the execution of *a* no buttons can be pressed. As soon as the execution of *a* is completed the light switches off, so that the observer knows that the process is ready for a new trial. Reactive machines as described above originate from MILNER [134, 135].

Next I will discuss the equivalences that originate from the various reactive machines. First consider the reactive machine that resembles the failure trace machine, thus without menu-lights and *undo*-button. An observation on such a machine consists of a sequence of accepted and refused actions. Such a sequence can be modelled as a failure trace where all refusal sets are singletons. For finitely branching processes the resulting equivalence is exactly the equivalence that originates from PHILLIPS notion of *refusal testing* [150]. There it is called *refusal equivalence*. The following proposition shows that for finitely branching processes refusal equivalence coincides with failure equivalence.

PROPOSITION 1.6: Let $p \in \mathbb{A}$, $\sigma \in FT(p)$ and $Cont(\sigma) = \{a \in Act \mid \sigma a \in FT(p)\}$.

- i. Then, for $X \subseteq Act$, $\sigma X \rho \in FT(p) \Leftrightarrow \sigma(X \cap Cont(\sigma)) \rho \in FT(p)$.
- ii. If p is finitely branching then $Cont(\sigma)$ is finite.
- iii. $\sigma(X \cup Y) \rho \in FT(p) \Leftrightarrow \sigma XY \rho \in FT(p)$.

PROOF: Straightforward. □

If the menu-lights are added to the reactive failure trace machine considered above one can observe ready trace sets, and the green light is redundant. If the green light (as well as the menu-lights) are removed one can only test trace equivalence, since any refusal may be caused by the last action not being ready yet. Reactive machines seem to be unsuited for testing completed trace and failure equivalence. If the menu-lights and the *undo*-button are added to the reactive failure trace machine one gets ready simulation again and if only the *undo*-button is added one obtains an equivalence that may be called *refusal simulation equivalence* and coincides with ready simulation equivalence on the domain of finitely branching processes. The following *refusal simulation formulas* originate from BLOOM, ISTRAIL & MEYER [35].

DEFINITION: The *refusal simulation formulas* and the corresponding satisfaction relation are defined by adding to the definitions of Subsection 1.9 the following extra clauses:

- 4. If $a \in Act$ then $\neg a \in \mathcal{L}_{CS}$.
- 4. $p \models \neg a$ if $a \notin I(p)$.

An alternative family of testing scenarios with reactive machines can be obtained by allowing the observer to try to depress more than one button at a time. In order to influence a particular choice, the observer could already start exercising pressure on buttons during the execution of the preceding action (when no button can go down). When this preceding action is finished, at most one of the buttons will go down. These testing scenarios are equipotent with the generative ones: putting pressure on a button is equivalent to setting the corresponding switch on 'free'.

1.12. 2-nested simulation semantics. *2-nested simulation equivalence* popped up naturally in GROOTE & VAANDRAGER [107] as the coarsest congruence with respect to a large and general class of operators that is finer than completed trace equivalence. In order to obtain a testing scenario for this equivalence one has to introduce the rather unnatural notion of a *lookahead* [107]: The *2-nested simulation machine* is a variant of the ready trace machine with duplicator, where in an idle state the machine not only tells which actions are on the menu, but even which simulation formulas are satisfied in the current state.

DEFINITION: A *2-nested simulation* is a simulation contained in simulation equivalence ($\leq$). p and q are *2-nested simulation equivalent* if there exists a 2-nested simulation R with pRq and a 2-nested simulation S with qSp .

1.13. Bisimulation semantics. The testing scenario for bisimulation semantics, as presented in MILNER [134] is the oldest and most powerful testing scenario, from which most others have been derived by omitting some of its features. It was based on a reactive failure trace machine with duplicator, but additionally the observer is equipped with the capacity of *global testing*. Global testing is described in ABRAMSKY [1] as: "the ability to enumerate all (of finitely many) possible 'operating environments' at each stage of the test, so as to guarantee that all nondeterministic branches will be pursued by various copies of the subject process". MILNER [134] implemented global testing by assuming that

- (i) It is the *weather* which determines in each state which *a*-move will occur in response of pressing the *a*-button (if the process under investigation is capable of doing an *a*-move at all);
- (ii) The weather has only finitely many states - at least as far as choice-resolution is concerned;
- (iii) We can control the weather.

Now it can be ensured that all possible moves a process can perform in reaction on an *a*-experiment will be investigated by simply performing the experiment in all possible weather conditions. Unfortunately, as remarked in MILNER [135], the second assumption implies that the amount of different *a*-moves an investigated process can perform is bounded by the number of possible weather conditions; so for general application this condition has to be relaxed.

A different implementation of global testing is given in LARSEN & SKOU [128]. They assumed that every transition in a transition system has a certain probability of being taken. Therefore an observer can with an arbitrary high degree of confidence assume that all transitions have been examined, simply by repeating an experiment many times.

As argued among others in BLOOM, ISTRAIL & MEYER [35], global testing in the above sense is a rather unrealistic testing ability. Once you assume that the observer is really as powerful as in the described scenarios, in fact more can be tested than only bisimulation equivalence: in the testing scenario of Milner also the correlation between weather conditions and transitions being taken by the investigated process can be recovered, and in that of Larsen & Skou one can determine the relative probabilities of the various transitions.

An observation in the global testing scenario can be represented as a formula in *Hennessy-Milner logic* [111] (*HML*). An HML formula is a simulation formula in which it is possible to indicate that certain branches are not present.

DEFINITION: The *HML-formulas* and the corresponding satisfaction relation are defined by adding to the definitions in Subsection 1.9 the clauses:

- 4. If $\phi \in \mathcal{L}$ then $\neg\phi \in \mathcal{L}$.
- 4. $p \models \neg\phi$ if $p \not\models \phi$.

Let $HML(p)$ denote the set of all HML-formula that are satisfied by the process p : $HML(p) = \{\phi \in \mathcal{L} \mid p \models \phi\}$. Two processes p and q are *HML-equivalent* if $HML(p) = HML(q)$.

For finitely branching processes HENNESSY & MILNER [111] provided the following characterization of this equivalence.

DEFINITION: Let $p, q \in \mathbf{A}$ be finitely branching processes. Then:

- $p \sim_0 q$ is always true.
- $p \sim_{n+1} q$ if for all $a \in \text{Act}$:
 - $p \xrightarrow{a} p'$ implies $\exists q': q \xrightarrow{a} q'$ and $p' \sim_n q'$;
 - $q \xrightarrow{a} q'$ implies $\exists p': p \xrightarrow{a} p'$ and $p' \sim_n q'$.
- p and q are *observational equivalent*, notation $p \sim q$, if $p \sim_n q$ for every $n \in \mathbb{N}$.

PROPOSITION 1.7: Let $p, q \in \mathbf{A}$ be finitely branching processes. Then $p \sim q \Leftrightarrow \text{HML}(p) = \text{HML}(q)$.

PROOF: In HENNESSY & MILNER [112]. □

As observed by PARK [145], for finitely branching processes observation equivalence can be reformulated as bisimulation equivalence.

DEFINITION: A *bisimulation* is a binary relation R on processes, satisfying, for $a \in \text{Act}$:

- if pRq and $p \xrightarrow{a} p'$, then $\exists q': q \xrightarrow{a} q'$ and $p'Rq'$;
- if pRq and $q \xrightarrow{a} q'$, then $\exists p': p \xrightarrow{a} p'$ and $p'Rq'$.

Two processes p and q are *bisimilar*, notation $p \Leftrightarrow q$, if there exists a bisimulation R with pRq .

The relation $\Leftrightarrow$ is again a bisimulation. As for similarity, one easily checks that bisimilarity is an equivalence on $\mathbf{A}$. Hence the relation will be called *bisimulation equivalence*. Finally note that the concept of bisimulation does not change if in the definition above the action relations $\xrightarrow{a}$ were replaced by generalized action relations $\xrightarrow{\sigma}$.

PROPOSITION 1.8: Let $p, q \in \mathbf{A}$ be finitely branching processes. Then $p \Leftrightarrow q \Leftrightarrow p \sim q$.

PROOF: " $\Rightarrow$ ": Straightforward with induction. " $\Leftarrow$ " follows from Theorem 5.6 in MILNER [134]. □

For infinitely branching processes $\sim$ is coarser than $\Leftrightarrow$ and will be called *finitary bisimulation equivalence*.

Another characterization of bisimulation semantics can be given by means of ACZEL's universe $\mathcal{V}$ of non-well-founded sets [7]. This universe is an extension of the Von Neumann universe of well-founded sets, where the axiom of foundation (every chain $x_0 \ni x_1 \ni \dots$ terminates) is replaced by an *anti-foundation axiom*.

DEFINITION: Let B denote the unique function $\mathbb{B}:\mathbb{A}\rightarrow\mathcal{V}$ satisfying $\mathbb{B}(p)=\{<a, \mathbb{B}(q)> \mid p \xrightarrow{a} q\}$ for all $p\in\mathbb{A}$. Two processes p and q are *branching equivalent* if $B(p)=B(q)$.

It follows from Aczel's anti-foundation axiom that such a solution exists. In fact the axiom amounts to saying that systems of equations like the one above have unique solutions. In [7] there is also a section on communicating systems. There two processes are identified iff they are branching equivalent.

A similar idea underlies the semantics of DE BAKKER & ZUCKER [20], but there the domain of processes is a complete metric space and the definition of B above only works for finitely branching processes, and only if $=$ is interpreted as *isometry*, rather than equality, in order to stay in well-founded set theory. For finitely branching processes the semantics of De Bakker and Zucker coincides with the one of Aczel and also with bisimulation semantics. This is observed in VAN GLABBEK & RUTTEN [89], where also a proof can be found of the next proposition, saying that bisimulation equivalence coincides with branching equivalence.

PROPOSITION 1.9: Let $p, q \in \mathbb{A}$. Then $p \Leftrightarrow q \Leftrightarrow B(p)=B(q)$.

PROOF: " $\Leftarrow$ ". Let B be the relation, defined by pBq iff $B(p)=B(q)$, then it suffices to prove that B is a bisimulation. Suppose pBq and $p \xrightarrow{a} p'$. Then $<a, B(p')> \in B(p)=B(q)$. So by the definition of $B(q)$ there must be a process q' with $B(p')=B(q')$ and $q \xrightarrow{a} q'$. Hence $p'Bq'$, which had to be proved. The second requirement for B being a bisimulation follows by symmetry. " $\Rightarrow$ ". Let B^* denote the unique solution of

$$\mathbb{B}^*(p)=\{<a, \mathbb{B}^*(r')> \mid \exists r: r \Leftrightarrow p \text{ \& } r \xrightarrow{a} r'\}.$$

As for B it follows from the anti-foundation axiom that such a unique solution exists. From the symmetry and transitivity of $\Leftrightarrow$ it follows that

$$p \Leftrightarrow q \Rightarrow B^*(p)=B^*(q). \quad (*)$$

Hence it remains to be proven that $B^*=B$. This can be done by showing that B^* satisfies the equations $\mathbb{B}(p)=\{<a, \mathbb{B}(q)> \mid p \xrightarrow{a} q\}$, which have B as unique solution. So it has to be established that $B^*(p)=\{<a, B^*(q)> \mid p \xrightarrow{a} q\}$. The direction " $\supseteq$ " follows directly from the reflexivity of $\Leftrightarrow$. For " $\subseteq$ ", suppose $<a, X> \in B^*(p)$. Then $\exists r: r \Leftrightarrow p$, $r \xrightarrow{a} r'$ and $X=B^*(r')$. Since $\Leftrightarrow$ is a bisimulation, $\exists p': p \xrightarrow{a} p'$ and $r' \Leftrightarrow p'$. Now from (*) it follows that $X=B^*(r')=B^*(p')$. Therefore $<a, X> \in \{<a, B^*(q)> \mid p \xrightarrow{a} q\}$, which had to be established. $\square$

2. THE SEMANTIC LATTICE

2.1. Ordering the equivalences for finitely branching processes. In Section 1 twelve semantics were defined that are different for finitely branching processes. These will be abbreviated by T , CT , F , R , FT , RT , S , CS , RS , PF , $2S$ and B . Write $\mathfrak{S} \leq \mathfrak{T}$ if semantics $\mathfrak{S}$ makes at least as much identifications as semantics $\mathfrak{T}$. This is the case if the equivalence corresponding with $\mathfrak{S}$ is equal to or coarser than the one corresponding with $\mathfrak{T}$.

THEOREM 2.1: $T \leq CT \leq F \leq R \leq RT$, $F \leq FT \leq RT \leq RS \leq 2S \leq B$,
 $T \leq S \leq CS \leq RS$, $CT \leq CS$ and $R \leq PF \leq 2S$.

PROOF: The first statement is trivial. For the next five statements it suffices to show that $CT(p)$ can be expressed in terms of $F(p)$, $F(p)$ in terms of $R(p)$, $R(p)$ in terms of $RT(p)$, $F(p)$ in terms of $FT(p)$ and $FT(p)$ in terms of $RT(p)$.

- $CT(p) = \{\sigma \in Act^* \mid \langle \sigma, Act \rangle \in F(p)\}$.
- $\langle \sigma, X \rangle \in F(p) \Leftrightarrow \exists Y \subseteq Act: \langle \sigma, Y \rangle \in R(p) \ \& \ X \cap Y = \emptyset$.
- $\langle \sigma, X \rangle \in R(p) \Leftrightarrow \sigma X \in RT(p)$.
- $\langle \sigma, X \rangle \in F(p) \Leftrightarrow \sigma X \in FT(p)$.
- $\sigma = \sigma_1 \sigma_2 \cdots \sigma_n \in FT(p) \ (\sigma_i \in Act \cup \mathcal{P}(Act)) \Leftrightarrow$
 $\exists \rho = \rho_1 \rho_2 \cdots \rho_n \in RT(p) \ (\rho_i \in Act \cup \mathcal{P}(Act))$ such that for $i = 1, \dots, n$ either
 $\sigma_i = \rho_i \in Act$ or $\sigma_i, \rho_i \subseteq Act$ and $\sigma_i \cap \rho_i = \emptyset$.

The remaining statements are (also) trivial. □

Theorem 2.1 is illustrated in Figure 1. There, however, completed simulation semantics is missing, since it did not occur in the literature.

2.2. Ordering the equivalences for infinitely branching processes. When the restriction to finitely branching processes is dropped, there exists a finitary and an infinitary variant of each of these semantics, depending on whether or not infinite observations are taken into account. These versions will be notationally distinguished by means of superscripts ‘*’ and ‘ω’ respectively; the unsubscripted abbreviation will, for historical reasons, refer to the infinitary versions in case of ‘simulation’-like semantics and to the finitary versions otherwise. For the semantics that are based on refusal sets, there exists even a third version, namely when also the refusal sets are required to be finite. These will be denoted by means of a superscript ‘-’. So F^- denotes failure semantics as defined in [42] (see Subsection 1.4), R^- denotes acceptance-refusal semantics [159] (Subsection 1.7), FT^- denotes refusal semantics (Subsection 1.11), RS^- denotes refusal simulation semantics (also Subsection 1.11) and B^- denotes HML-semantics (Subsection 1.13). Finally, finitary failure simulation semantics (FS^*) can be distinguished from finitary ready simulation semantics (RS^*), although FS and RS , as well as FS^- and RS^- , coincide. Now the $\leq$ -relation is represented by arrows in Figure 7.




FIGURE 7. The infinitary linear time - branching time spectrum

THEOREM 2.2: Let $\mathcal{S}, \mathcal{T}$ be any two of the semantics mentioned above. Then $\mathcal{S} \leq \mathcal{T}$ whenever this is indicated in Figure 7.

Again the proof is straightforward. If the labelled transition system $\mathbb{A}$ on which these semantic equivalences are defined is large enough, then they are all different and $\mathcal{S} \leq \mathcal{T}$ holds only if this follows from Theorem 2.2 (and the fact that $\leq$ is a partial order), as will be shown in Subsection 2.8. However, for certain labelled transition systems much more identifications can be made. It has been remarked already that for finitely branching processes all semantics that are connected by dashed arrows in Figure 7 coincide. This result will be slightly strengthened in the next subsection. In the subsequent subsection a class of processes will be defined on which all the semantics coincide.

2.3. Image finite processes.

DEFINITION: A process $p \in \mathbb{A}$ is *image finite* if for each $\sigma \in \text{Act}^*$ the set $\{q \in \mathbb{A} \mid p \xrightarrow{\sigma} q\}$ is finite.

Note that finitely branching processes are image finite, but the reverse does not hold.

THEOREM 2.3: *On a domain of image finite processes, semantics that are connected with a dashed arrow in Figure 7 coincide.*

PROOF: For the upper two arrows, connecting HML-semantics with finitary bisimulation semantics and finitary bisimulation semantics with bisimulation semantics, the proof has been given in HENNESSY & MILNER [112]. For the other simulation-like semantics the proof goes likewise. For the trace-like semantics the correspondence between the finitary and infinitary versions (the arrows on the right) follows directly from König's lemma. Here I only prove the correspondence between F^- and F ; the remaining cases can be proved likewise.

It has to be established that, for image finite processes p and $q \in \mathbb{A}$, $F^-(p) = F^-(q) \Rightarrow F(p) = F(q)$, where $F^-(p)$ denotes the set of failure pairs $\langle \sigma, X \rangle$ of p with finite refusal set X . The reverse implication is trivial. For finitely branching processes $F(p)$ is completely determined by $F^-(p)$ (Proposition 1.1), from which the implication follows. For arbitrary image finite processes this is no longer the case, but the implication still holds.

Let p and $q \in \mathbb{A}$ be two image finite processes with $F(p) \neq F(q)$. Say there is a failure pair $\langle \sigma, X \rangle \in F(p) - F(q)$. By image finiteness of q there are only finitely many processes r_i with $q \xrightarrow{\sigma} r_i$, and for each of those there is an action $a_i \in I(r_i) \cap X$ (otherwise $\langle \sigma, X \rangle$ would be a failure pair of q). Let Y be the set of all those a_i 's. then Y is a finite subset of X , so $\langle \sigma, Y \rangle \in F^-(p)$. On the other hand $a_i \in I(r_i) \cap Y$ for all r_i , so $\langle \sigma, Y \rangle \notin F^-(q)$. $\square$

2.4. Deterministic processes.

DEFINITION: A process p is *deterministic* if $p \xrightarrow{\sigma} q$ & $p \xrightarrow{\sigma} r \Rightarrow q = r$.

REMARK: If p is deterministic and $p \xrightarrow{\sigma} p'$ then also p' is deterministic. Hence any domain of processes on which action relations are defined, has a subdomain of deterministic processes with the inherited action relations. (A similar remark can be made for image finite processes.)

PROOF: Suppose $p' \xrightarrow{\rho} q$ and $p' \xrightarrow{\rho} r$. Then $p \xrightarrow{\sigma\rho} q$ and $p \xrightarrow{\sigma\rho} r$, so $q = r$.

THEOREM 2.4 (PARK [145]): *On a domain of deterministic processes all semantics on the infinitary linear time - branching time spectrum coincide.*

PROOF: Because of Theorem 2.2 it suffices to show that $BS \leq T$. This is the case if $T(p) = T(q) \Rightarrow p \leq q$ for any two deterministic processes p and q . Let R be the relation, defined by $p R q$ iff $T(p) = T(q)$, then it suffices to prove that

R is a bisimulation. Suppose pRq and $p \xrightarrow{\sigma} p'$. Then $\sigma \in T(p) = T(q)$. So there is a process q' with $q \xrightarrow{\sigma} q'$. Now let $\rho \in T(p')$. Then $\exists r: p' \xrightarrow{\rho} r$. Hence $p \xrightarrow{\sigma\rho} r$ and $\sigma\rho \in T(p) = T(q)$. So there must be a process s with $q \xrightarrow{\sigma\rho} s$. By the definition of the generalized action relations $\exists t: q \xrightarrow{\sigma} t \xrightarrow{\rho} s$, and since q is deterministic, $t = q'$. Thus $\rho \in T(q')$, and from this it follows that $T(p') \subseteq T(q')$. Since also p is deterministic the converse can be established in the same way, and together this yields $T(p') = T(q')$, or $p'Rq'$. This finishes the proof. $\square$

2.5. Process graphs. In process theory it is common practice to represent processes as elements in a mathematical domain. The semantics of a process theory can then be modelled as an equivalence on such a domain. In Section 1 several semantic equivalences were defined on any domain of sequential processes which is provided with action relations. Such a domain was called a labelled transition system. In Section 3 a term domain $\mathbb{P}$ with action relations will be presented for which these definitions apply. The present subsection introduces one of the most popular labelled transition systems: the domain $\mathbb{G}$ of *process graphs* or *state transition diagrams*.

DEFINITION: A *process graph* over a given alphabet Act is a rooted, directed graph whose edges are labelled by elements of Act . Formally, a process graph g is a triple $(\text{NODES}(g), \text{EDGES}(g), \text{ROOT}(g))$, where

- $\text{NODES}(g)$ is a set, whose elements are called the *nodes* or *states* of g ,
- $\text{ROOT}(g) \in \text{NODES}(g)$ is a special node: the *root* or *initial state* of g ,
- and $\text{EDGES}(g) \subseteq \text{NODES}(g) \times Act \times \text{NODES}(g)$ is a set of triples (s, a, t) with $s, t \in \text{NODES}(g)$ and $a \in Act$: the *edges* or *transitions* of g .

If $e = (s, a, t) \in \text{EDGES}(g)$, one says that e goes from s to t . A (finite) *path* π in a process graph is an alternating sequence of nodes and edges, starting and ending with a node, such that each edge goes from the node before it to the node after it. If $\pi = s_0(s_0, a_1, s_1)s_1(s_1, a_2, s_2) \cdots (s_{n-1}, a_n, s_n)s_n$, also denoted as $\pi: s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \cdots \xrightarrow{a_n} s_n$, one says that π goes from s_0 to s_n ; it starts in s_0 and ends in $\text{end}(\pi) = s_n$. Let $\text{PATHS}(g)$ be the set of paths in g starting from the root. If s and t are nodes in a process graph then t can be reached from s if there is a path going from s to t . A process graph is said to be *connected* if all its nodes can be reached from the root; it is a *tree* if each node can be reached from the root by exactly one path. Let $\mathbb{G}$ be the domain of connected process graphs over a given alphabet Act .

DEFINITION: For $g \in \mathbb{G}$ and $s \in \text{NODES}(g)$, let g_s be the process graph defined by

- $\text{NODES}(g_s) = \{t \in \text{NODES}(g) \mid \text{there is a path going from } s \text{ to } t\}$,
- $\text{ROOT}(g_s) = s \in \text{NODES}(g_s)$,
- and $(t, a, u) \in \text{EDGES}(g_s)$ iff $t, u \in \text{NODES}(g_s)$ and $(t, a, u) \in \text{EDGES}(g)$.

Of course $g_s \in \mathbb{G}$. Remark that $g_{\text{root}(g)} = g$. Now on $\mathbb{G}$ action relations $\xrightarrow{a}$ for $a \in \text{Act}$ are defined by $g \xrightarrow{a} h$ iff $(\text{root}(g), a, s) \in \text{EDGES}(g)$ and $h = g_s$. This makes $\mathbb{G}$ into a labelled transition system. Hence all semantic equivalences of Section 1 are well-defined on $\mathbb{G}$. Below the sets of observations $O(g)$ for $O \in \{T, CT, R, F, RT, FT\}$ and $g \in \mathbb{G}$, are characterized in terms of the paths of g , rather than the generalized action relations between graphs.

DEFINITION: Let $g \in \mathbb{G}$ and let $\pi: s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} s_n \in \text{PATHS}(g)$. Consider the following notions:

- the *trace* associated to π : $T(\pi) = a_1 a_2 \dots a_n \in \text{Act}^*$;
- the *menu* of a node $s \in \text{NODES}(g)$: $I(s) = \{a \in \text{Act} \mid \exists t: (s, a, t) \in \text{EDGES}(g)\}$;
- the *ready pair* associated to π : $R(\pi) = \langle T(\pi), I(s_n) \rangle$;
- the *failure set* of π : $F(\pi) = \{ \langle T(\pi), X \rangle \mid I(s_n) \cap X = \emptyset \}$;
- the *ready trace set* of π : $RT(\pi)$ is the smallest subset of $(\text{Act} \cup \mathcal{P}(\text{Act}))^*$ satisfying
 - $I(s_0) a_1 I(s_1) a_2 \dots a_n I(s_n) \in RT(\pi)$,
 - $\sigma X \rho \in RT(\pi) \Rightarrow \sigma \rho \in RT(\pi)$,
 - $\sigma X \rho \in RT(\pi) \Rightarrow \sigma X X \rho \in RT(\pi)$;
- and the *failure trace set* of π : $FT(\pi)$ is the smallest subset of $(\text{Act} \cup \mathcal{P}(\text{Act}))^*$ satisfying
 - $(A - I(s_0)) a_1 (A - I(s_1)) a_2 \dots a_n (A - I(s_n)) \in FT(\pi)$,
 - $\sigma X \rho \in FT(\pi) \Rightarrow \sigma \rho \in FT(\pi)$,
 - $\sigma X \rho \in FT(\pi) \Rightarrow \sigma X X \rho \in FT(\pi)$,
 - $\sigma X \rho \in FT(\pi) \wedge Y \subseteq X \Rightarrow \sigma Y \rho \in FT(\pi)$.

PROPOSITION 2.5:

$$\begin{aligned}
 T(g) &= \{T(\pi) \mid \pi \in \text{PATHS}(g)\} \\
 CT(g) &= \{T(\pi) \mid \pi \in \text{PATHS}(g) \wedge I(\text{end}(\pi)) = \emptyset\} \\
 R(g) &= \{R(\pi) \mid \pi \in \text{PATHS}(g)\} \\
 F(g) &= \bigcup_{\pi \in \text{PATHS}(g)} F(\pi) \\
 RT(g) &= \bigcup_{\pi \in \text{PATHS}(g)} RT(\pi) \\
 FT(g) &= \bigcup_{\pi \in \text{PATHS}(g)} FT(\pi)
 \end{aligned}$$

PROOF: Straightforward. □

Analogously, the simulation-like equivalences can be characterized by means of simulation relations between the nodes of two process graphs, rather than between process graphs themselves. Below this is done for bisimulation equivalence.

DEFINITION: Let $g, h \in \mathbb{G}$. A *bisimulation* between g and h is a binary relation $R \subseteq \text{NODES}(g) \times \text{NODES}(h)$, satisfying:

1. $\text{ROOT}(g) R \text{ROOT}(h)$.
2. If $s R t$ and $(s, a, s') \in \text{EDGES}(g)$, then there is an edge $(t, a, t') \in \text{EDGES}(h)$ such that $s' R t'$.
3. If $s R t$ and $(t, a, t') \in \text{EDGES}(h)$, then there is an edge $(s, a, s') \in \text{EDGES}(g)$ such that $s' R t'$.

This definition is illustrated in Figure 8. Now it follows easily that two graphs g and h are bisimilar iff there exists a bisimulation between them.




FIGURE 8. A bisimulation

Proposition 2.5 yields a technique for deciding that two process graphs are ready trace equivalent, c.q. failure trace equivalent, without calculating their entire ready trace or failure trace set.

Let $g, h \in \mathbb{G}$, $\pi: s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} s_n \in \text{PATHS}(g)$ and $\pi': t_0 \xrightarrow{a_1} t_1 \xrightarrow{a_2} \dots \xrightarrow{a_m} t_m \in \text{PATHS}(h)$. Path π' is a *failure trace augmentation* of π , notation $\pi \leq_{FT} \pi'$, if $FT(\pi) \subseteq FT(\pi')$. This is the case exactly when $n = m$ and $I(t_i) \subseteq I(s_i)$ for $i = 1, \dots, n$. Write $\pi =_{FT} \pi'$ for $\pi \leq_{FT} \pi' \wedge \pi' \leq_{FT} \pi$. It follows that $\pi =_{FT} \pi' \Leftrightarrow FT(\pi) = FT(\pi') \Leftrightarrow RT(\pi) = RT(\pi')$. From this the following can be concluded.

COROLLARY 2.5: Two process graphs $g, h \in \mathbb{G}$ are ready trace equivalent iff

- for any path $\pi \in \text{PATHS}(g)$ in g there is a $\pi' \in \text{PATHS}(h)$ such that $\pi =_{FT} \pi'$
- and for any path $\pi \in \text{PATHS}(g)$ in h there is a $\pi' \in \text{PATHS}(g)$ such that $\pi =_{FT} \pi'$.

They are failure trace equivalent iff

- for any path $\pi \in \text{PATHS}(g)$ in g there is a $\pi' \in \text{PATHS}(h)$ such that $\pi \leq_{FT} \pi'$
- and for any path $\pi \in \text{PATHS}(g)$ in h there is a $\pi' \in \text{PATHS}(g)$ such that $\pi \leq_{FT} \pi'$.

If g and h are moreover without infinite paths, then it suffices to check the requirements above for maximal paths.

2.6. Drawing process graphs.

DEFINITION: Let $g, h \in \mathbb{G}$. A *graph isomorphism* between g and h is a bijective function $f : \text{NODES}(g) \rightarrow \text{NODES}(h)$ satisfying

- $f(\text{ROOT}(g)) = \text{ROOT}(h)$ and
- $(s, a, t) \in \text{EDGES}(g) \Leftrightarrow (f(s), a, f(t)) \in \text{EDGES}(h)$.

Graphs g and h are *isomorphic*, notation $g \cong h$, if there exists a graph isomorphism between them.

In this case g and h differ only in the identity of their nodes. Remark that graph isomorphism is an equivalence on $\mathbb{G}$.

PROPOSITION 2.6: For $g, h \in \mathbb{G}$, $g \cong h$ iff there exists a bisimulation R between g and h , satisfying

4. If sRt and uRv then $s = u \Leftrightarrow t = v$.

PROOF: Suppose $g \cong h$. Let $f : \text{NODES}(g) \rightarrow \text{NODES}(h)$ be a graph isomorphism. Define $R \subseteq \text{NODES}(g) \times \text{NODES}(h)$ by sRt iff $f(s) = t$. Then it is routine to check that R satisfies clauses 1, 2, 3 and 4. Now suppose R is a bisimulation between g and h , satisfying 4. Define $f : \text{NODES}(g) \rightarrow \text{NODES}(h)$ by $f(s) = t$ iff sRt . Since g is connected it follows from the definition of a bisimulation that for each s such a t can be found. Furthermore direction " $\Rightarrow$ " of clause 4 implies that $f(s)$ is uniquely determined. Hence f is well-defined. Now direction " $\Leftarrow$ " of clause 4 implies that f is injective. From the connectedness of h it follows that f is also surjective, and hence a bijection. Finally clauses 1, 2 and 3 imply that f is a graph isomorphism. $\square$

COROLLARY: If $g \cong h$ then g and h are equivalent according to all semantic equivalences of Section 1.

Finitely branching connected process graphs can be pictured by using open dots ($\circ$) to denote nodes, and labelled arrows to denote edges, as can be seen in Subsection 2.8. There is no need to mark the root of such a process graph if it can be recognized as the unique node without incoming edges, as is the case in all my examples. These pictures determine process graphs only up to graph isomorphism, but usually this suffices since it is virtually never needed to distinguish between isomorphic graphs.

2.7. Embedding labelled transition systems in $\mathbb{G}$. Let $\mathbb{A}$ be an arbitrary labelled transition system and let $p \in \mathbb{A}$. The *canonical graph* $G(p)$ of p is defined as follows:

- $\text{NODES}(G(p)) = \{q \in \mathbb{A} \mid \exists \sigma \in A^* : p \xrightarrow{\sigma} q\},$
- $\text{ROOT}(G(p)) = p \in \text{NODES}(G(p)),$
- and $(q, a, r) \in \text{EDGES}(G(p))$ iff $q, r \in \text{NODES}(G(p))$ and $q \xrightarrow{a} r.$

Of course $G(p) \in \mathbb{G}$. This means G is a function from $\mathbb{A}$ to $\mathbb{G}$.

PROPOSITION 2.7: $G: \mathbb{A} \rightarrow \mathbb{G}$ is an injective function, satisfying, for $a \in \text{Act}$:
 $G(p) \xrightarrow{a} G(q) \Leftrightarrow p \xrightarrow{a} q.$

PROOF: Trivial. □

COROLLARY: For $p \in \mathbb{A}$ and $O \in \{T, CT, F, R, FT, RT, S, CS, RS, PF, 2S, B\}$,
 $O(G(p)) = O(p).$

Proposition 2.7 says that G is an *embedding* of $\mathbb{A}$ in $\mathbb{G}$. It implies that any labelled transition system over Act can be represented as a subclass $G(\mathbb{A}) = \{G(p) \in \mathbb{G} \mid p \in \mathbb{A}\}$ of $\mathbb{G}$.

Since $\mathbb{G}$ is also a labelled transition system, G can be applied to $\mathbb{G}$ itself. The following proposition says that the function $G: \mathbb{G} \rightarrow \mathbb{G}$ leaves its arguments intact up to graph isomorphism.

PROPOSITION 2.8: For $g \in \mathbb{G}$, $G(g) \cong g.$

PROOF: Remark that $\text{NODES}(G(g)) = \{g_s \mid s \in \text{NODES}(g)\}.$ Now the function $f: \text{NODES}(G(g)) \rightarrow \text{NODES}(g)$ defined by $f(g_s) = s$ is a graph isomorphism. □

2.8. Counterexamples. In this subsection a number of examples will be presented, showing that on $\mathbb{G}$ all semantic notions mentioned in Theorem 2.2 are different and $\mathbb{S} \leq \mathbb{T}$ holds only if this follows from that theorem. Moreover, apart from the examples needed to show the difference between semantics that are connected by a dashed arrow in Figure 7, all examples will use finite processes only. Thus it follows that neither the ordering of Theorem 2.1 nor the ordering of Theorem 2.2 can be further expanded. Let $\mathbb{H}$ be the set of finite connected process graphs. Here a process graph g is *finite* if $\text{PATHS}(g)$ is finite. Finite graphs are acyclic and have only finitely many nodes and edges. They represent finite processes.

THEOREM 2.9: Let $\mathbb{S}$ and $\mathbb{T}$ be semantics on $\mathbb{H}$ from the series $T, CT, F, R, FT, RT, S, CS, RS, PF, 2S, B$. Then $\mathbb{S} \leq \mathbb{T}$ only if this follows from Theorem 2.1. (and the fact that $\leq$ is a partial order).

PROOF: The following counterexamples provide for any statement $\mathbb{S} \leq \mathbb{T}$, not following from Theorem 2.1 and the fact that $\leq$ is a partial order, two finite connected process graphs that are identified in $\mathbb{T}$, but distinguished in $\mathbb{S}$.




FIGURE 9

1. $T \not\equiv CT$. For the graphs of Figure 9, $T(left) = T(right) = \{\epsilon, a, ab\}$, whereas $CT(left) \neq CT(right)$ (since $a \in CT(left) - CT(right)$). Hence they are identified in trace semantics but distinguished in completed trace semantics. Furthermore the two graphs are simulation equivalent (the construction of the two simulations is left to the reader). Since $\leq$ is a partial order, the same example shows that $\mathcal{S} \not\equiv \mathcal{T}$ for $\mathcal{S} \in \{CT, CS, F, R, FT, RT, RS, PF, 2S, B\}$ and $\mathcal{T} \in \{T, S\}$.




FIGURE 10

2. $CT \not\equiv F$. For the graphs of Figure 10, $CT(left) = CT(right) = \{ab, ac\}$, whereas $F(left) \neq F(right)$ (since $\langle a, \{b\} \rangle \in F(left) - F(right)$). Hence they are identified in completed trace semantics but distinguished in failure semantics. Furthermore the two graphs are completed simulation equivalent (the construction of the two completed simulations is again left to the reader). Since $\leq$ is a partial order, the same example shows that $\mathcal{S} \not\equiv \mathcal{T}$ for $\mathcal{S} \in \{F, R, FT, RT, RS, PF, 2S, B\}$ and $\mathcal{T} \in \{CT, CS\}$.




FIGURE 11

3. $FT \not\approx R$. For the graphs of Figure 11, $FT(left) = FT(right)$, whereas $R(left) \neq R(right)$. The first statement follows from Corollary 2.5, since the new maximal paths at the right-hand side are both failure trace augmented by the two maximal paths both sides have in common. The second one follows since $\langle a, \{b, c\} \rangle \in R(right) - R(left)$. Hence these processes are identified in failure trace semantics but distinguished in readiness semantics. Since $\leq$ is a partial order, the same example shows that $\mathcal{S} \not\approx \mathcal{T}$ for any $\mathcal{S} \leq FT$ and $\mathcal{T} \geq R$, so in particular $F \not\approx R$ and $FT \not\approx RT$.




FIGURE 12

4. $R \not\approx FT$. For the graphs of Figure 12, $R(left) = R(right)$, whereas $FT(left) \neq FT(right)$. The first statement follows since in the second graph only 4 ready pairs swapped places. The second one follows since $a\{b\}ce \in FT(left) - FT(right)$. Hence these processes are identified in readiness semantics but distinguished in failure trace semantics. Since $\leq$ is a partial order, the same example shows that $\mathcal{S} \not\approx \mathcal{T}$ for any $\mathcal{S} \leq R$ and $\mathcal{T} \geq FT$, so in particular $F \not\approx FT$ and $R \not\approx RT$. Since $PF(left) \neq PF(right)$ this example does not show that $PF \not\approx FT$. It is left as an exercise to the reader to adapt the example so that also that is established.




FIGURE 13

5. $RT \not\approx S$. For the graphs of Figure 13, $RT(left) = RT(right)$, whereas $S(left) \neq S(right)$. The first statement follows immediately from Corollary 2.5. The second one follows since $a(bcT \wedge bdT) \in S(right) - S(left)$. Hence these processes are identified in ready trace semantics but distinguished in simulation semantics. Since $\leq$ is a partial order, the same example shows that $S \not\approx T$ for any $S \leq RT$ and $T \geq S$, so in particular $T \not\approx S$, $CT \not\approx CS$ and $RT \not\approx RS$.




FIGURE 14

6. $RS \not\approx 2S$. The graphs of Figure 14 are ready simulation equivalent, but not 2-nested simulation equivalent. There exists exactly one simulation from *right* by *left*, namely the one mapping *right* on the right-hand side of *left*, and this simulation is a ready simulation as well as a 2-nested simulation. There also exists exactly one simulation from *left* by *right*, which maps the black node on the *left* on the black node on the *right*. This simulation is a ready simulation (related nodes have the same menu of initial actions) but not a 2-nested simulation (the two subgraphs originating from the two black nodes are not

simulation equivalent). Hence $RS \not\approx 2S$. Furthermore $PF(left) \neq PF(right)$, since $\langle a, \{\epsilon, b, bc\} \rangle \in PF(left) - PF(right)$. Hence $\mathfrak{S} \not\approx PF$ for any $\mathfrak{S} \leq RS$.




FIGURE 15

7. $2S \not\approx B$. The graphs of Figure 15 are 2-nested simulation equivalent, but not bisimulation equivalent. There now exists 2-nested simulations in both directions since the two subgraphs originating from the two black nodes are simulation equivalent. However, $a \neg b \neg cT \in HML(left) - HML(right)$. $\square$

THEOREM 2.10: Let $\mathfrak{S}$ and $\mathfrak{T}$ be semantics on $\mathbb{G}$ mentioned in Subsection 2.2. Then $\mathfrak{S} \leq \mathfrak{T}$ only if this follows from Theorem 2.2. (and the fact that $\leq$ is a partial order).

PROOF: The following counterexamples, together with the ones used in the previous proof, provide for any statement $\mathfrak{S} \leq \mathfrak{T}$, not following from Theorem 2.2 and the fact that $\leq$ is a partial order, two connected process graphs that are identified in $\mathfrak{T}$, but distinguished in $\mathfrak{S}$.

8. $B^* \not\approx T^\omega$. The graphs of Figure 4 in Chapter III are finitary bisimulation equivalent (as follows straightforward with induction) but not infinitary trace equivalent (since only the graph at the right has an infinite trace). Since $\leq$ is a partial order it follows that $\mathfrak{S} \not\approx \mathfrak{T}$ for $\mathfrak{S} \leq B^*$ and $\mathfrak{T} \geq T^\omega$.

9. $B^- \not\approx CT$. For the graphs of Figure 16, $HML(left) = HML(right)$, whereas $CT(left) \neq CT(right)$. The first statement follows since by means of HML-formulas one can only say that a *finite* set of actions can not take place in a certain state. The second one follows since $a \in CT(left) - CT(right)$. Since $\leq$ is a partial order it follows that $\mathfrak{S} \not\approx \mathfrak{T}$ for $\mathfrak{S} \leq B^-$ and $\mathfrak{T} \geq CT$.

10. $FS^* \not\approx R$. The counterexample showing $FS^* \not\approx R$ and hence $FS^* \not\approx RS^*$ is a bit complicated to draw, but can be easily described in terms of the previous




FIGURE 16

example. It consists of two trees, not containing paths with more than two edges, the upper parts of which are the same as the upper parts of the two trees in Figure 16. But any node with an incoming a -edge, has an outgoing edge labeled b_i for any $i \geq 1$, except for the i for which it had an outgoing b_i -edge in Figure 16. Now $FS^*(left) = FS^*(right)$, but $\langle a, \{b_1, b_2, \dots\} \rangle \in R(left) - R(right)$. $\square$

One could say that a semantics $\mathbb{S}$ respects deadlock behaviour iff $\mathbb{S} \geq CT$. Example 9 above then shows that none of the semantics on the left in Figure 7 respects deadlock behaviour; only the left-hand process of Figure 16 can deadlock after an a -move.

3. COMPLETE AXIOMATIZATIONS

3.1. A language for finite, concrete, sequential processes. Consider the following basic CCS- and CSP-like language BCCSP for finite, concrete, sequential processes over a given alphabet Act :

- inaction*: 0 (called *nil* or *stop*) is a constant, representing a process that refuses to do any action.
- action*: a is a unary operator for any action $a \in Act$. The expression ap represents a process, starting with an a -action and proceeding with p .
- choice*: $+$ is a binary operator. $p + q$ represents a process, first being involved in a choice between its summands p and q , and then proceeding as the chosen process.

The set $\mathbb{P}$ of (closed) *process expressions* or terms over this language is defined as usual:

- $0 \in \mathbb{P}$,
- $ap \in \mathbb{P}$ for any $a \in Act$ and $p \in \mathbb{P}$,
- $p + q \in \mathbb{P}$ for any $p, q \in \mathbb{P}$.

Subterms $a0$ may be abbreviated by a .

On $\mathbb{P}$ action relations $\xrightarrow{a}$ for $a \in Act$ are defined as the predicates on $\mathbb{P}$ gen-

erated by the *action rules* of Table 1. Here a ranges over Act and p and q over $\mathbb{P}$.

$ap \xrightarrow{a} p$	$\frac{p \xrightarrow{a} p'}{p + q \xrightarrow{a} p'}$	$\frac{q \xrightarrow{a} q'}{p + q \xrightarrow{a} q'}$
------------------------	---	---

TABLE 1

Now all semantic equivalences of Section 1 are well-defined on $\mathbb{P}$, and for each of the semantics it is determined when two process expressions denote the same process.

3.2. Axioms. In Table 2, complete axiomatizations can be found for ten of the twelve semantics of this chapter that differ on BCCSP. Axioms for 2-nested simulation and possible-futures semantics are more cumbersome, and the corresponding testing notions are less plausible. Therefore they have been omitted. In order to formulate the axioms, variables have to be added to the language as usual. In the axioms they are supposed to be universally quantified. Most of the axioms are axiom schemes, in the sense that there is one axiom for each substitution of actions from Act for the parameters a, b, c . Some of the axioms are conditional equations, using an auxiliary operator I . Thus provability is defined according to the standards of either first-order logic with equality or conditional equational logic. I is a unary operator that calculates the set of initial actions of a process expression, coded as a process expression again.

THEOREM 3.1: *For each of the semantics $O \in \{T, S, CT, CS, F, R, FT, RT, RS, B\}$ two process expressions $p, q \in \mathbb{P}$ are O -equivalent iff they can be proved equal from the axioms marked with ‘+’ in the column for O in Table 2. The axioms marked with ‘v’ or ‘ω’ are valid in O -semantics but not needed for the proof.*

PROOF: For F, R and B the proof is given in BERGSTRA, KLOP & OLDEROG [30] by means of *graph transformations*. A similar proof for RT can be found in BAETEN, BERGSTRA & KLOP [14]. For the remaining semantics a proof can be given along the same lines. $\square$

THEOREM 3.2: *Suppose Act is infinite. For each of the semantics $O \in \{T, S, CT, F, R, FT, RT, RS, B\}$ two BCCSP expressions with variables are O -equivalent iff they can be proved equal from the axioms marked with ‘+’ or ‘ω’ in the column for O in Table 2. It follows that the axioms marked with ‘v’ are derivable.*

PROOF: For $O \in \{T, CT, F, R, FT, RT, B\}$ this has been established by JAN

	<i>B</i>	<i>R</i>	<i>S</i>	<i>RT</i>	<i>FT</i>	<i>R</i>	<i>F</i>	<i>CS</i>	<i>CT</i>	<i>S</i>	<i>T</i>
$x + y = y + x$	+	+	+	+	+	+	+	+	+	+	+
$(x + y) + z = x + (y + z)$	+	+	+	+	+	+	+	+	+	+	+
$x + x = x$	+	+	+	+	+	+	+	+	+	+	+
$x + 0 = x$	+	+	+	+	+	+	+	+	+	+	+
$I(x) = I(y) \Rightarrow ax + a(x + y) = a(x + y)$			+	v	v	v	v	v	v	v	v
$I(x) = I(y) \Rightarrow ax + ay = a(x + y)$				+	+	v	v		v		v
$ax + ay = ax + ay + a(x + y)$					+		v		v		v
$a(bx + u) + a(by + v) = a(bx + by + u) + a(bx + by + v)$						+	+		v		v
$ax + a(y + z) = ax + a(x + y) + a(y + z)$							+		ω		v
$a(bx + u + y) = a(bx + u) + a(bx + u + y)$								+	v	v	v
$a(bx + u) + a(cy + v) = a(bx + cy + u + v)$									+		v
$ax + a(x + y) = a(x + y)$										+	v
$ax + ay = a(x + y)$											+
$I(0) = 0$	+	+	+	+	+	+	+	+	+	+	+
$I(ax) = a0$	+	+	+	+	+	+	+	+	+	+	+
$I(x + y) = I(x) + I(y)$	+	+	+	+	+	+	+	+	+	+	+

TABLE 2

FRISO GROOTE [104]. His proof for F , R , FT and RT can be applied to S and RS as well. $\square$

GROOTE also showed that if Act is finite, Theorem 3.2 does not hold for F , R , FT and RT . But for B and CT it suffices to assume that Act is nonempty, and for T it suffices to assume that Act has at least two elements. I don't know which cardinality restriction on Act is needed in the cases of S and RS . A complete axiomatization for open terms for completed simulation semantics has so far not been provided.

CONCLUDING REMARKS

In this chapter various semantic equivalences for concrete sequential processes are defined, motivated, compared and axiomatized. Of course many more equivalences can be given than the ones presented here. The reason for selecting just these, is that they can be motivated rather nicely and/or play a role in the literature on semantic equivalences. In ABRAMSKY & VICKERS [2] the observations which underly many of the semantics in this chapter are placed in a uniform algebraic framework, and some general completeness criteria are stated and proved. They also introduce *acceptance semantics*, which can be obtained from acceptance-refusal semantics (Section 1.7) by dropping the refusals, and analogously *acceptance trace semantics*. I'm not aware of any

reasonable testing scenario for these notions.

In Section 1.10 I remarked that a testing scenario for simulation and ready simulation semantics can be obtained by adding an *undo*-button to the scenario's for trace and ready trace semantics. Likewise SCHNOEBELEN [163] investigates the addition of an *undo*-button to the testing scenarios for completed trace, readiness, failure and failure trace semantics, thereby obtaining 3 new equivalences $CT_{\#}$, $R_{\#}$ and $F_{\#}$. *Undo*-failure trace equivalence coincides with finitary failure simulation equivalence, just like *undo*-trace and *undo*-ready trace equivalence coincide with finitary simulation and finitary ready simulation equivalence, and for image finite processes $R_{\#}$ coincides with $F_{\#}$. Furthermore $R \leq R_{\#} \leq RS^*$, $F \leq F_{\#} \leq FS^*$, $CT \leq CT_{\#} \leq CS^*$ and $S^* \leq CT_{\#} \leq F_{\#} \leq R_{\#}$.

It is left for a future occasion to give (and apply) criteria for selecting between these equivalences for particular applications (such as the complexity of deciding if two finite-state processes are equivalent, or the range of useful operators for which they are congruences). The work in this direction reported so far includes [35] and [107]. More research in this direction will be reported in [83].

An interesting topic is the generalization of this work to a setting with silent moves and/or with parallelism. In Chapter III the generalization of bisimulation semantics to a setting with silent steps is considered; in Chapters IV-VII bisimulation and trace semantics will be considered in a setting with parallelism. In both cases there turn out to be many interesting variations. The generalization of the rest of the spectrum to a setting with invisible actions will be tackled in [81]. Generalizing the entire spectrum to a setting with parallelism remains as of yet to be done. However, in many papers parts of a classification can be found already (see for instance [153]).

A generalization to preorders, instead of equivalences, can be obtained by replacing conditions like $O(p) = O(q)$ by $O(p) \subseteq O(q)$. Since preorders are often useful for verification purposes, it seems to be worthwhile to classify them as well. Also this will be done in [81].

Furthermore it would be interesting to give explicit representations of the equivalences, by representing processes as sets of observations instead of equivalence classes of process graphs, and defining operators like action prefixing and choice directly on these representations, as has been done for failure semantics in [42] and for readiness semantics in [144].

Chapter II

Modular Specifications in Process Algebra

With Curious Queues

Rob van Glabbeek & Frits Vaandrager

In recent years a wide variety of process algebras has been proposed in the literature. Often these process algebras are closely related: they can be viewed as homomorphic images, submodels or restrictions of each other. The aim of this chapter is to show how the semantical reality, consisting of a large number of closely related process algebras, can be reflected, and even used, on the level of algebraic specifications and in process verifications. This is done by means of the notion of a module. The simplest modules are building blocks of operators and axioms, each block describing a feature of concurrency in a certain semantical setting. These modules can then be combined by means of a union operator $+$, an export operator $\square$, allowing to forget some operators in a module, an operator H , changing semantics by taking homomorphic images, and an operator S which takes subalgebras. These operators enable us to combine modules in a subtle way, when the direct combination would be inconsistent. We show how auxiliary process algebra operators can be hidden when this is needed. Moreover it is demonstrated how new process combinators can be defined in terms of the more elementary ones in a clean way. As an illustration of our approach, a methodology is presented that can be used to specify FIFO-queues, and that facilitates verification of concurrent systems containing these queues.

TABLE OF CONTENTS

Introduction	56
1. Module Logic	61
2. Process Algebra	66
3. Applications of the Module Approach in Process Algebra	80
4. Queues	89
5. A Protocol Verification	102
Conclusions and Open Problems	109
Appendix: Logics	111

INTRODUCTION

During the last decade, a lot of research has been done on *process algebra*: the branch of theoretical computer science concerned with the modelling of concurrent systems as elements of an algebra. Besides the Calculus of Communicating Systems (CCS) of MILNER [134, 137], several related formalisms have been developed, such as the theory of Communicating Sequential Processes (CSP) of HOARE [116], the MEIJER calculus of AUSTRY & BOUDOL [10] and the Algebra of Communicating Processes (ACP) of BERGSTRÄ & KLOP [25, 26, 28].

When work on process algebra started, many people hoped that it would be possible to come up, eventually, with the ‘ultimate’ process algebra, leading to a ‘Church thesis’ for concurrent computation. This process algebra, one imagined, should contain only a few fundamental operators and it should be suited to model all concurrent computational processes. Moreover there should be a calculus for this model making it possible to prove the identity of processes algebraically, thus proving correctness of implementations with respect to specifications. As far as we know, the ultimate process algebra has not yet been found, but we will not exclude that it will be discovered in the near future.

Two things however, have become clear in the meantime: (1) it is doubtful whether algebraic system verification, as envisaged in [134], will be possible in this model, and (2) even if the ultimate process algebra exists, this certainly does not mean that all other process algebras are no longer interesting. We elaborate on this below.

A central idea in process algebra is that two processes which cannot be distinguished by observation should preferably be identified: the process semantics should be fully abstract with respect to some notion of testing (see [64, 134] and the first chapter of this thesis). This means that the choice of a suitable process algebra may depend on the tools an environment has to distinguish between certain processes. In different applications the tools of the environment may be different, and therefore different applications may require different process algebras. A large number of process semantics are not fully abstract with respect to any (reasonable) notion of testing (bisimulation semantics and partial order semantics, for instance). Still these semantics can be very interesting because they have simple definitions or correspond to some strong operational intuition. Our hypothetical ultimate process algebra will make very few identifications, because it should be resistant against all forms of testing. Therefore not many algebraic laws will be valid in this model and algebraic system verification will presumably not be possible (specification and implementation correspond to different processes in the model).

Another factor which plays a role has to do with the operators of process algebras. For theoretical purposes it is in general desirable to work with a single, small set of fundamental operators. We doubt however that a unique optimal and minimal collection exists. What is optimal depends on the type of results one likes to prove. This becomes even more clear if we look towards practical applications. Some operators in process algebra can be used for a wide range of applications, but we agree with JIFENG & HOARE [118] that we

may have to accept that each application will require derivation of specialised laws (and operators) to control its complexity.

Many people are embarrassed by the multitude of process algebras occurring in the literature. They should be aware of the fact that there are close relationships between the various process algebras: often one process algebra can be viewed as a homomorphic image, subalgebra or restriction of another one. The aim of this chapter is to show how the semantical reality, consisting of a large number of closely related process algebras, can be reflected, and even used, on the level of algebraic specifications and in process verifications.

This chapter is about process algebras, their mutual relationships, and strategies to prove that a formula is valid in a process algebra. Still, we do not present any particular process algebra here. In the other chapters of this thesis several process algebras are discussed. However we neither define all the operations we use in this chapter nor all the semantical notions that will be considered here. In this chapter we only define classes of models of process modules. One reason for doing this is that a detailed description of all particular process algebras we use would make this thesis too long. Another reason is that there is often no clear argument for selecting a particular process algebra. In such situations we are interested in assertions saying that a formula is valid in all algebras satisfying a certain theory. A number of times we need results stating that some formulas *cannot* be proven from a certain module. A standard way to prove this is to give a model of the module where the formulas are not true. For this reason we will often refer to particular process algebras which have been described elsewhere in the literature.

The discussion of this chapter takes place in the setting of ACP. We think however that the results can be carried over to CCS, CSP, MEJE, or any other process algebra formalism.

Modularisation.

The creation of an algebraic framework suitable to deal with realistic applications, gives rise to the construction of building blocks, or modules, of operators and axioms, each block describing a feature of concurrency in a certain semantical setting. These modules can then be combined by means of a module combinator $+$. We give some examples:

- i) A kernel module, that expresses some basic features of concurrent processes, is the module ACP. For a lot of applications however, ACP does not provide enough operators. Often the use of *renaming operators* makes specifications shorter and more comprehensible. These renaming operators can be defined in a separate module RN. Now the module $ACP + RN$ combines the specification and verification power of modules ACP and RN.
- ii) The axioms of module ACP correspond to the semantical notion of bisimulation. For some applications bisimulation semantics does not make enough identifications. In these cases one would like to deal with processes on the level of, for example, failure semantics. Now one can define a module F, corresponding to the identifications made in failure

semantics on top of the identifications of bisimulation semantics. The module $ACP + F$ then corresponds to the failure model.

Once a number of modules have been defined, they can be combined in a lot of ways. Some combinations are interesting (for example the module $ACP + RN + F$), for other combinations no interesting applications exist (the module $RN + F$). Didactical aspects aside, a major advantage of the modular approach is that results which have been proved from a module M , can also be proved from a module $M + N$. This means that process verifications become *reusable*.

It turns out that certain pairs of modules are incompatible in a very strong sense: with the combination of two modules strange and counter-intuitive identities can be derived. In BAETEN, BERGSTRA & KLOP [14], for example, it is shown that the combination of failure semantics and the priority operator is inconsistent in the sense that an identity can be derived which says that a particular process that can do a b -action after it has done an a -action, equals a process that cannot do this. Another example can be found in BERGSTRA, KLOP & OLDEROG [29], where it is pointed out that the combination of failure semantics and Koomen's Fair Abstraction Rule (KFAR) is inconsistent.

In the first section of this chapter we present, besides the combinator $+$, some other operators on modules. We discuss an export operator $\square$, allowing to forget some operators in a module, an operator H , changing semantics by taking homomorphic images, and an operator S which takes subalgebras. These operators enable us to combine modules in a subtle way, when the direct combination would be inconsistent. In Section 2 we describe a large number of process modules which play a role in the ACP framework. Section 3 contains two examples of applications of the new module operators in process algebra:

1. The axiom system ACP contains auxiliary operators $\parallel$ and $|$ (left-merge and communication-merge) which drastically simplify computations and have some desirable 'metamathematical' consequences (finite axiomatisability¹; greater suitability for term rewriting analysis). These auxiliary operators can be defined in a large class of process algebras. However, it turns out that in a setting with the silent step τ the left-merge cannot be added consistently to all algebras (for instance not to the usual variants of failure semantics). Now one may think that this result means that someone who is doing failure semantics with τ 's cannot profit from the nice properties of the left-merge. However, we will show in this chapter that use of the module approach makes it possible to do failure semantics with τ 's but still benefit from the left-merge in verifications. The idea is that verifications take place on two levels: the level of bisimulation semantics where the left-merge can be used, and a level of for instance failure semantics, where no left-merge is present. The failure model can be

1. Recently, FARON MOLLER [139] from Edinburgh showed that in bisimulation semantics the merge operator cannot be finitely axiomatised without auxiliary operators.

obtained from the bisimulation model by removing the auxiliary operators and taking a homomorphic image. Now we use the observation that certain formulas (the ‘positive’ ones without auxiliary operators) are preserved under this procedure. A consequence of this application is that even if bisimulation semantics is not considered to be an appropriate process semantics (since it is not fully abstract with respect to any reasonable notion of testing), it still can be useful as an expedient for proving formulas in failure semantics.

2. As already pointed out above, one would like to have, from a theoretical point of view, as few operators or combinators as possible. On the other hand, when dealing with applications, it is often very rewarding to introduce new operators. This paradox can be resolved if the new operators are definable in terms of the more elementary ones. In that case the new operators can be considered as notations which are useful, but do not complicate the underlying theory. A problem with defining operators in terms of other operators is that often auxiliary atomic actions are needed in the definition. These auxiliary actions can then not be used in any other place, because that would disturb the intended semantics of the operator. In the laws that can be derived for the defined operator, the auxiliary actions occur prominently. These ‘side effects’ are often quite unpleasant. One may think that side effects are unavoidable and that someone who really does not like them should define new operators directly in the algebras (even though this is in conflict with the desire to have as few operators as possible). However, we will show that the module approach can be used to solve also this problem: with the restriction operator we remove the auxiliary actions from the signature and then we apply the subalgebra operator in order to ‘move’ to algebras where the auxiliary actions are not present at all.

The concept of hiding auxiliary operators in a module in some formal way is quite familiar in the literature (see BERGSTRA, HEERING & KLINT [23] for example), but the use of module operators H and S , and their application in combining modules that would be incompatible otherwise, is, as far as we know, new. The H and S operations are in spirit related to the **abstract** operation of SANNELLA & WIRSING [162] and SANNELLA & TARLECKI [161], which also extends the model class of a module.

In previous papers on ACP, the underlying logic used in process verifications was not made explicit. The reason for this was that a long definition of the logic would distract the reader’s attention from the more essential parts of the paper. It was felt that filling in the details of the logic would not be too difficult and that moreover different options were equivalent. In this chapter we generalise the classical notion of a formal proof of a formula from a theory to the notion of a formal proof of a formula from a module. The definition of this last notion is parametrised by the underlying logic. What is provable from a module really depends on the logic that is used, and this makes it necessary to consider in more detail the issue of logics. In an appendix we present three alternatives: (1) Equational logic. This logic is

suited for dealing with finite processes, but not strong enough for handling infinite processes; (2) Infinitary conditional equational logic. This is the logic used in most process verifications in the ACP framework until now; (3) First order logic with equality.

Our investigations into the precise nature of the calculi used in process algebra, led us to alternative formulations of some of the proof principles in ACP which fit better in our formal setup. We present a reformulation of the Recursive Specification Principle (RSP) and also an alphabet operator which returns a process instead of a set of actions.

Queues.

As an illustration of the techniques developed in Sections 1 to 3, we present in Section 4 an algebraic treatment of FIFO-queues. FIFO-queues play an important role in the description of languages with asynchronous message passing, the modelling of communication channels occurring in computer networks and the implementation of languages with synchronous communication. We show how the chaining operator can be used to give short specifications of various (faulty) queues and simple proofs of numerous identities, for example of the fact that the chaining of a queue with unbounded capacity and a one datum buffer is again a queue.

We give an example of an identity that holds intuitively (there is no experiment that distinguishes between the two processes) but is not valid in bisimulation semantics. We use the machinery developed in Section 1-3 to extend the axiom system in a neat way (avoiding inconsistencies) so that we can prove the processes identical.

A protocol verification.

The usefulness of the proof technique for queues is illustrated in Section 5, where a modular verification is presented of a concurrent alternating bit protocol. This verification takes 4 pages (or 5 if the proof of the standard facts about the queues is included) and is thereby considerably shorter than the proof of similar protocols in papers by KOYMANS & MULDER [123] and LARSEN & MILNER [127] (15 and 11 pages respectively). The verification shows that the protocol is correct if the channels behave as faulty FIFO-queues with unbounded capacity. However, a minor change in the proof is enough to show that the protocol also works if the channels behave as n -buffers, faulty n -buffers, etc. In our view the basic merit of our way of dealing with queues is that it becomes possible to use inductive arguments when dealing with the length of queues in protocol systems.

1. MODULE LOGIC

In this chapter, as in many other papers about process algebra, we use formal calculi to prove statements about concurrent systems. In this section we answer the following questions:

- Which kind of calculi do we use?
- What do we understand by a proof?

In the next sections we will apply this general setup to the setting of concurrent systems.

1.1. Statements about concurrent systems. In many theories of concurrency it is common practice to represent processes - the behaviours of concurrent systems - as elements in an *algebra*. This is a mathematical domain, on which some operators and predicates are defined. Algebras, which are suitable for the representation of processes are called *process algebras*. Thus a statement about the behaviour of concurrent systems can be regarded as a statement about the elements of a certain process algebra. Such a statement can be represented by a formula in a suitable language which is interpreted in this process algebra. Sometimes we consider several process algebras at the same time and want to formulate a statement about concurrent processes without choosing one of these algebras. In this case we represent the statement by a formula in a suitable language which has an interpretation in all these process algebras. Hence we are interested in assertions of the form: 'Formula ϕ holds in the process algebra $\mathcal{A}$ ', notation $\mathcal{A} \models \phi$, or 'Formula ϕ holds in the class of process algebras $\mathcal{C}$ ', notation $\mathcal{C} \models \phi$. Now we can formulate the goal that is pursued in the present section: to propose a method for proving assertions $\mathcal{A} \models \phi$, or $\mathcal{C} \models \phi$.

1.2. Proving formulas from theories. Classical logic gave us the notion of a formal proof of a formula ϕ from a theory T . Here a theory is a set of formulas. We write $T \vdash \phi$ if such a proof exists. The use of this notion is revealed by the following soundness theorem: *If $T \vdash \phi$ then ϕ holds in all algebras satisfying T .* Here an algebra $\mathcal{A}$ satisfies T , notation $\mathcal{A} \models T$, if all formulas of T hold in this algebra. Thus if we want to prove $\mathcal{A} \models \phi$ it suffices to prove $T \vdash \phi$ and $\mathcal{A} \models T$ for a suitable theory T . Likewise, if we want to prove $\mathcal{C} \models \phi$, with $\mathcal{C}$ a class of algebras, it suffices to prove $T \vdash \phi$ and $\mathcal{C} \models T$.

At first sight the method of proving $\mathcal{A} \models \phi$ by means of a formal proof of ϕ out of T seems very inefficient. Instead of verifying $\mathcal{A} \models \phi$, one has to verify $\mathcal{A} \models \psi$ for all $\psi \in T$, and moreover the formal proof has to be constructed. However, there are two circumstances in which this method is efficient, and in most applications both of them apply. First of all it may be the case that ϕ is more complicated than the formulas of T and that a direct verification of $\mathcal{A} \models \phi$ is much more work than the formal proof and all verifications $\mathcal{A} \models \psi$ together. Secondly, it may occur that a single theory T with $\mathcal{A} \models T$ is used to prove many formulas ϕ , so that many verifications $\mathcal{A} \models \phi$ are balanced against many formal proofs of ϕ out of T and a single set of verifications $\mathcal{A} \models \psi$. Especially when constructing formal proofs is considered easier than making verifications $\mathcal{A} \models \phi$, this reusability argument is very powerful. It also indicates that for a given

algebra $\mathcal{A}$ we want to find a theory T from which most interesting formulas ϕ with $\mathcal{A} \models \phi$ can be proved.

Often there are reasons for representing processes in an algebra that satisfies a particular theory T , but there is no clear argument for selecting one of these algebras. In this situation we are interested in assertions $\mathcal{C} \models \phi$ with $\mathcal{C}$ the class of all algebras satisfying T . Of course assertions of this type can be conveniently proved by means of a formal proof of ϕ from T .

1.3. Proving formulas from modules. In process algebra we often want to modify the process algebra currently used to represent processes. Such a modification might be as simple as the addition of another operator, needed for the proper modelling of yet another feature of concurrency, but it can also be a more involved modification, such as factoring out a congruence, in order to identify processes that should not be distinguished in a certain application. It is our explicit concern to organise proofs of statements about concurrent systems in such a way that, whenever possible, our results carry over to modifications of the process algebra for which they were proved.

Now suppose $\mathcal{A}$ is a process algebra satisfying the theory T and a statement $\mathcal{A} \models \phi$ has been proved by means of a formal proof of ϕ out of T . Furthermore suppose that $\mathcal{B}$ is obtained from $\mathcal{A}$ by factoring out a congruence relation on $\mathcal{A}$ (so $\mathcal{B}$ is a *homomorphic image* of $\mathcal{A}$) and for a certain application $\mathcal{B}$ is considered to be a more suitable model of concurrency than $\mathcal{A}$. Then in general $\mathcal{B} \models \phi$ cannot be concluded, but if ϕ belongs to a certain class of formulas (the *positive* ones) it can. So if ϕ is positive we can use the following theorem: 'If $\mathcal{A} \models T$, $T \vdash \phi$, ϕ is positive, and $\mathcal{B}$ is a homomorphic image of $\mathcal{A}$, then $\mathcal{B} \models \phi$ '. This saves us the trouble of finding another theory U , verifying that $\mathcal{B} \models U$ and proving $U \vdash \phi$ for many formulas ϕ that have been proved from T already. Another way of formulating the same idea is to introduce a module $H(T)$. We postulate that one may derive ' $H(T) \vdash \phi$ ' from ' $T \vdash \phi$ ' and ' ϕ is positive', and $H(T) \vdash \phi$ implies that ϕ holds in all homomorphic images of algebras satisfying T .

Thus we propose a generalisation of the notion of a formal proof. Instead of theories we use the more general notion of *modules*. Like a theory a module characterises a class $\mathcal{C}$ of algebras, but besides the class of all algebras satisfying a given set of formulas, $\mathcal{C}$ can for instance also be the class of homomorphic images or subalgebras of a class of algebras specified earlier. Now a proof in the framework of module algebra is a sequence or tree of assertions $M \vdash \phi$ such that in each step either the formula ϕ is manipulated, as in classical proofs, or the module M is manipulated. Of course we will establish a soundness theorem as before, and then an assertion $\mathcal{A} \models \phi$ can be proved by means of a module M with $\mathcal{A} \models M$ and a formal proof of ϕ out of M . We will now turn to the formal definitions.

1.4. Signatures. Let NAMES be a given set of names.

A *sort declaration* is an expression $\mathbb{S}:S$ with $S \in \text{NAMES}$.

A *function declaration* is an expression $\mathbb{F}:f:S_1 \times \cdots \times S_n \rightarrow S$ with $f, S_1, \dots, S_n, S \in \text{NAMES}$.

A *predicate declaration* is an expression $\mathbb{R}:p \subseteq S_1 \times \cdots \times S_n$ with $p, S_1, \dots, S_n \in \text{NAMES}$.

A *signature* σ is a set of sort, function and predicate declarations, satisfying:

$$(\mathbb{F}:f:S_1 \times \cdots \times S_n \rightarrow S) \in \sigma \Rightarrow (\mathbb{S}:S_i) \in \sigma \ (i=1, \dots, n) \wedge (\mathbb{S}:S) \in \sigma$$

$$(\mathbb{R}:p \subseteq S_1 \times \cdots \times S_n) \in \sigma \Rightarrow (\mathbb{S}:S_i) \in \sigma \ (i=1, \dots, n)$$

A function declaration $\mathbb{F}:f \rightarrow S$ of arity 0 is sometimes called a *constant declaration* and written as $\mathbb{F}:f \in S$.

1.5. σ -Algebras. Let σ be a signature. A σ -*algebra* $\mathcal{A}$ is a function on σ that maps

$(\mathbb{S}:S) \in \sigma$ to a set $S^{\mathcal{A}}$,

$(\mathbb{F}:f:S_1 \times \cdots \times S_n \rightarrow S) \in \sigma$ to a function $f_{S_1 \times \cdots \times S_n \rightarrow S}^{\mathcal{A}}:S_1^{\mathcal{A}} \times \cdots \times S_n^{\mathcal{A}} \rightarrow S^{\mathcal{A}}$,

$(\mathbb{R}:p \subseteq S_1 \times \cdots \times S_n) \in \sigma$ to a predicate $p_{S_1 \times \cdots \times S_n}^{\mathcal{A}} \subseteq S_1^{\mathcal{A}} \times \cdots \times S_n^{\mathcal{A}}$.

Let $\mathcal{A}$ and $\mathcal{B}$ be σ -algebras. $\mathcal{B}$ is a *subalgebra* of $\mathcal{A}$ if $S^{\mathcal{B}} \subseteq S^{\mathcal{A}}$ for all $(\mathbb{S}:S) \in \sigma$, if moreover $f_{S_1 \times \cdots \times S_n \rightarrow S}^{\mathcal{A}}$ restricted to $S_1^{\mathcal{B}} \times \cdots \times S_n^{\mathcal{B}} \rightarrow S^{\mathcal{B}}$ is just $f_{S_1 \times \cdots \times S_n \rightarrow S}^{\mathcal{B}}$ for all $\mathbb{F}:f:S_1 \times \cdots \times S_n \rightarrow S$ in σ , and if $p_{S_1 \times \cdots \times S_n}^{\mathcal{A}}$ restricted to $S_1^{\mathcal{B}} \times \cdots \times S_n^{\mathcal{B}}$ is just $p_{S_1 \times \cdots \times S_n}^{\mathcal{B}}$ for all $\mathbb{R}:p \subseteq S_1 \times \cdots \times S_n$ in σ .

A *homomorphism* $h:\mathcal{A} \rightarrow \mathcal{B}$ consists of mappings $h_S:S^{\mathcal{A}} \rightarrow S^{\mathcal{B}}$ for all $\mathbb{S}:S$ in σ , such that

$$h_S(f_{S_1 \times \cdots \times S_n \rightarrow S}^{\mathcal{A}}(x_1, \dots, x_n)) = f_{S_1 \times \cdots \times S_n \rightarrow S}^{\mathcal{B}}(h_{S_1}(x_1), \dots, h_{S_n}(x_n))$$

$$\text{for all } (\mathbb{F}:f:S_1 \times \cdots \times S_n \rightarrow S) \in \sigma \text{ and all } x_i \in S_i^{\mathcal{A}} (i=1, \dots, n)$$

$$p_{S_1 \times \cdots \times S_n}^{\mathcal{A}}(x_1, \dots, x_n) \Leftrightarrow p_{S_1 \times \cdots \times S_n}^{\mathcal{B}}(h_{S_1}(x_1), \dots, h_{S_n}(x_n))$$

$$\text{for all } (\mathbb{R}:p \subseteq S_1 \times \cdots \times S_n) \in \sigma \text{ and all } x_i \in S_i^{\mathcal{A}} (i=1, \dots, n)$$

$\mathcal{B}$ is a *homomorphic image* of $\mathcal{A}$ if there exists a surjective homomorphism $h:\mathcal{A} \rightarrow \mathcal{B}$.

Let $\mathcal{A}$ be a σ -algebra. The *restriction* $\rho \sqcap \mathcal{A}$ of $\mathcal{A}$ to the signature ρ is the $\rho \cap \sigma$ -algebra $\mathcal{B}$, defined by

$$S^{\mathcal{B}} = S^{\mathcal{A}} \text{ for all } (\mathbb{S}:S) \in \rho \cap \sigma$$

$$f_{S_1 \times \cdots \times S_n \rightarrow S}^{\mathcal{B}} = f_{S_1 \times \cdots \times S_n \rightarrow S}^{\mathcal{A}} \text{ for all } (\mathbb{F}:f:S_1 \times \cdots \times S_n \rightarrow S) \in \rho \cap \sigma$$

$$p_{S_1 \times \cdots \times S_n}^{\mathcal{B}} = p_{S_1 \times \cdots \times S_n}^{\mathcal{A}} \text{ for all } (\mathbb{R}:p \subseteq S_1 \times \cdots \times S_n) \in \rho \cap \sigma$$

1.6. *Logics.* A logic $\mathcal{L}$ is a complex of prescriptions, defining for any signature σ

- a set $F_\sigma^\mathcal{L}$ of *formulas* over σ such that $F_\sigma^\mathcal{L} \cap F_\rho^\mathcal{L} = F_{\sigma \cap \rho}^\mathcal{L}$,
- a binary relation $\models_\sigma^\mathcal{L}$ on σ -algebras $\times F_\sigma^\mathcal{L}$ such that for all ρ -algebras $\mathcal{Q}$ and $\phi \in F_{\sigma \cap \rho}^\mathcal{L}$: $\sigma \sqsubseteq \mathcal{Q} \models_\sigma^\mathcal{L} \phi \iff \mathcal{Q} \models_\rho^\mathcal{L} \phi$
- and a set $I_\sigma^\mathcal{L}$ of *inference rules* $\frac{H}{\phi}$ with $H \subseteq F_\sigma^\mathcal{L}$ and $\phi \in F_\sigma^\mathcal{L}$.

If $\mathcal{Q} \models_\sigma^\mathcal{L} \phi$ we say that the σ -algebra $\mathcal{Q}$ *satisfies* the formula ϕ , or that ϕ *holds* in $\mathcal{Q}$. A *theory* over σ is a set of formulas over σ . If T is a theory over σ and $\mathcal{Q} \models_\sigma^\mathcal{L} \phi$ for all $\phi \in T$ we say that $\mathcal{Q}$ *satisfies* T , notation $\mathcal{Q} \models_\sigma^\mathcal{L} T$. We also say that $\mathcal{Q}$ is a *model* of T .

A logic $\mathcal{L}$ is *sound* if $\frac{H}{\phi} \in I_\sigma^\mathcal{L}$ implies $\mathcal{Q} \models_\sigma^\mathcal{L} H \implies \mathcal{Q} \models_\sigma^\mathcal{L} \phi$ for any σ -algebra $\mathcal{Q}$.

A formula $\phi \in F_\sigma^\mathcal{L}$ is *preserved under subalgebras* if $\mathcal{Q} \models_\sigma^\mathcal{L} \phi$ implies $\mathcal{B} \models_\sigma^\mathcal{L} \phi$, for any subalgebra $\mathcal{B}$ of $\mathcal{Q}$.

A formula $\phi \in F_\sigma^\mathcal{L}$ is *preserved under homomorphisms* if $\mathcal{Q} \models_\sigma^\mathcal{L} \phi$ implies $\mathcal{B} \models_\sigma^\mathcal{L} \phi$, for any homomorphic image $\mathcal{B}$ of $\mathcal{Q}$.

Without doubt, the definition of a ‘logic’ as presented above is too general for most applications. However, it is suited for our purposes and anyone can substitute his/her favourite (and more restricted) definition whenever he/she likes.

In the process algebra verifications of this chapter we will use infinitary conditional equational logic. The definition of this logic can be found in the appendix. For comparison, the definitions of equational logic and first order logic with equality are included too.

1.7. Classical logic.

DERIVABILITY. A σ -*proof* of a formula $\phi \in F_\sigma^\mathcal{L}$ from a theory $T \subseteq F_\sigma^\mathcal{L}$ using the logic $\mathcal{L}$, is a well-founded, upwardly branching tree of which the nodes are labelled by σ -formulas, such that

- the root is labelled by ϕ
- and if ψ is the label of a node q and H is the set of labels of the nodes directly above q then
 - either $\psi \in T$ and $H = \emptyset$,
 - or $\frac{H}{\psi} \in I_\sigma^\mathcal{L}$.

If a σ -proof of ϕ from T using $\mathcal{L}$ exists, we say that ϕ is σ -*provable* from T by means of $\mathcal{L}$, notation $T \vdash_\sigma^\mathcal{L} \phi$.

TRUTH. Let $\mathcal{C}$ be a class of σ -algebras and $\phi \in F_\sigma^\mathcal{L}$. Then ϕ is said to be *true* in $\mathcal{C}$, notation $\mathcal{C} \models_\sigma^\mathcal{L} \phi$, if ϕ holds in all σ -algebras $\mathcal{Q} \in \mathcal{C}$. Let $\text{Alg}(\sigma, T)$ be the class of all σ -algebras satisfying T .

SOUNDNESS THEOREM. *If $\mathcal{L}$ is sound then $T \vdash_\sigma^\mathcal{L} \phi$ implies $\text{Alg}(\sigma, T) \models_\sigma^\mathcal{L} \phi$.*

PROOF. Straightforward with induction. □

If no confusion is likely to result, the sub- and superscripts of $\models$ and $\vdash$ may be dropped without further warning.

1.8. Module logic. The set $\mathfrak{M}$ of modules is defined inductively as follows:

- If σ is a signature and T a theory over σ , then $(\sigma, T) \in \mathfrak{M}$,
- If M and $N \in \mathfrak{M}$ then $M + N \in \mathfrak{M}$,
- If σ is a signature and $M \in \mathfrak{M}$ then $\sigma \sqcap M \in \mathfrak{M}$,
- If $M \in \mathfrak{M}$ then $H(M) \in \mathfrak{M}$,
- If $M \in \mathfrak{M}$ then $S(M) \in \mathfrak{M}$.

Here $+$ is the composition operator, allowing to organise specifications in a modular way, and $\sqcap$ is the export operator, restricting the visible signature of a module, thereby hiding auxiliary items. These operators occur in some form or other frequently in the literature on software engineering. Our notation is taken from BERGSTRA, HEERING & KLINT [23] in which also additional references can be found. The homomorphism operator H and the subalgebra operator S are, as far as we know, new in the context of algebraic specifications. Of course they are well known in model theory, see for instance MONK [140].

The *visible signature* $\Sigma(M)$ of a module M is defined inductively by:

- $\Sigma(\sigma, T) = \sigma$,
- $\Sigma(M + N) = \Sigma(M) \cup \Sigma(N)$,
- $\Sigma(\sigma \sqcap M) = \sigma \cap \Sigma(M)$,
- $\Sigma(H(M)) = \Sigma(M)$,
- $\Sigma(S(M)) = \Sigma(M)$.

TRUTH. The class $Alg(M)$ of models of a module M is defined inductively by:

- $\mathcal{A}$ is a model of (σ, T) if it is a σ -algebra, satisfying T ;
- $\mathcal{A}$ is a model of $M + N$ if it is a $\Sigma(M + N)$ -algebra, such that $\Sigma(M) \sqcap \mathcal{A}$ is a model of M and $\Sigma(N) \sqcap \mathcal{A}$ is a model of N ;
- $\mathcal{A}$ is a model of $\sigma \sqcap M$ if it is the restriction of a model $\mathfrak{B}$ of M to the signature σ ;
- $\mathcal{A}$ is a model of $H(M)$ if it is a homomorphic image of a model $\mathfrak{B}$ of M ;
- $\mathcal{A}$ is a model of $S(M)$ if it is a subalgebra of a model $\mathfrak{B}$ of M .

Note that $Alg(M)$ is a generalisation of $Alg(\sigma, T)$ as defined earlier. All the elements of $Alg(M)$ are $\Sigma(M)$ -algebras. A $\Sigma(M)$ -algebra $\mathcal{A} \in Alg(M)$ is said to *satisfy* M . A formula $\phi \in F_{\Sigma(M)}^c$ is *satisfied* by a module M , notation $M \models^c \phi$, if $Alg(M) \models_{\Sigma(M)}^c \phi$, thus if ϕ holds in all $\Sigma(M)$ -algebras satisfying M .

DERIVABILITY. A *proof* of a formula $\phi \in F_{\Sigma(M)}^c$ from a module M using the logic $\mathcal{L}$, is a well-founded, upwardly branching tree of which the nodes are labelled by assertions $N \vdash \psi$, such that

- the root is labelled by $M \vdash \phi$
- if $N \vdash \psi$ is the label of a node q and H is the set of labels of the nodes directly above q then $\frac{H}{N \vdash \psi}$ is one of the inference rules of Table 1.

$(\sigma, T) \vdash \phi$	if $\phi \in T$
$\frac{M \vdash \phi_j \ (j \in J)}{M \vdash \phi}$	whenever $\frac{\phi_j \ (j \in J)}{\phi} \in I_{\Sigma(M)}^{\mathcal{L}}$
$\frac{M \vdash \phi}{M + N \vdash \phi}$	$\frac{N \vdash \phi}{M + N \vdash \phi}$
$\frac{M \vdash \phi}{\sigma \Box M \vdash \phi}$	if $\phi \in F_{\sigma}^{\mathcal{L}}$
$\frac{M \vdash \phi}{H(M) \vdash \phi}$	if ϕ is <i>positive</i>
$\frac{M \vdash \phi}{S(M) \vdash \phi}$	if ϕ is <i>universal</i>

TABLE 1

Here *positive* and *universal* are syntactic criteria, to be defined for each logic $\mathcal{L}$ separately, ensuring that a formula is preserved under homomorphisms and subalgebras respectively. We write $N \vdash \psi$ for $\frac{\emptyset}{N \vdash \psi}$, and omit braces in the conditions of inference rules. If a proof of ϕ from M using $\mathcal{L}$ exists, we say that ϕ is *provable* from M by means of $\mathcal{L}$, notation $M \vdash^{\mathcal{L}} \phi$.

LEMMA. *If $M \vdash^{\mathcal{L}} \phi$ then $\phi \in F_{\Sigma(M)}^{\mathcal{L}}$.*

PROOF. With induction. The only nontrivial cases are the rules for $+$ and $\Box$. These follow from $F_{\sigma}^{\mathcal{L}} \subseteq F_{\sigma \cup \rho}^{\mathcal{L}}$ and $F_{\sigma}^{\mathcal{L}} \cap F_{\rho}^{\mathcal{L}} \subseteq F_{\sigma \cap \rho}^{\mathcal{L}}$ respectively. $\square$

SOUNDNESS THEOREM. *If $\mathcal{L}$ is sound then $M \vdash^{\mathcal{L}} \phi$ implies $M \models^{\mathcal{L}} \phi$.*

PROOF. With induction. Again the only nontrivial cases are the rules for $+$ and $\Box$. These follow since for all ρ -algebras $\mathcal{Q}$ and $\phi \in F_{\sigma \cap \rho}^{\mathcal{L}}$: $\sigma \Box \mathcal{Q} \models \phi \Rightarrow \mathcal{Q} \models \phi$ and $\sigma \Box \mathcal{Q} \models \phi \Leftarrow \mathcal{Q} \models \phi$ respectively. $\square$

2. PROCESS ALGEBRA

This thesis does not contain an introductory chapter on process algebra. We only give a listing of some important process modules. For an introduction to the ACP formalism we refer the reader to [25, 26, 28] and [17].

2.1. ACP_τ . In this chapter a central role will be played by the module ACP_τ , the Algebra of Communicating Processes with abstraction. A first parameter of ACP_τ is a finite set A of *actions*. For each action $a \in A$ there is a constant a in the language, representing the process, starting with an a -action and terminating (successfully) after some time.

The first two composition operations we consider are $;$, denoting *sequential composition*, and $+$ for *alternative composition*. If x and y are two processes, then $x;y$ is the process that starts execution of y after successful completion of x , and $x+y$ is the process that either behaves like x or like y . We do not specify whether the choice between x and y is made by the process itself, or by the environment.

We have a special constant δ , denoting *deadlock*, *inaction*, a process that cannot do anything at all. In particular δ does not terminate successfully. We write $A_\delta = A \cup \{\delta\}$.

Next we have a *parallel composition* operator $\parallel$. $x \parallel y$ denotes the process corresponding to the parallel execution of x and y . Execution of $x \parallel y$ either starts with a step from x , or with a step from y , or with a *synchronisation* of an action from x and an action from y . Synchronisation of actions is described by the second parameter of ACP_τ , which is a binary communication function $\gamma: A_\delta \times A_\delta \rightarrow A_\delta$ that is commutative, associative and has δ as zero element:

$$\gamma(a,b) = \gamma(b,a) \quad \gamma(a,\gamma(b,c)) = \gamma(\gamma(a,b),c) \quad \gamma(a,\delta) = \delta$$

If $\gamma(a,b) = c \neq \delta$ this means that actions a and b can synchronise. The synchronous performance of a and b is then regarded as a performance of the communication action c . Formally we should add the parameters to the name of a module: $ACP_\tau(A, \gamma)$. However, in order to keep notation simple, we will always omit the parameters if this can be done without causing confusion. In order to axiomatise the $\parallel$ -operator we use two auxiliary operators $\sqcup$ (*left-merge*) and $\mid$ (*communication merge*). $x \sqcup y$ is $x \parallel y$, but with the restriction that the first step comes from x , and $x \mid y$ is $x \parallel y$ but with a synchronisation action as the first step.

Next we have for each $H \subseteq A$ an *encapsulation* operator ∂_H . The operator ∂_H blocks actions from H . The operator is used to encapsulate a process, i.e. to block synchronisation with the environment.

When describing concurrent systems and reasoning about their behaviour, it is often useful to have a distinguished action that cannot synchronise with any other action. Such an action is denoted by the constant $\tau \notin A_\delta$. The fact that τ cannot synchronise makes that in some sense this action is not observable. Therefore it is often called the *silent* action. For each $I \subseteq A$ the language contains an *abstraction* or *hiding* operator τ_I . This operator hides actions in I by renaming them into τ , thus expressing that certain actions in a system behaviour cannot be observed.

In Table 2 we summarize the signature of module ACP_τ .

S (sort):	P	the set of processes
F (functions):	$+$:	$P \times P \rightarrow P$ alternative composition (sum)
	$\cdot$:	$P \times P \rightarrow P$ sequential composition (product)
	$\parallel$:	$P \times P \rightarrow P$ parallel composition (merge)
	$\sqcup$:	$P \times P \rightarrow P$ left-merge
	$\mid$:	$P \times P \rightarrow P$ communication-merge
	∂_H :	$P \rightarrow P$ encapsulation, for any $H \subseteq A$
	τ_I :	$P \rightarrow P$ abstraction, for any $I \subseteq A$
	a	$\in P$ for any atomic action $a \in A$
	δ	$\in P$ inaction, deadlock
	τ	$\in P$ silent action

TABLE 2

Table 3 contains the theory of the module ACP_τ . In this chapter we present ACP_τ as a monolithic module. In [28] however, it is shown that ACP_τ can be viewed as the sum of a large number of sub-modules which are interesting in their own right. The module consisting of axioms A1-5 only is called BPA (from Basic Process Algebra). If we add axioms A6-7 we obtain BPA_δ , and BPA_δ plus axioms T1-3 gives $BPA_{\tau,\delta}$. The module ACP consists of the axioms A1-7, CF, CM1-9 and D1-4, i.e. the left column of Table 3. All axioms in Table 3 are in fact axiom schemes in a , b , H and I . Here a and b range over A_δ (unless further restrictions are made in the table) and $H, I \subseteq A$. In a product $x \cdot y$ we will often omit the $\cdot$. We take $\cdot$ to be more binding than other operations and $+$ to be less binding than other operations. In case we are dealing with an associative operator, we also leave out parentheses.

2.1.1. Note. Let $n > 0$. Let $D = \{d_1, \dots, d_n\}$ be a finite set. Let $t_{d_1}, \dots, t_{d_n}$ be process expressions. We use the notation $\sum_{d \in D} t_d$ for the sum $t_{d_1} + \dots + t_{d_n}$.

$\sum_{d \in \emptyset} t_d = \delta$ by definition.

2.1.2. Summand inclusion. In process verifications the summand inclusion predicate $\subseteq$ turns out to be a useful notation. It is defined by: $x \subseteq y \Leftrightarrow x + y = y$. From the ACP_τ -axioms A1, A2 and A3 respectively it follows that $\subseteq$ is antisymmetrical, transitive and reflexive, and hence a partial order.

2.1.3. PROPOSITION. $ACP_\tau \vdash \tau x \parallel y = \tau(x \parallel y)$.

PROOF. $\tau x \parallel y \supseteq \tau x \sqcup y = \tau(x \parallel y) = \tau x \sqcup y = \tau \tau x \sqcup y = \tau(\tau x \parallel y) \supseteq \tau x \parallel y$.

Now use the fact that $\subseteq$ is a partial order. □

$x + y = y + x$	A1	$x\tau = x$	T1
$x + (y + z) = (x + y) + z$	A2	$\tau x + x = \tau x$	T2
$x + x = x$	A3	$a(\tau x + y) = a(\tau x + y) + ax$	T3
$(x + y)z = xz + yz$	A4		
$(xy)z = x(yz)$	A5		
$x + \delta = x$	A6		
$\delta x = \delta$	A7		
$a b = \gamma(a, b)$	CF		
$x y = x _y + y _x + x y$	CM1	$\tau _x = \tau x$	TM1
$a _x = ax$	CM2	$(\tau x) _y = \tau(x y)$	TM2
$(ax) _y = a(x y)$	CM3	$\tau x = \delta$	TC1
$(x + y) _z = x _z + y _z$	CM4	$x \tau = \delta$	TC2
$(ax) b = (a b)x$	CM5	$(\tau x) y = x y$	TC3
$a (bx) = (a b)x$	CM6	$x (\tau y) = x y$	TC4
$(ax) (by) = (a b)(x y)$	CM7		
$(x + y) z = x z + y z$	CM8		
$x (y + z) = x y + x z$	CM9		
		$\partial_H(\tau) = \tau$	DT
		$\tau_I(\tau) = \tau$	TI1
$\partial_H(a) = a \quad \text{if } a \notin H$	D1	$\tau_I(a) = a \quad \text{if } a \notin I$	TI2
$\partial_H(a) = \delta \quad \text{if } a \in H$	D2	$\tau_I(a) = \tau \quad \text{if } a \in I$	TI3
$\partial_H(x + y) = \partial_H(x) + \partial_H(y)$	D3	$\tau_I(x + y) = \tau_I(x) + \tau_I(y)$	TI4
$\partial_H(xy) = \partial_H(x) \cdot \partial_H(y)$	D4	$\tau_I(xy) = \tau_I(x) \cdot \tau_I(y)$	TI5

TABLE 3

2.1.4. Monotony. Most of the operators of ACP_τ are monotonous with respect to the summand inclusion ordering. Using essentially the distributivity of the operators over $+$, one can show that if $x \subseteq y$, ACP_τ proves:

- $x + z \subseteq y + z$,
- $x \cdot z \subseteq y \cdot z$,
- $x||_z \subseteq y||_z$,
- $x|z \subseteq y|z$,
- $\partial_H(x) \subseteq \partial_H(y)$,
- $\tau_I(x) \subseteq \tau_I(y)$.

Due to branching time, in general $z \cdot x \not\subseteq z \cdot y$, $x||z \not\subseteq y||z$ and $z||_x \not\subseteq z||_y$. However, we do have monotony of the merge for the case where x is of the form $\tau x'$. If $\tau x' \subseteq y$, then $ACP_\tau \vdash \tau x' || z \subseteq y || z$:

$$\stackrel{2.1.3}{\tau x' || z} = \tau(x' || z) = \tau x' ||_z \subseteq y ||_z \subseteq y || z.$$

2.2. *Standard Concurrency.* Often one adds to ACP_τ the following module SC of Standard Concurrency ($a \in A_\delta$), which is parametrised by A . A proof that these axioms hold for all closed recursion-free terms can be found in [26].

SC	$(x \parallel y) \parallel z = x \parallel (y \parallel z)$	SC1
	$(x \mid ay) \parallel z = x \mid (ay \parallel z)$	SC2
	$x \mid y = y \mid x$	SC3
	$x \parallel y = y \parallel x$	SC4
	$x \mid (y \mid z) = (x \mid y) \mid z$	SC5
	$x \parallel (y \parallel z) = (x \parallel y) \parallel z$	SC6

TABLE 4

2.3. *Renamings.* Let $A_{\tau\delta} = A_\delta \cup \{\tau\}$. For every function $f: A_{\tau\delta} \rightarrow A_{\tau\delta}$ with the property that $f(\delta) = \delta$ and $f(\tau) = \tau$, we introduce an operator $\rho_f: P \rightarrow P$. Axioms for ρ_f are given in Table 5 (Here $a \in A_{\tau\delta}$ and id is the identity). Module RN is parametrised by A .

RN	$\rho_f(a) = f(a)$	RN1
	$\rho_f(x + y) = \rho_f(x) + \rho_f(y)$	RN2
	$\rho_f(xy) = \rho_f(x) \cdot \rho_f(y)$	RN3
	$\rho_{id}(x) = x$	RN4
	$\rho_f \circ \rho_g(x) = \rho_{f \circ g}(x)$	RN5

TABLE 5

For $t \in A_{\tau\delta}$ and $H \subseteq A$ we define mappings $r_{t,H}: A_{\tau\delta} \rightarrow A_{\tau\delta}$ as follows:

$$r_{t,H}(a) = \begin{cases} t & \text{if } a \in H \\ a & \text{otherwise} \end{cases}$$

In the following we will implicitly identify the operators ∂_H and $\rho_{r_{\delta,H}}$, and also the operators τ_I and $\rho_{r_{\tau,I}}$: encapsulation is just renaming of actions into δ , and abstraction is renaming of actions into the silent step τ .

2.4. *Chaining operators.* A basic situation we will encounter is one in which processes input and output values in a domain D . Often we want to ‘chain’ two processes in such a way that the output of the first one becomes the input of the second. In order to describe this, we define *chaining* operators $\ggg$ and $\gg$. In the process $x \ggg y$ the output of process x serves as input of process y . Operator $\gg$ is identical to operator $\ggg$, but hides in addition the communications that take place at the internal communication port. The reason for introducing two operators is a technical one: the operator $\gg$ (in which we are

interested most) often leads to the possibility of an infinite sequence of internal actions corresponding to hidden synchronisations between the two arguments of the operator (a form of *unguarded recursion*, cf. Sections 2.8.1 and 2.12.1). In order to deal with such behaviours, it is useful to view $\gg$ as the composition of two operators: the $\ggg$ operator and an abstraction operator that hides the communications of $\ggg$. We will define the chaining operators in terms of the operators of $ACP_\tau + RN$. In this way we obtain a simple, finite axiomatisation of the operators. The operator $\gg$ occurs (in a different notation) already in HOARE [115] and MILNER [134].

Let for $d \in D$, $\downarrow d$ be the action of reading d , and $\uparrow d$ be the action of sending d . Furthermore let $ch(D)$ be the following set:

$$ch(D) = \{\uparrow d, \downarrow d, s(d), r(d), c(d) \mid d \in D\}.$$

Here $r(d)$, $s(d)$ and $c(d)$ ($d \in D$) are auxiliary actions which play a role in the definition of the chaining operators. The module for the chaining operators is parametrised by an action alphabet A satisfying $ch(D) \subseteq A$. The module should occur in a context with a module $ACP_\tau(A, \gamma)$ where

$$range(\gamma) \cap \{\downarrow d, \uparrow d, s(d), r(d) \mid d \in D\} = \emptyset$$

and communication on $ch(D)$ is defined by

$$\gamma(s(d), r(d)) = c(d)$$

(all other communications give δ). The renaming functions $\uparrow s$ and $\downarrow r$ are defined by

$$\uparrow s(\uparrow d) = s(d) \quad \text{and} \quad \downarrow r(\downarrow d) = r(d) \quad (d \in D)$$

and $\uparrow s(a) = \downarrow r(a) = a$ for every other $a \in A_{\tau\delta}$. Now the ‘concrete’ chaining of processes x and y , notation $x \ggg y$, is defined by means of the axiom ($H = \{s(d), r(d) \mid d \in D\}$):

$$x \ggg y = \partial_H(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(y)) \quad \text{CH1}$$

The ‘abstract’ chaining of processes x and y , notation $x \gg y$, is defined by means of the axiom ($I = \{c(d) \mid d \in D\}$):

$$x \gg y = \tau_I(x \ggg y) \quad \text{CH2}$$

The module CH^+ consists of axioms CH1 and CH2, and is parametrised by A . The ‘+’ in CH^+ refers to the auxiliary actions in the module, which will be removed in Section 3.

2.4.1. EXAMPLE. Let $D = \{0, 1\}$. Process *AND* reads two bits and then outputs 1 if both are 1, and 0 otherwise:

$$AND = \downarrow 0 \cdot (\downarrow 0 \cdot \uparrow 0 + \downarrow 1 \cdot \uparrow 0) + \downarrow 1 \cdot (\downarrow 0 \cdot \uparrow 0 + \downarrow 1 \cdot \uparrow 1)$$

Process *OR* reads two bits, outputs 0 if both are 0, and 1 otherwise:

$$OR = \downarrow 0 \cdot (\downarrow 0 \cdot \uparrow 0 + \downarrow 1 \cdot \uparrow 1) + \downarrow 1 \cdot (\downarrow 0 \cdot \uparrow 1 + \downarrow 1 \cdot \uparrow 1)$$

Process *NEG* reads a bit *b* and outputs $1-b$:

$$NEG = \downarrow 0 \cdot \uparrow 1 + \downarrow 1 \cdot \uparrow 0$$

These processes can be composed using chaining operators. It is not too hard to prove:

$$(NEG \cdot NEG \gg AND) \gg NEG = OR$$

Note however that we do not have

$$(NEG \cdot NEG \ggg AND) \ggg NEG = OR$$

since in the LHS process internal computation steps are still visible.

2.5. Recursion. A *recursive specification* *E* is a set of equations $\{x = t_x \mid x \in V_E\}$ with V_E a set of variables and t_x a process expression for $x \in V_E$. Only the variables of V_E may appear in t_x . A solution of *E* is an interpretation of the variables of V_E as processes (in a certain domain), such that the equations of *E* are satisfied. Recursive specifications are used to define (or specify) infinite processes.

For each recursive specification *E* and $x \in V_E$, the module REC introduces a constant $\langle x \mid E \rangle$, denoting the *x*-component of a solution of *E*.

In most applications the variables $X \in V_E$ in a recursive specification *E* will be chosen fresh, so that there is no need to repeat *E* in each occurrence of $\langle X \mid E \rangle$. Therefore the convention will be adopted that once a recursive specification has been declared, $\langle X \mid E \rangle$ can be abbreviated by *X*. If this is done, *X* is called a *formal variable*. Formal variables are denoted by capital letters. So after the declaration $X = aX$, a statement $X = aaX$ should be interpreted as an abbreviation of $\langle X \mid X = aX \rangle = aa \langle X \mid X = aX \rangle$.

Let $E = \{x = t_x \mid x \in V_E\}$ be a recursive specification, and *t* a process expression. Then $\langle t \mid E \rangle$ denotes the term *t* in which each free occurrence of $x \in V_E$ is replaced by $\langle x \mid E \rangle$. In a recursive language we have for each *E* as above and $x \in V_E$ an axiom

$\langle x \mid E \rangle = \langle t_x \mid E \rangle$	REC
---	-----

If the above convention is used, these formulas seem to be just the equations of *E*. The module REC is parametrised by the signature in which the recursive equations are written. In the presence of module REC each system of recursion equations over this signature has a solution.

2.6. Projection. The operator $\pi_n : P \rightarrow P$ ($n \in \mathbb{N}$) *stops* processes after they have performed *n* atomic actions, with the understanding that τ -steps are transparent. The axioms for π_n are given in Table 6. Module PR is parametrised by *A*.

PR	$\pi_n(\tau) = \tau$	PR1
	$\pi_0(ax) = \delta$	PR2
	$\pi_{n+1}(ax) = a \cdot \pi_n(x)$	PR3
	$\pi_n(\tau x) = \tau \cdot \pi_n(x)$	PR4
	$\pi_n(x + y) = \pi_n(x) + \pi_n(y)$	PR5

TABLE 6

In this chapter, as in other papers on process algebra, we have an infinite collection of unary projection operators. Another option, which we do not pursue here, but which might be more fruitful if one is interested in finitary process algebra proofs, is to introduce a single binary projection operator $\mathbb{F} : \pi : \mathbb{N} \times P \rightarrow P$.

2.7. Boundedness. The predicate $B_n \subseteq P$ ($n \in \mathbb{N}$) states that the nondeterminism displayed by a process before its n^{th} atomic steps is bounded. If for all $n \in \mathbb{N}$: $B_n(x)$, we say x is bounded. Axioms for B_n are in Table 7 ($a \in A_\delta$). Module B is parametrised by A .

B	$B_0(x)$	B1
	$B_n(\tau)$	B2
	$\frac{B_n(x)}{B_n(\tau x)}$	B3
	$\frac{B_n(x)}{B_{n+1}(ax)}$	B4
	$\frac{B_n(x), B_n(y)}{B_n(x + y)}$	B5

TABLE 7

Boundedness predicates were introduced in [77].

2.8. Approximation Induction Principle. AIP^- is a proof rule which is vital if we want to prove things about infinite processes. The rule expresses the idea that if two processes are equal to any depth, and one of them is bounded then they are equal.

(AIP ⁻)	$\frac{\forall n \in \mathbb{N} \quad \pi_n(x) = \pi_n(y), B_n(x)}{x = y}$
---------------------	--

The ‘-’ in AIP^- , distinguishes the rule from a variant without predicates B_n .

2.8.1. DEFINITION. Let t be an open ACP_τ -term without abstraction operators. An occurrence of a variable X in t is *guarded* if t has a subterm of the form $a \cdot M$, with $a \in A_\delta$, and this X occurs in M . Otherwise, the occurrence is *unguarded*.

Let $E = \{x = t_x \mid x \in V_E\}$ be a recursive specification in which all t_x are ACP_τ -terms without abstraction operators. For $X, Y \in V_E$ we define:

$$X \xrightarrow{u} Y \Leftrightarrow Y \text{ occurs unguarded in } t_X.$$

We call E *guarded* if relation $\xrightarrow{u}$ is well-founded (i.e. there is no infinite sequence $X \xrightarrow{u} Y \xrightarrow{u} Z \xrightarrow{u} \dots$).

2.8.2. THEOREM (*Recursive Specification Principle (RSP)*).

$ACP_\tau + REC + PR + B + AIP^- \vdash$

$(RSP) \quad \frac{E}{x = \langle x \mid E \rangle} \quad E \text{ guarded}$
--

In plain English the RSP rule says that every guarded recursive specification has at most one solution.

2.8.3. EXAMPLE. Let $E = \{X = (a + b) \cdot X\}$ and $F = \{Y = a \cdot (a + b) \cdot Y + b \cdot Y\}$ be two recursive specifications. Since

$$\begin{aligned} \langle X \mid E \rangle &= (a + b) \cdot \langle X \mid E \rangle = a \cdot \langle X \mid E \rangle + b \cdot \langle X \mid E \rangle = \\ &= a \cdot (a + b) \cdot \langle X \mid E \rangle + b \cdot \langle X \mid E \rangle, \end{aligned}$$

the constant $\langle X \mid E \rangle$ satisfies the equation of F . Because the specification F is guarded, RSP now gives that $\langle X \mid E \rangle = \langle Y \mid F \rangle$.

2.9. *Koomen's Fair Abstraction Rule (KFAR)*. In the verification of communication protocols one often uses the following rule, called Koomen's Fair Abstraction Rule ($I \subseteq A$). Module KFAR is parametrised by A .

$(KFAR) \quad \frac{x = ix + y \quad (i \in I)}{\tau_I(x) = \tau \tau_I(y)}$
--

Fair abstraction here means that $\tau_I(x)$ will eventually exit the hidden i -cycle. Below we will formulate a generalisation of KFAR, the Cluster Fair Abstraction Rule (CFAR), which can be derived from KFAR.

2.9.1. DEFINITION. Let $E = \{X = t_X \mid X \in V_E\}$ be a recursive specification, and let $I \subseteq A$. A subset C of V_E is called a *cluster (of I) in E* iff for all $X \in C$:

$$t_X = \sum_{k=1}^m i_k \cdot X_k + \sum_{l=1}^n Y_l$$

(For $m \geq 0$, $i_1, \dots, i_m \in I \cup \{\tau\}$, $X_1, \dots, X_m \in C$, $n \geq 0$ and $Y_1, \dots, Y_n \in V_E - C$).

Variables $X \in C$ are called *cluster variables*. For $X \in C$ and $Y \in V_E$ we say that

$$X \rightsquigarrow Y \Leftrightarrow Y \text{ occurs in } t_X.$$

We define

$$e(C) = \{Y \in V_E - C \mid \exists X \in C: X \rightsquigarrow Y\}$$

Variables in $e(C)$ are called *exits*. $\rightsquigarrow^*$ is the transitive and reflexive closure of $\rightsquigarrow$. Cluster C is *conservative* iff every exit can be reached from every cluster variable via a path in the cluster:

$$\forall X \in C \forall Y \in e(C): X \rightsquigarrow^* Y.$$

2.9.2. EXAMPLE. The transition diagram of Figure 1 represents a cluster in a recursive specification. The nodes represent variables in the recursive specification, labelled edges represent summands, and the triangles denote exits. The sets $\{1,2,3\}$, $\{4,5,6,7\}$, $\{8\}$ and $\{1,2,3,4,5,6,7,8\}$ are examples of conservative clusters. Cluster $\{1,2,3,4,5,6,7\}$ is not conservative since exit Z cannot be reached from cluster variables 4, 5, 6 and 7.




FIGURE 1

2.9.3. DEFINITION. The *Cluster Fair Abstraction Rule (CFAR)* reads as follows:

(CFAR) Let E be a guarded recursive specification; let $I \subseteq A$ with $|I| \geq 2$; let C be a finite conservative cluster of I in E ; and let $X, X' \in C$ with $X \rightsquigarrow X'$. Then: $\tau_I(X) = \tau \cdot \sum_{Y \in e(C)} \tau_I(Y)$

2.9.4. THEOREM. $\text{ACP}_\tau + \text{RN} + \text{REC} + \text{RSP} + \text{KFAR} \vdash \text{CFAR}$.

PROOF. See [167]. $\square$

2.10. *Alphabets.* Intuitively the alphabet of a process is the set of atomic actions which it can perform. This idea is formalised in [12], where an operator $\alpha: P \rightarrow 2^A$ is introduced, with axioms such as:

$$\alpha(\delta) = \emptyset$$

$$\alpha(ax) = \{a\} \cup \alpha(x)$$

$$\alpha(x+y) = \alpha(x) \cup \alpha(y)$$

In this approach the question arises what axioms should be adopted for the set-operators $\cup$, $\cap$, etc. One option, which is implicitly adopted in previous papers on process algebra, is to take the equalities which are true in set theory. This collection is unstructured and too large for our purposes. Therefore we propose a different, more algebraic solution. We view the alphabet of a process as a *process*; the alphabet operator α goes from sort P to sort P . Process $\alpha(x)$ is the alternative composition of the actions which can be performed by x . In this way we represent a set of actions by a process. A set B of actions is represented by the process expression $B =_{\text{def}} \sum_{b \in B} b$. So the empty set is

represented by δ , a singleton-set $\{a\}$ by the expression a , and a set $\{a, b\}$ by expression $a + b$. Set union corresponds to alternative composition. The process algebra axioms A1-3 and A6 correspond to similar axioms for the set union operator. The notation $\subseteq$ for summand inclusion between processes (Section 2.1.2), fits with the notation for the subset predicate on sets.

The following axioms in Table 8 define the alphabet of finite processes ($a \in A$). Module AB is parametrised by A .

AB	$\alpha(\delta) = \delta$	AB1
	$\alpha(ax) = a + \alpha(x)$	AB2
	$\alpha(x+y) = \alpha(x) + \alpha(y)$	AB3
	$\alpha(\tau) = \delta$	AB4
	$\alpha(\tau x) = \alpha(x)$	AB5

TABLE 8

In order to compute the alphabet of infinite processes, we introduce an additional module AA which is parametrised by A . It is not hard to see that the axioms of AA hold for all closed recursion-free terms.

2.10.1. EXAMPLE. (from [12]). Let $p = \langle X \mid \{X = aX\} \rangle$, and define $q = \tau_{\{a\}}(p)$, $r = q \cdot b$ (with $b \neq a$). What is the alphabet of r ? We derive:

$$\begin{aligned}
 \alpha(r) &= \alpha(qb) = \alpha(\tau_{\{a\}}(p) \cdot b) = \alpha(\tau_{\{a\}}(p) \cdot \tau_{\{a\}}(b)) = \\
 &= \alpha(\tau_{\{a\}}(pb)) \stackrel{\text{AA3}}{\subseteq} \tau_{\{a\}} \circ \partial_{\{a\}} \circ \alpha(pb) \stackrel{\text{RN5}}{=} \partial_{\{a\}} \circ \alpha(pb).
 \end{aligned}$$

AA	$\alpha(x) \subseteq A$	AA1
	$\alpha(x \parallel y) = \alpha(x) + \alpha(y) + \alpha(x) \mid \alpha(y)$	AA2
	$\alpha \circ \rho_f(x) \subseteq \rho_f \circ \partial_H \alpha(x)$ (where $H = \{a \in A \mid f(a) = \tau\}$)	AA3
	$\frac{\forall n \in \mathbb{N} \quad \alpha(\pi_n(x)) \subseteq \gamma}{\alpha(x) \subseteq \gamma}$	AA4

TABLE 9

Since

$$\alpha(pb) = \alpha(apb) \stackrel{AB2}{=} a + \alpha(pb),$$

we have that $a \subseteq \alpha(pb)$. On the other hand we derive for $n \in \mathbb{N}$:

$$\alpha(\pi_n(pb)) = \alpha(a^n \cdot \delta) \subseteq a$$

and therefore, by application of axiom AA4, $\alpha(pb) \subseteq a$.

Consequently $\alpha(pb) = a$ and

$$\alpha(r) = \partial_{\{a\}} \circ \alpha(pb) = \partial_{\{a\}}(a) = \delta.$$

Information about alphabets must be available if we want to apply the following rules. These rules, which are a generalisation of the conditional axioms of [12], occur in a slightly different form also in [168]. Rules like these are an important tool in system verifications based on process algebra. Module RR is parametrised by A and γ .

$\frac{\alpha(x) \subseteq B}{\rho_f(x) = x} \forall b \in B : f(b) = b$	RR1
$\frac{\alpha(x) \subseteq B, \alpha(y) \subseteq C}{\rho_f(x \parallel y) = \rho_f(x) \parallel \rho_f(y)} \forall c \in C. f(c) = f^2(c) \wedge (\forall b \in B. f \circ \gamma(b, c) = f \circ \gamma(b, f(c)))$	RR2

TABLE 10

Observe that axioms AA1 and RR1 together imply axiom RN4 of Table 5. Axiom RR2, which describes the interaction between renaming and parallel composition, looks complicated, but that is only because it is so general. The axioms RR are derivable for closed recursion-free terms.

2.10.2. **LEMMA: (Conditional Axioms (CA)):** Let CA be the theory consisting of the conditional axioms in Table 11. Then: $ACP_\tau + RN + AB + RR \vdash CA$.

$\frac{\alpha(x) (\alpha(y) \cap H) \subseteq H}{\partial_H(x y) = \partial_H(x \partial_H(y))}$	CA1	$\frac{\alpha(x) (\alpha(y) \cap I) = \emptyset}{\tau_I(x y) = \tau_I(x \tau_I(y))}$	CA2
$\frac{\alpha(x) \cap H = \emptyset}{\partial_H(x) = x}$	CA3	$\frac{\alpha(x) \cap I = \emptyset}{\tau_I(x) = x}$	CA4
$\frac{H = H_1 \cup H_2}{\partial_H(x) = \partial_{H_1} \circ \partial_{H_2}(x)}$	CA5	$\frac{I = I_1 \cup I_2}{\tau_I(x) = \tau_{I_1} \circ \tau_{I_2}(x)}$	CA6
$\frac{H \cap I = \emptyset}{\tau_I \circ \partial_H(x) = \partial_H \circ \tau_I(x)}$			

TABLE 11

PROOF: We prove three of the rules. The others can be dealt with similarly.

CA3: Choose $a \in \alpha(x)$. Then $a \notin H$. This means that $r_{\delta, H}(a) = a$. Because a was chosen arbitrarily, we can apply rule RR1, which gives $\rho_{r_{\delta, H}}(x) = \partial_H(x) = x$.

CA5: Follows immediately from the observation

$$r_{\delta, H} = r_{\delta, H_1} \circ r_{\delta, H_2}$$

and application of axiom RN5 of Table 5.

CA1: Choose $c \in \alpha(y)$. We have:

$$r_{\delta, H}(c) = r_{\delta, H} \circ r_{\delta, H}(c)$$

Choose $b \in \alpha(x)$. If $c \notin H$ then $r_{\delta, H}(c) = c$ and the condition of rule RR2 is fulfilled. If $c \in H$ then either $\gamma(b, c)$ equals δ (so that we have $r_{\delta, H} \circ \gamma(b, c) = \delta$), or $\gamma(b, c) \in H$, so that again $r_{\delta, H} \circ \gamma(b, c) = \delta$. But in case $c \in H$ we also have

$$r_{\delta, H} \circ \gamma(b, r_{\delta, H}(c)) = r_{\delta, H} \circ \gamma(b, \delta) = \delta$$

This means that we can apply rule RR2. $\square$

2.10.3. **REMARK.** In most of the situations where we want to apply axiom CA1, H does not contain results of communications: $(A | A) \cap H = \emptyset$. Further actions from $\alpha(x)$ will not communicate with actions from H . In these cases the following weakened version of axiom CA1 is already strong enough:

$$\frac{\alpha(x) | H = \emptyset}{\partial_H(x || y) = \partial_H(x || \partial_H(y))} \text{ CA1*}$$

2.11. $ACP_\tau^\#$. The combination of all modules presented thus far, except for KFAR, will be called $ACP_\tau^\#$ (the system $ACP_\tau^\#$ as presented here slightly differs from a system with the same name occurring in [28]). The module is defined by:

$$ACP_\tau^\# = ACP_\tau + SC + RN + CH^+ + REC + PR + B + AIP^- + AB + AA + RR$$

Bisimulation semantics, as described in for instance [13], gives a model for the module $ACP_\tau^\# + KFAR$. Work of BERGSTRA, KLOP & OLDEROG [29] showed that in a large number of interesting models KFAR is not valid. Therefore we have chosen not to include KFAR in the ‘standard’ module $ACP_\tau^\#$.

2.12. *Generalised Recursive Specification Principle*. For many applications the RSP is too restrictive. Therefore we will present below a more general version of this rule, called RSP^+ .

2.12.1. **DEFINITION.** Let $\mathcal{P}$ be the set of closed expressions in the signature of $ACP_\tau^\#$. A process expression $p \in \mathcal{P}$ is called *guardedly specifiable* if there exists a guarded recursive specification F with $Y \in V_F$ such that

$$ACP_\tau^\# \vdash p = \langle Y | F \rangle.$$

We have the following theorem:

2.12.2. **THEOREM** (*Generalised Recursive Specification Principle* (RSP^+)).
 $ACP_\tau^\# \vdash$

$(RSP^+) \quad \frac{E}{x = \langle x E \rangle} \quad \langle x E \rangle \text{ guardedly specifiable}$

2.12.3. *Remarks.* In the definition of the notion ‘guardedly specifiable’, it is essential that the identity $p = \langle Y | F \rangle$ is *provable*. If we would only require that $p = \langle Y | F \rangle$, then the corresponding version of RSP^+ would not be provable from $ACP_\tau^\#$, since this rule would then not be valid in the action relation model of [77]. In this model we have the identity $\langle X | \{X = X\} \rangle = \delta$.¹ Hence $\langle X | \{X = X\} \rangle = \langle Y | \{Y = \delta\} \rangle = \delta$. Since the specification $\{Y = \delta\}$ is guarded, this would mean that expression $\langle X | \{X = X\} \rangle$ is guardedly specifiable. But then RSP^+ gives that for arbitrary x : $x = \langle X | \{X = X\} \rangle = \delta$. This is clearly false.

1. Strictly speaking, this is not correct. In [77], a recursion construct $\langle X | E \rangle$ is viewed as a kind of variable which ranges over the X -components of the solutions of E . Since any process X satisfies $X = X$, the identity $\langle X | \{X = X\} \rangle = \delta$ does not hold under this interpretation. However, if one interprets the construct $\langle X | E \rangle$ as a constant in the model of [77], then the most natural choice is to relate to $\langle X | E \rangle$ the bisimulation equivalence class of the term $\langle X | E \rangle$. Under this interpretation $\langle X | \{X = X\} \rangle = \delta$.

We conjecture that an expression p is guardably specifiable iff it is provably bounded, i.e. for all $n \in \mathbb{N}$: $\text{ACP}_\tau^\# \vdash B_n(x)$.

3. APPLICATIONS OF THE MODULE APPROACH IN PROCESS ALGEBRA

3.1. The auxiliary status of the left-merge.

3.1.1. Semantics. Sometimes it happens that our ‘customers’ complain that they do not succeed in proving the identity of two processes in $\text{ACP}_\tau^\#$, whose behaviour is considered ‘intuitively the same’. Often, this is because there are many intuitions possible, and $\text{ACP}_\tau^\#$ happens not to represent the particular intuitions of these customers. Therefore we have defined some auxiliary modules that should bridge the gaps between intuitions.

In general a user of process algebra wants that his system proves $p = q$ (here p and q are closed process expressions in the signature of $\text{ACP}_\tau^\#$), whenever p and q have the same interesting properties. So it depends on what properties are interesting for a particular user, whether his system should be designed to prove the equality of p and q or not. For this reason the semantical branch of process algebra research generated a variety of process algebras in which different identification strategies were pursued. In *bisimulation semantics* we find algebras that distinguish between any two processes that differ in the precise timing of internal choices; in *trace semantics* only processes are distinguished which can perform different sequences of actions; and, somewhere in between, the algebras of *failure semantics* identify processes if they have the same traces (can perform the same sequences of actions) and have the same deadlock behaviour in any context. A lot of these process algebras can be organised as homomorphic images of each other, as indicated in Figure 2. For concrete process algebra (without τ -moves) these process algebras have been defined in Chapter I. If two process expressions p and q represent the same process in bisimulation semantics with explicit divergence, they have many properties in common; if they only represent the same process in trace semantics, this only guarantees that they share some of these properties; and, descending from bisimulation semantics with explicit divergence to trace semantics, less and less distinctions are made. Now a user should state exactly in which properties of processes (s)he is interested. Suppose (s)he is only interested in traces and deadlock behaviour, then we can tell that for this purpose failure semantics suffices. This means that if processes p and q are proven equal in failure semantics, this guarantees that they have the same relevant properties. If they are only identified in trace semantics (somewhere in the lattice below failure semantics) such a conclusion cannot be drawn, but if they are identified in a semantics finer than failure semantics (such as bisimulation semantics with explicit divergence), then they certainly have the same interesting properties, and probably some uninteresting ones as well. Hence a proof in bisimulation semantics with explicit divergence is just as good as one in failure semantics (or even better).




FIGURE 2. The linear time - branching time spectrum

This is the reason that we do our proofs mostly in bisimulation semantics: the entire module $ACP_{\tau}^{\#}$ is sound with respect to bisimulation semantics with explicit divergence. However, if two processes are different in bisimulation semantics, we will never succeed in proving them equal from $ACP_{\tau}^{\#}$. In such a case we might add some axioms to the system, that represent the extra identifications made in a less discriminating semantics. If we find a proof from this enriched module, it can be used by anyone satisfied with the properties of this coarser semantics.

It is in the light of the above considerations that one should judge the appearance of the following module T4:

$$\text{T4} \quad \boxed{\tau(\tau x + y) = \tau x + y}$$

The law of this module does not hold in bisimulation semantics, but it does hold in all other semantics of Figure 2. Thus any identity derived from $ACP_{\tau}^{\#} + \text{T4}$ holds in ready trace semantics and hence also in the courser ones like failure and trace semantics, or so it seems

3.1.2. An inconsistency.

3.1.2.1. DEFINITION. Let M be a process module with $\Sigma(M) \supseteq \Sigma(\text{BPA}_{\tau\delta})$. We call M *consistent* if for all closed expressions x and y in the signature of $\text{BPA}_{\tau\delta}$ with

$$M \vdash x = y,$$

the sets of complete traces agree:

$$\text{trace}(x) = \text{trace}(y).$$

A *complete trace* is a finite sequence of actions, ending with a symbol $\sqrt{}$ or δ indicating successful resp. unsuccessful termination. A formal definition of the set $\text{trace}(x)$ is given in [29]. Here we only give some examples, which should make the notion sufficiently clear:

$$\text{trace}(abc + ad\delta + a(\tau bc + d)) = \{abc\sqrt{}, ad\delta, ad\sqrt{}\}$$

$$\text{trace}(\tau) = \{\sqrt{}\} \neq \{\delta, \sqrt{}\} = \text{trace}(\tau + \tau\delta)$$

A model $\mathcal{Q}$ of M is *consistent* if for all closed expressions x and y in the signature of $\text{BPA}_{\tau\delta}$ with

$$\mathcal{Q} \models x = y,$$

the sets of complete traces agree. The module $\text{ACP}_{\tau}^{\#} + \text{KFAR}$ is consistent because bisimulation semantics with fair abstraction, as described in [13], gives a consistent model for this module. However, KFAR is not valid in any of the other semantics of Figure 2.

3.1.2.2. PROPOSITION.

$\text{ACP}_{\tau} + \text{T4} \vdash \tau(ac + ca) + bc = \tau(\tau(ac + ca) + bc + c(\tau a + b))$.

PROOF.

$$\tau(\tau a + b) \llcorner c = (\tau a + b) \llcorner c = \tau(a \parallel c) + bc = \tau(ac + ca) + bc$$

$$\tau(\tau a + b) \llcorner c = \tau((\tau a + b) \parallel c) = \tau(\tau(ac + ca) + bc + c(\tau a + b)) \quad \square$$

Proposition 3.1.2.2 shows that module $\text{ACP}_{\tau} + \text{T4}$ is not consistent. This sudden inconsistency must be the result of a serious misunderstanding. And indeed, what's wrong is the use of ACP_{τ} in the less discriminating models (say in failure semantics). It happens that, in a setting with τ , failure equivalence (or ready trace equivalence for that matter) is not a congruence for the left-merge $\llcorner$, and this causes all the trouble.

3.1.3. Solution. In applications we do not use the operators $\llcorner$ and $\mid$ directly. In specifications we use the merge operator $\parallel$, and $\llcorner$ and $\mid$ are only auxiliary operators, needed to give a complete axiomatisation of the merge.

Let sap_{τ} be the signature obtained from $\Sigma(\text{ACP}_{\tau})$ by stripping the left-merge and communication-merge:

$$\text{sacp}_\tau = \Sigma(\text{ACP}_\tau) - \{\mathbb{F} : \llcorner : P \times P \rightarrow P, \mathbb{F} : | : P \times P \rightarrow P\}$$

Failure equivalence as in [29], etc. are congruences for the operators of sacp_τ . However, the operators $\llcorner$ and $|$ in ACP_τ are needed to axiomatise the $\parallel$ -operator, and without them even the most elementary equations cannot be derived. Our solution to this problem is based on the following idea. Suppose we want to prove an equation $p = q$ in the signature sacp_τ that holds in ready trace semantics (and hence in failure semantics) but not in bisimulation semantics. Then we first prove an intermediate result from ACP_τ : one or more equations holding in bisimulation semantics (with explicit divergence) and in which no $\llcorner$ and $|$ appear. This intermediate result is preserved after mapping the bisimulation model homomorphically on the ready trace or failure model, and can be combined consistently with the axiom T4. Thus the proof of $p = q$ can be completed. In our language of modules we can describe this as follows. The module

$$\text{SACP}_\tau = H(\text{sacp}_\tau \square (\text{ACP}_\tau + \text{SC}))$$

does not contain the operators $\llcorner$ and $|$ in its visible signature and since failure semantics can be obtained as a homomorphic image of bisimulation semantics, considering that $\text{ACP}_\tau + \text{SC}$ is sound w.r.t. bisimulation semantics and that the operators of sacp_τ carry over to failure semantics, we conclude that this module is sound w.r.t. failure semantics. Hence it can be combined consistently with T4, and SACP_τ is a suitable framework for proving statements in failure semantics.

We would like to stress that the use of the H -operator is essential here. The H -operator makes that from module SACP_τ only *positive* formulas are provable. The following example shows what goes wrong if we also allow non-positive formulas. From the proof of Proposition 3.1.2.2 it follows that:

$$\text{sacp}_\tau \square (\text{ACP}_\tau + \text{SC}) \vdash \frac{\tau(\tau a + b) = \tau a + b}{c(\tau a + b) \subseteq \tau(ac + ca) + bc}$$

Consequently we can prove an inconsistency if we add law T4:

$$\text{sacp}_\tau \square (\text{ACP}_\tau + \text{SC}) + \langle \tau(\tau x + y) = \tau x + y \rangle \vdash c(\tau a + b) \subseteq \tau(ac + ca) + bc$$

So although the formulas provable from module $\text{sacp}_\tau \square (\text{ACP}_\tau + \text{SC})$ contain no left-merge, some of them (which are non-positive) cannot be combined consistently with the laws of ready trace semantics and failure semantics.

3.2. Associativity of the chaining operator. ACP_τ is a universal specification formalism in the sense that in bisimulation semantics every finitely branching, effectively presented process can be specified in ACP_τ by a finite system of recursion equations (see [13]). Still it often turns out that adding new operators to the theory facilitates specification and verification of concurrent systems. In general, adding new operators and laws can have far reaching consequences for the underlying mathematical theory. Often however, new operators are *definable* in terms of others operators and the axioms are *derivable*

from the other axioms. In that case the new operators can be considered as notations which are useful, but do not complicate the underlying theory in any way. Examples of definable operators are the projection operators and the process creation operator of [22].

Just like the left-merge and the communication-merge are needed in order to axiomatise the parallel composition operator, new atomic actions are often needed if we want to define a new operator in terms of more elementary operators. As an example we mention the actions $s(d)$ and $r(d)$ which we need in the definition of the chaining operators. These auxiliary atoms will never be used in process specifications. Unfortunately they have the unpleasant property that they occur in some important algebraic laws for the new operators. One of the properties of the chaining operators we use most is that they are ‘associative’. However, due to the auxiliary actions, the chaining operators are not associative in general. We do not have general associativity in the model of bisimulation semantics. Counterexample:

$$\begin{aligned} (r(d) \ggg (s(d) + s(e))) \ggg r(e) &= c(d) \cdot \delta \\ r(d) \ggg ((s(d) + s(e)) \ggg r(e)) &= c(e) \cdot \delta \end{aligned}$$

However, we *do* have associativity under some very weak assumptions. In the model of bisimulation semantics, the following law is valid (here $H = \{s(d), r(d) \mid d \in D\}$):

$$\boxed{\frac{\partial_H(x)=x, \partial_H(y)=y, \partial_H(z)=z}{(x \ggg y) \ggg z = x \ggg (y \ggg z)} \text{ CC}}$$

It would be much nicer if we somehow could ‘hide’ the auxiliary atoms, and, for the $\ggg$ -operator, have associativity in general. In this section we will see how this can be accomplished by means of the module approach.

3.2.1. The associativity of the chaining operators. Although the rule CC holds in the model of bisimulation semantics, we have not been able to prove it algebraically from module $ACP^\#$. However, we can prove algebraically a weaker version of rule CC if we make some additional assumptions about the alphabet. We assume that besides actions $ch(D)$, the alphabet A contains actions:

$$\bar{H} = \{\bar{s}(d), \bar{r}(d) \mid d \in D\} \text{ en } \underline{H} = \{\underline{s}(d), \underline{r}(d) \mid d \in D\}$$

One may think about these actions as special fresh atoms which are added to A only in order to prove the associativity of the chaining operators.¹ Let $\underline{H} = \{\underline{r}(d), \underline{s}(d) \mid d \in D\}$ and let $\bar{H} = H \cup \underline{H} \cup \bar{H}$. We assume that actions from $\bar{H}$ do not synchronise with the other actions in the alphabet, and that

1. The *Fresh Atom Principle* (FAP) says that we can use new (or ‘fresh’) atomic actions in proofs. In [16], it is shown that FAP holds in bisimulation semantics. We have not included FAP in the theoretical framework of this chapter. Therefore, if we need certain ‘fresh’ atoms in a proof, we have to assume that they were in the alphabet right from the beginning.

$\text{range}(\gamma) \cap \hat{H} = \emptyset$. On $\hat{H}$ communication is given by ($d \in D$):

$$\begin{aligned} \gamma(\bar{s}(d), \bar{r}(d)) &= \gamma(\bar{s}(d), r(d)) = \gamma(s(d), \bar{r}(d)) = \gamma(s(d), r(d)) = \\ &= \gamma(\underline{s}(d), \underline{r}(d)) = \gamma(\underline{s}(d), r(d)) = \gamma(s(d), \underline{r}(d)) = c(d) \end{aligned}$$

We define for $v, w \in \{\uparrow, \downarrow, s, r, \bar{s}, \bar{r}, \underline{s}, \underline{r}\}$ the renaming function vw :

$$vw(a) = \begin{cases} w(d) & \text{if } a = v(d) \text{ for some } d \in D \\ a & \text{otherwise} \end{cases}$$

3.2.1.1. LEMMA. $\text{SACP}_\tau + \text{RN} + \text{CH}^+ + \text{AB} + \text{AA} + \text{RR} \vdash$

$$\frac{\partial_{\hat{H}}(x) = x, \partial_{\hat{H}}(y) = y, \partial_{\hat{H}}(z) = z}{\partial_{\bar{H}}(\rho_{\uparrow\bar{s}}(x) \parallel \rho_{\downarrow\bar{r}}(y)) = x \ggg y = \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(y))}$$

PROOF. We only prove the first equality. The second one follows by symmetry.

$$\begin{aligned} \partial_{\bar{H}}(\rho_{\uparrow\bar{s}}(x) \parallel \rho_{\downarrow\bar{r}}(y)) &= (\text{Note 1 below, RR1}) \\ &= \partial_{\bar{H}} \circ \rho_{s\bar{s}} \circ \rho_{r\bar{r}}(\rho_{\uparrow\bar{s}}(x) \parallel \rho_{\downarrow\bar{r}}(y)) = (\text{RN5, } y = \partial_{\hat{H}}(y)) \\ &= \partial_{\bar{H}} \circ \rho_{s\bar{s}} \circ \rho_{r\bar{r}}(\rho_{\uparrow\bar{s}}(x) \parallel \rho_{\bar{r}\bar{r}} \circ \rho_{\downarrow r}(y)) = (\text{Note 2, RR2}) \\ &= \partial_{\bar{H}} \circ \rho_{s\bar{s}} \circ \rho_{r\bar{r}}(\rho_{\uparrow\bar{s}}(x) \parallel \rho_{\downarrow r}(y)) = (\text{SC4, RN5, } x = \partial_{\hat{H}}(x)) \\ &= \partial_{\bar{H}} \circ \rho_{\bar{r}\bar{r}} \circ \rho_{s\bar{s}}(\rho_{\downarrow r}(y) \parallel \rho_{s\bar{s}} \circ \rho_{\uparrow s}(x)) = (\text{as in Note 2, RR2}) \\ &= \partial_{\bar{H}} \circ \rho_{\bar{r}\bar{r}} \circ \rho_{s\bar{s}}(\rho_{\downarrow r}(y) \parallel \rho_{\uparrow s}(x)) = (\text{RN5}) \\ &= \partial_{\bar{H}} \circ \partial_{\bar{H}}(\rho_{\downarrow r}(y) \parallel \rho_{\uparrow s}(x)) = (\text{Note 3, RR1, SC4}) \\ &= \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(y)) \stackrel{\text{CH1}}{=} x \ggg y \end{aligned}$$

Note 1. Let $B = A - H$. We claim $\alpha(\rho_{\uparrow\bar{s}}(x) \parallel \rho_{\downarrow\bar{r}}(y)) \subseteq B$ (recall that $B =_{\text{def}} \sum_{b \in B} b$).

PROOF: $\alpha(\rho_{\uparrow\bar{s}}(x) \parallel \rho_{\downarrow\bar{r}}(y)) =$

$$\stackrel{\text{AA2}}{=} \alpha \circ \rho_{\uparrow\bar{s}}(x) + \alpha \circ \rho_{\downarrow\bar{r}}(y) + \alpha \circ \rho_{\uparrow\bar{s}}(x) \mid \alpha \circ \rho_{\downarrow\bar{r}}(y) \subseteq$$

(Use that $x \subseteq y \Rightarrow x \mid z \subseteq y \mid z$. Use further $x = \partial_{\hat{H}}(x) \stackrel{\text{RN5}}{=} \partial_{\bar{H}} \circ \partial_{\hat{H}}(x) = \partial_{\bar{H}}(x)$.)

$$\stackrel{\text{AA1}}{\subseteq} \alpha \circ \rho_{\uparrow\bar{s}} \circ \partial_{\bar{H}}(x) + \alpha \circ \rho_{\downarrow\bar{r}} \circ \partial_{\bar{H}}(y) + A \mid A \subseteq$$

(Use that $\text{range}(\gamma) \cap H = \emptyset$.)

$$\stackrel{\text{RN5}}{\subseteq} \alpha \circ \partial_{\bar{H}} \circ \rho_{\uparrow\bar{s}}(x) + \alpha \circ \partial_{\bar{H}} \circ \rho_{\downarrow\bar{r}}(y) + B \subseteq$$

$$\stackrel{\text{AA3, RN4}}{\subseteq} \partial_{\bar{H}} \circ \alpha \circ \rho_{\uparrow\bar{s}}(x) + \partial_{\bar{H}} \circ \alpha \circ \rho_{\downarrow\bar{r}}(y) + B \subseteq$$

(Use that $x \subseteq y$ implies $\rho_f(x) \subseteq \rho_f(y)$.)

$$\begin{array}{l} \text{AA1} \\ \subseteq \partial_H(A) + \partial_H(A) + B = B \end{array}$$

This finishes the proof of the claim.

Note 2. Application of axiom AA1 gives: $\alpha \circ \rho_{\uparrow s}(x) \subseteq A$ and $\alpha \circ \rho_{\downarrow r}(y) \subseteq A$. In order to apply axiom RR2, we first have to check that for all $c \in A$: $\bar{r}\bar{r}(c) = \bar{r}\bar{r} \circ \bar{r}\bar{r}(c)$. This is obviously the case. Because $\text{range}(\gamma) \cap H = \emptyset$, we have for all $b, c \in A$: $\bar{r}\bar{r} \circ \gamma(b, c) = \gamma(b, c)$. Now the last thing to be checked is that for $b, c \in A$: $\gamma(b, c) = \gamma(b, \bar{r}\bar{r}(c))$. This turns out to be the case.

Note 3. Let $C = A - \bar{H}$. We claim: $\alpha(\rho_{\downarrow r}(y) \parallel \rho_{\uparrow s}(x)) \subseteq C$. The proof is similar to the proof in Note 1.

This finishes the proof of the lemma. □

3.2.1.2. THEOREM. $\text{SACP}_r + \text{RN} + \text{CH}^+ + \text{AB} + \text{AA} + \text{RR} +$

$$\frac{\partial_H(x) = x, \partial_H(y) = y, \partial_H(z) = z}{x \ggg (y \ggg z) = (x \ggg y) \ggg z}$$

PROOF. This is essentially Theorem 1.12.2 of [168]. We give a sketch of the proof.

$$\begin{aligned} x \ggg (y \ggg z) &= \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r} \circ \partial_{\bar{H}}(\rho_{\uparrow s}(y) \parallel \rho_{\downarrow r}(z))) = \\ &\stackrel{\text{RN5}}{=} \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \partial_{\bar{H}} \circ \rho_{\downarrow r}(\rho_{\uparrow s}(y) \parallel \rho_{\downarrow r}(z))) = \\ &\stackrel{\text{RR1}}{=} \partial_{\bar{H}} \circ \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \partial_{\bar{H}} \circ \rho_{\downarrow r}(\rho_{\uparrow s}(y) \parallel \rho_{\downarrow r}(z))) = \\ &\stackrel{\text{RR2}}{=} \partial_{\bar{H}} \circ \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(\rho_{\uparrow s}(y) \parallel \rho_{\downarrow r}(z))) = \\ &\stackrel{\text{RR2}}{=} \partial_{\bar{H}} \circ \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(\rho_{\downarrow r} \circ \rho_{\uparrow s}(y) \parallel \rho_{\downarrow r}(z))) = \\ &\stackrel{\text{RR1}}{=} \partial_{\bar{H}} \circ \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r} \circ \rho_{\uparrow s}(y) \parallel \rho_{\downarrow r}(z)) = \\ &\stackrel{\text{RN5}}{=} \partial_{\bar{H}} \circ \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \rho_{\uparrow s} \circ \rho_{\downarrow r}(y) \parallel \rho_{\downarrow r}(z)) = \\ &\stackrel{\text{RR1}}{=} \partial_{\bar{H}} \circ \partial_{\bar{H}}(\rho_{\uparrow s}(\rho_{\uparrow s}(x) \parallel \rho_{\uparrow s} \circ \rho_{\downarrow r}(y)) \parallel \rho_{\downarrow r}(z)) = \\ &\stackrel{\text{RR2}}{=} \partial_{\bar{H}} \circ \partial_{\bar{H}}(\rho_{\uparrow s}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(y)) \parallel \rho_{\downarrow r}(z)) = \\ &\stackrel{\text{RR2}}{=} \partial_{\bar{H}} \circ \partial_{\bar{H}}(\partial_{\bar{H}} \circ \rho_{\uparrow s}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(y)) \parallel \rho_{\downarrow r}(z)) = \\ &\stackrel{\text{RR1}}{=} \partial_{\bar{H}}(\partial_{\bar{H}} \circ \rho_{\uparrow s}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(y)) \parallel \rho_{\downarrow r}(z)) = \\ &\stackrel{\text{RN5}}{=} \partial_{\bar{H}}(\rho_{\uparrow s} \circ \partial_{\bar{H}}(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(y)) \parallel \rho_{\downarrow r}(z)) = (x \ggg y) \ggg z \end{aligned}$$

□

3.2.1.3. THEOREM. $\text{SACP}_\tau + \text{RN} + \text{CH}^+ + \text{AB} + \text{AA} + \text{RR} \vdash$

$$\frac{\partial_{\hat{H}}(x)=x, \partial_{\hat{H}}(y)=y, \partial_{\hat{H}}(z)=z}{x \gg (y \gg z) = (x \gg y) \gg z}$$

PROOF. Let $I = \{c(d) \mid d \in D\}$.

$$\begin{aligned} x \gg (y \gg z) &\stackrel{\text{CH2}}{=} \tau_I(x \gg (\tau_I(y \gg z))) = \\ &\stackrel{\text{CH1}}{=} \tau_I \circ \partial_H(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r} \circ \tau_I(y \gg z)) = \\ &\stackrel{\text{RN5}}{=} \partial_H \circ \tau_I(\rho_{\uparrow s}(x) \parallel \tau_I \circ \rho_{\downarrow r}(y \gg z)) = \\ &\stackrel{\text{RR2}}{=} \partial_H \circ \tau_I(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(y \gg z)) = \\ &\stackrel{\text{RN5}}{=} \tau_I \circ \partial_H(\rho_{\uparrow s}(x) \parallel \rho_{\downarrow r}(y \gg z)) = \\ &\stackrel{\text{CH1}}{=} \tau_I(x \gg (y \gg z)) = \\ &\stackrel{3.2.1.2}{=} \tau_I((x \gg y) \gg z) = \dots = (x \gg y) \gg z \quad \square \end{aligned}$$

3.2.2. *Removing auxiliary atoms.* We will now apply the module approach to remove completely the auxiliary atoms which were used in the definition of the chaining operators. What we want to obtain is a module where ‘inside’ the auxiliary atoms are used to define the chaining operator but where ‘outside’ they are no longer visible and moreover chaining is associative in general. Below we will employ the notation:

$$\sigma \Delta M \equiv (\Sigma(M) - \sigma) \square M.$$

Consider the module:

$$\begin{aligned} \text{CH}^- &= (\{F : a \in P \mid a \in \hat{H}\} \cup \{F : \rho_f : P \rightarrow P \mid f : A_{\tau\delta} \rightarrow A_{\tau\delta}\}) \\ &\Delta(\text{SACP}_\tau + \text{RN} + \text{CH}^+ + \text{AB} + \text{AA} + \text{RR}). \end{aligned}$$

This module cannot be used to prove any formula containing atoms in $\hat{H}$. But unfortunately module CH^- still does not prove the general associativity of the chaining operators:

$$\text{CH}^- \not\vdash x \gg (y \gg z) = (x \gg y) \gg z$$

The reason is that the auxiliary atoms, although removed from the language, are still present in the models of module CH^- . Thus the counterexample $(r(d) \gg (s(d) + s(e))) \gg r(e)$ still works in the models. Let $A^- = A - \hat{H}$. We are interested in consistent models which only contain actions of A^- . The module $\text{CH}^- + \langle \alpha(x) \subseteq A^- \rangle$ does not denote such models: all consistent models of CH^- contain the process A with $\alpha(A) = A \not\subseteq A^-$. Adding the law $\alpha(x) \subseteq A^-$ therefore throws away all consistent models. The right class of models can be denoted with the help of operator S . We consider the module

$$CH = S(CH^-) + \langle \alpha(x) \subseteq A^- \rangle.$$

Some models of module CH^- have consistent submodels which do not contain auxiliary atoms. In these models the law $\alpha(x) \subseteq A^-$ holds. Thus module CH has consistent models.

From Theorems 3.2.1.2 and 3.2.1.3, together with axiom RR1, it follows that:

$$\begin{aligned} CH^- \vdash & \frac{\alpha(x) \subseteq A^-, \alpha(y) \subseteq A^-, \alpha(z) \subseteq A^-}{(x \ggg y) \ggg z = x \ggg (y \ggg z)} \quad \text{and} \\ CH^- \vdash & \frac{\alpha(x) \subseteq A^-, \alpha(y) \subseteq A^-, \alpha(z) \subseteq A^-}{(x \gg y) \gg z = x \gg (y \gg z)}. \end{aligned}$$

From this we can easily see that module CH proves the general associativity of the chaining operators:

$$\begin{aligned} CH \vdash x \ggg (y \ggg z) &= (x \ggg y) \ggg z \quad \text{and} \\ CH \vdash x \gg (y \gg z) &= (x \gg y) \gg z. \end{aligned}$$

3.2.3. The following laws can be easily proven from module CH (here $d, e \in D$):

$$\begin{aligned} \uparrow d \cdot x \gg (\sum_{e \in D} \downarrow e \cdot y^e) &= \tau \cdot (x \gg y^d) & \text{L1} \\ \uparrow d \cdot x \gg \uparrow e \cdot y &= \uparrow e \cdot (\uparrow d \cdot x \gg y) & \text{L2} \\ (\sum_{d \in D} \downarrow d \cdot x^d) \gg (\sum_{e \in D} \downarrow e \cdot y^e) &= \sum_{d \in D} \downarrow d \cdot (x^d \gg (\sum_{e \in D} \downarrow e \cdot y^e)) & \text{L3} \\ (\sum_{d \in D} \downarrow d \cdot x^d) \gg \uparrow e \cdot y &= \sum_{d \in D} \downarrow d \cdot (x^d \gg \uparrow e \cdot y) + \uparrow e \cdot ((\sum_{d \in D} \downarrow d \cdot x^d) \gg y) & \text{L4} \end{aligned}$$

The laws are equally valid when the operator $\gg$ is replaced by $\ggg$, except for law L1 where in addition the τ has to be replaced by $c(d)$.

3.3. $SACP^\#_\tau$. Module $SACP^\#_\tau$ is an ‘improved’ version of module $ACP^\#_\tau$. It is defined by:

$$SACP^\#_\tau = SACP_\tau + RN + CH + REC + PR + B + AIP^- + AB + AA + RR.$$

If modules in the above defining equation have an alphabet as parameter, this is A^- , and if they are parametrised by a communication function this is the restriction γ^- of γ to $(A^- \cup \{\delta\}) \times (A^- \cup \{\delta\})$. The rules RSP, RSP⁺ and CFAR can still be used in a setting with module $SACP^\#_\tau$. We have $SACP^\#_\tau \vdash RSP$, $SACP^\#_\tau \vdash RSP^+$ and $SACP^\#_\tau + KFAR \vdash CFAR$.

4. QUEUES

In the specification of concurrent systems FIFO queues with unbounded capacity often play an important role. We give some examples:

- The semantical description of languages with asynchronous message passing such as CHILL (see Recommendation Z.200 (CHILL language definition), CCITT Study Group XI, 1980),
- The modelling of communication channels occurring in computer networks (see LARSEN & MILNER [127] and VAANDRAGER [167]),
- The implementation of languages with many-to-one synchronous communication, such as POOL (see AMERICA [9] and VAANDRAGER [168]).

Consequently the questions how queues can be specified, and how one can prove properties of systems containing queues, are important. For a nice sample of queue-specifications we refer to the solutions of the first problem of the STL/SERC workshop [69]. Some other references are BROY [44], HOARE [116] and PRATT [154].

4.1. Also in the setting of ACP a lot of attention has been paid to the specification of queues. Below we give an infinite specification of the process behaviour of a queue. Here D is a finite set of data, D^* is the set of finite sequences σ of elements from D , the empty sequence is ϵ . Sequence $\sigma\sigma'$ is the concatenation of sequences σ and σ' . The sequence, only consisting of $d \in D$ is denoted by d as well.

$$\boxed{\begin{aligned} QUEUE &= Q_\epsilon = \sum_{d \in D} \downarrow d \cdot Q_d \\ Q_{\sigma\star d} &= \sum_{e \in D} \downarrow e \cdot Q_{e\star\sigma\star d} + \uparrow d \cdot Q_\sigma \end{aligned}}$$

Note that this infinite specification uses only the signature of BPA_δ (see Section 2.1). We have the following fact:

4.1.1. THEOREM: *Using read/send communication, the process QUEUE cannot be specified in ACP by finitely many recursion equations.*

PROOF: See BAETEN & BERGSTRA [11] and BERGSTRA & TIURYN [31]. □

It turns out that if one allows an arbitrary communication function, or extends the signature with an (almost) arbitrary additional operator, the process *QUEUE* can be specified by finitely many recursion equations. For some nice examples we refer to BERGSTRA & KLOP [28].

4.2. *Definition of the queue by means of chaining.* A problem we had with all ACP-specifications of the queue is that they are difficult to deal with in process verifications. For example, let *BUF* 1 be a buffer with capacity one:

$$\boxed{\begin{aligned} BUF 1 &= \sum_{d \in D} \downarrow d \cdot BUF 1^d \\ BUF 1^d &= \uparrow d \cdot BUF 1 \end{aligned}}$$

In process verifications we need propositions like $QUEUE \gg BUF 1 = QUEUE$ (in Section 5 we present a protocol verification where a similar fact is actually used). However, the proof of this fact starting from the infinite specification is rather complicated. Now the following specification of a queue by means of the (abstract) chaining operator allows for a simple proof of the proposition and numerous other useful identities involving queues. This specification is also described by HOARE [116] (p. 158).

$$\boxed{Q = \sum_{d \in D} \downarrow d \cdot (Q \gg BUF 1^d)}$$

The first thing we have to prove is that the process described above really is a queue.

4.2.1. THEOREM: $Q = QUEUE$.

PROOF: Define for every $n \in \mathbb{N}$ and $\sigma = d_1, \dots, d_m \in D^*$ processes D_σ^n as follows:

$$D_\sigma^n = Q \gg BUF 1 \cdots \text{\textit{n times}} \gg BUF 1^{d_1} \cdots \gg BUF 1^{d_m}$$

So by definition $D_\epsilon^0 = Q$. Using the laws of Section 3.2.3, we derive the following recursion equations:

$$\begin{aligned} D_\epsilon^0 &= Q = \sum_{d \in D} \downarrow d \cdot (Q \gg BUF 1^d) = \sum_{d \in D} \downarrow d \cdot D_d^0 \\ D_{\sigma^*d}^n &= Q \gg BUF 1 \cdots \text{\textit{n times}} \gg BUF 1^{d_1} \cdots \gg BUF 1^{d_m} \gg BUF 1^d = \\ &= \sum_{e \in D} \downarrow e \cdot (Q \gg BUF 1^e \gg BUF 1 \cdots \text{\textit{n times}} \gg BUF 1^{d_1} \cdots \gg BUF 1^{d_m} \gg BUF 1^d) \\ &\quad + \uparrow d \cdot (Q \gg BUF 1 \cdots \text{\textit{n times}} \gg BUF 1^{d_1} \cdots \gg BUF 1^{d_m} \gg BUF 1) = \\ &\stackrel{\text{note}}{=} \sum_{e \in D} \downarrow e \cdot (Q \gg BUF 1 \cdots \text{\textit{n times}} \gg BUF 1^e \gg BUF 1^{d_1} \cdots \gg BUF 1^{d_m} \gg BUF 1^d) \\ &\quad + \uparrow d \cdot (Q \gg BUF 1 \cdots \text{\textit{n+1 times}} \gg BUF 1^{d_1} \cdots \gg BUF 1^{d_m}) = \\ &= \sum_{e \in D} \downarrow e \cdot D_{e^*\sigma^*d}^n + \uparrow d \cdot D_\sigma^{n+1} \end{aligned}$$

Note. In the second last step we moved the data in the sequence of 1-datum-buffers to the right as far as possible. It is easy to see that this is allowed. Suppose that not all data are moved to the right. By applying the associativity of the chaining operator we can rewrite the expression in such a way that we get

a subterm of the form $BUF 1^d \gg BUF 1$. This subterm can be rewritten into $\tau \cdot (BUF 1 \gg BUF 1^d)$. Next we move the initial τ to the front of the sequence using the identity $\tau x \parallel y = \tau(x \parallel y)$ of Proposition 2.1.3, and remove it by means of axiom T1 ($x\tau = x$) of ACP_τ . Now we have moved one datum one place to the right in the queue. We can iterate this procedure until the desired result is obtained.

Define the process Q_ϵ^0 by:

$$Q_\epsilon^0 = \sum_{d \in D} \downarrow d \cdot Q_d^0$$

$$Q_{\sigma^* d}^n = \sum_{e \in D} \downarrow e \cdot Q_{e^* \sigma^* d}^n + \uparrow d \cdot Q_\sigma^{n+1}$$

The specification of process Q_ϵ^0 is clearly guarded. Applying RSP gives us on the one hand that $QUEUE = Q_\epsilon^0$, and on the other hand that $Q = D_\epsilon^0 = Q_\epsilon^0$. Consequently $QUEUE = Q$. $\square$

The proof above shows the ‘view of a queue’ that lies behind the specification of Q . During execution there is a long chain of 1-datum buffers passing messages from ‘the left to the right’. After the input of a new datum on the left, a new buffer is created, containing the new datum and placed at the leftmost position in the chain. Because no buffer is ever removed from the system, the number of empty buffers increases after every output of a datum.

4.2.2. LEMMA: $Q \gg BUF 1 = Q$.

PROOF:

$$\begin{aligned} Q \gg BUF 1 &= \sum_{d \in D} \downarrow d \cdot ((Q \gg BUF 1^d) \gg BUF 1) = \\ &= \sum_{d \in D} \downarrow d \cdot (Q \gg (BUF 1^d \gg BUF 1)) = \\ &= \sum_{d \in D} \downarrow d \cdot (Q \gg \tau \cdot (BUF 1 \gg BUF 1^d)) = \\ &\stackrel{2.1.3}{=} \sum_{d \in D} \downarrow d \cdot (Q \gg (BUF 1 \gg BUF 1^d)) = \\ &= \sum_{d \in D} \downarrow d \cdot ((Q \gg BUF 1) \gg BUF 1^d) \end{aligned}$$

Now apply RSP^+ (from the proof of Theorem 4.2.1 it follows that Q is guardedly specifiable). $\square$

By means of an inductive argument we can easily prove the following corollary of Lemma 4.2.2.

4.2.3. COROLLARY: Let for $\sigma \in D^*$, Q^σ be a queue with content σ :

$$\begin{array}{l} Q^\epsilon = Q \\ Q^{\sigma*d} = Q^\sigma \gg \text{BUF } 1^d \end{array}$$

Then: $\tau \cdot (Q^\sigma \gg \text{BUF } 1) = \tau \cdot Q^\sigma$.

4.2.4. PROPOSITION: $Q \gg Q = Q$.

PROOF:

$$\begin{aligned} Q \gg Q &= \sum_{d \in D} \downarrow d \cdot ((Q \gg \text{BUF } 1^d) \gg Q) = \\ &= \sum_{d \in D} \downarrow d \cdot (Q \gg (\text{BUF } 1^d \gg Q)) = \\ &= \sum_{d \in D} \downarrow d \cdot (Q \gg \tau \cdot (\text{BUF } 1 \gg (Q \gg \text{BUF } 1^d))) = \\ &\stackrel{2.1.3}{=} \sum_{d \in D} \downarrow d \cdot (Q \gg (\text{BUF } 1 \gg (Q \gg \text{BUF } 1^d))) = \\ &= \sum_{d \in D} \downarrow d \cdot ((Q \gg \text{BUF } 1) \gg (Q \gg \text{BUF } 1^d)) = \\ &\stackrel{4.2.2}{=} \sum_{d \in D} \downarrow d \cdot (Q \gg (Q \gg \text{BUF } 1^d)) = \\ &= \sum_{d \in D} \downarrow d \cdot ((Q \gg Q) \gg \text{BUF } 1^d) \end{aligned}$$

Now apply RSP⁺. □

4.2.5. COROLLARY: Let $\sigma, \rho \in D^*$. Then: $\tau(Q^\sigma \gg Q^\rho) = \tau Q^{\sigma*\rho}$.

4.2.6. Remark. It will be clear that the implementation which is suggested by the specification of process Q is not very efficient: at each time the number of empty storage elements equals the number of data that have left the queue. But we can do it even more inefficiently: the following queue doubles the number of empty storage elements each time a datum is written.

$$\bar{Q} = \sum_{d \in D} \downarrow d \cdot (\bar{Q} \gg \uparrow d \cdot \bar{Q})$$

A standard proof gives that $\bar{Q} = \text{QUEUE}$. From the point of view of process algebra this specification is very efficient. It is the shortest specification of a FIFO-queue known to the authors, except for a 5-character specification of PRATT [154]: $\downarrow \uparrow \times D^*$. A problem with Pratt's specification is that a neat axiomatisation of the orthocurrence operator $\times$ is not available. Our $\bar{Q}$ -specification has the disadvantage that it does not allow for simple proofs of identities like $\bar{Q} \gg \bar{Q} = \bar{Q}$.

4.3. *Bags*. In [24] a bag over data domain D is defined by:

$$BAG = \sum_{d \in D} \downarrow d \cdot (\uparrow d \parallel BAG)$$

4.3.1. THEOREM: $Q \gg BAG = BAG$.

PROOF:

$$\begin{aligned} Q \gg BAG &= \sum_{d \in D} \downarrow d \cdot ((Q \gg \uparrow d \cdot BUF\ 1) \gg BAG) = \\ &= \sum_{d \in D} \downarrow d \cdot (Q \gg (\uparrow d \cdot BUF\ 1 \gg BAG)) = \\ &= \sum_{d \in D} \downarrow d \cdot (Q \gg \tau \cdot (BUF\ 1 \gg (BAG \parallel \uparrow d))) = \\ &= \sum_{d \in D} \downarrow d \cdot (Q \gg (BUF\ 1 \gg (BAG \parallel \uparrow d))) = \\ &= \sum_{d \in D} \downarrow d \cdot ((Q \gg BUF\ 1) \gg (BAG \parallel \uparrow d)) = \\ &= \sum_{d \in D} \downarrow d \cdot (Q \gg (BAG \parallel \uparrow d)) = \\ &\stackrel{note}{=} \sum_{d \in D} \downarrow d \cdot ((Q \gg BAG) \parallel \uparrow d) \end{aligned}$$

Now apply RSP.

Note. We claim that $SACP_\tau + RN + CH^+ + AB + AA + RR \vdash (Q \gg (BAG \parallel \uparrow d)) = ((Q \gg BAG) \parallel \uparrow d)$. Let $I = \{c(d) \mid d \in D\}$ and $H = \{r(d), s(d) \mid d \in D\}$. Then:

$$(Q \gg (BAG \parallel \uparrow d)) = \tau_I \circ \partial_H (\rho_{\uparrow s}(Q) \parallel \rho_{\downarrow r}(BAG \parallel \uparrow d)) =$$

(straightforward application of axioms of $AB + AA + RR + SC6$)

$$\begin{aligned} &= \tau_I \circ \partial_H (\rho_{\uparrow s}(Q) \parallel \rho_{\downarrow r}(BAG)) \parallel \uparrow d = \\ &= ((Q \gg BAG) \parallel \uparrow d) \end{aligned}$$

From the claim it follows that $CH \vdash (Q \gg (BAG \parallel \uparrow d)) = ((Q \gg BAG) \parallel \uparrow d)$ and consequently $SACP_\tau^\# \vdash (Q \gg (BAG \parallel \uparrow d)) = ((Q \gg BAG) \parallel \uparrow d)$. $\square$

4.3.2. *Remark.* The identity $BAG \gg Q = BAG$ does not hold. The intuitive argument for this is as follows: if a bag contains an apple and an orange, and the environment wants an apple, then it can just take this apple from the bag. In the case where a system, consisting of the chaining of a bag and a queue, contains an apple and an orange, it can occur that the first element in the queue is an orange. In this situation the environment *has* to take the orange first. The argument that processes $Q \gg BAG$ and BAG are different, because in the first process the environment is not able to pick an apple that is still in the

queue, does not hold. In ACP_τ we abstract from the real-time behaviour of concurrent systems. If the environment waits long enough then the apple will be in the bag.

4.4. A queue that can lose data. In the specification of communication protocols, we often encounter transmission channels that can make errors: they can lose, damage or duplicate data. All process algebra specifications of these channels we have seen thus far were lengthy and often incomprehensible. Consequently it was difficult to prove properties of systems containing these queues. Now, interestingly, the same idea that was used to specify the normal queue by means of the chaining operator, can also be used to specify the various faulty queues. One just has to replace the process $BUF1$ in the definition by a process that behaves like a buffer but can lose, damage or duplicate data.

First we describe a queue FQ that can lose every datum contained in it at every moment, without any possibilities for the environment to prevent this from happening. The basic component of this queue is the following Faulty Buffer with capacity one:

$$\begin{aligned} FBUF1 &= \sum_{d \in D} \downarrow d \cdot FBUF1^d \\ FBUF1^d &= (\uparrow d + \tau) \cdot FBUF1 \end{aligned}$$

If the faulty buffer contains a datum, then this can get lost at any moment through the occurrence of a τ -action. In the equation for $FBUF1^d$ there is no τ -action before the $\uparrow d$ -action because this would make it possible for the buffer to reach a state where datum d could not get lost.

We use the above specification in the definition of the faulty queue FQ :

$$FQ = \sum_{d \in D} \downarrow d \cdot (FQ \gg FBUF1^d)$$

The idea behind this specification of the faulty queue is illustrated in Figure 3.




FIGURE 3. *The faulty queue*

4.4.1. LEMMA: $FBUF1^d \gg FBUF1 = \tau \cdot (FBUF1 \gg FBUF1^d)$.

PROOF:

$$\begin{aligned} FBUF1^d \gg FBUF1 &= \tau \cdot (FBUF1 \gg FBUF1^d) + \tau \cdot (FBUF1 \gg FBUF1) = \\ &= \tau \cdot (FBUF1 \gg FBUF1^d) \end{aligned}$$

In the last step we use that:

$$\tau \cdot (FBUF1 \gg FBUF1) \subseteq FBUF1 \gg FBUF1^d \subseteq \tau \cdot (FBUF1 \gg FBUF1^d). \quad \square$$

Compare the simple definition of FQ with the following $BPA_{\tau\delta}$ -specification of the same process.

4.4.2. Let $\sigma, \rho \in D^*$. We write $\sigma \rightarrow \rho$ if ρ can be obtained from σ by deleting one datum. Let $R(\sigma) = \{\rho \mid \sigma \rightarrow \rho\}$ be the finite set of residues of σ after one deletion. Now $FQUEUE$ is the following process.

$$\begin{aligned} FQUEUE &= FQ_\epsilon = \sum_{d \in D} \downarrow d \cdot FQ_d \\ FQ_{\sigma * d} &= \sum_{e \in D} \downarrow e \cdot FQ_{e * \sigma * d} + \uparrow d \cdot FQ_\sigma + \sum_{\rho \in R(\sigma * d)} \tau \cdot FQ_\rho \end{aligned}$$

4.4.3. THEOREM: $FQ = FQUEUE$.

PROOF: Analogously to the proof of Theorem 4.2.1. Use Lemma 4.4.1. $\square$

Analogous versions of the identities we derived for the normal queue can be derived for the faulty queue in the same way. In the proofs we use Lemma 4.4.1.

4.4.4. PROPOSITION:

- i) $FQ \gg FBUF1 = FQ$,
- ii) Let for $\sigma \in D^*$, FQ^σ be a faulty queue with content σ :

$$FQ^\epsilon = FQ$$

$$FQ^{\sigma^*d} = FQ^\sigma \gg FBUF1^d$$

- Then: $\tau \cdot (FQ^\sigma \gg FBUF1) = \tau \cdot FQ^\sigma$,
- iii) $FBUF1^d \gg FQ = \tau \cdot (FBUF1 \gg (FQ \gg FBUF1^d))$,
- iv) $Q \gg FQ = FQ \gg FQ = FQ$,
- v) Let $\sigma, \rho \in D^*$. Then: $\tau \cdot (FQ^\sigma \gg FQ^\rho) = \tau \cdot FQ^{\sigma^*\rho}$.

4.5. An identity that does not hold. In this subsection we will discuss the identity

$$FQ = Q \gg FBUF1.$$

'Intuitively' the processes FQ and $Q \gg FBUF1$ are equal since both behave like a FIFO-queue that can lose data. Furthermore, with both processes the environment cannot prevent in any way that a datum gets lost. Unlike the situation with the processes $BAG \gg Q$ and BAG which we discussed in Section 4.3, we can think of no 'experiment' that distinguishes between the two processes. Still the identity cannot be proved with the axioms presented thus far.

4.5.1. THEOREM: If parameter D of operator $\gg$ contains more than one element, then $SACP^\# \not\vdash FQ = Q \gg FBUF1$.

PROOF: We show that the identity is not valid in the model of process graphs modulo bisimulation congruence as presented in BAETEN, BERGSTRÄ & KLOP [13]. Suppose that there exists a bisimulation between processes FQ and $Q \gg FBUF1$. Suppose that process FQ reads successively two different data, starting from the initial state. Because of the bisimulation it must be possible for the process $Q \gg FBUF1$ to read the same data in such a way that the resulting state is bisimilar to the state process FQ has reached. Now process FQ executes a τ -step and forgets the second datum. We claim that process $Q \gg FBUF1$ is not capable to perform a corresponding sequence of zero or more τ -step. This is because there are only two possibilities:

- 1) $Q \gg FBUF1$ forgets the second datum. But this means that also the first datum is forgotten. In the resulting state $Q \gg FBUF1$ cannot output any datum (before reading one), whereas process FQ can do this.
- 2) $Q \gg FBUF1$ does not forget the second datum. In the resulting state $Q \gg FBUF1$ can output this datum. Process FQ cannot do that.

The argument is illustrated in Figure 4.




FIGURE 4.

The next theorem shows that, if we add law T4, the two faulty queues can be proven equivalent.

4.5.2. THEOREM: $SACP_r^\# + T4 \vdash FQ = Q \gg FBUF 1$.

PROOF: Define the process QF by:

$$QF = QF_\epsilon = \sum_{d \in D} \downarrow d \cdot QF_d$$

$$QF_{\sigma*d} = \sum_{e \in D} \downarrow e \cdot QF_{e*\sigma*d} + (\uparrow d + \tau) \cdot QF_\sigma$$

Analogous to the proof of Theorem 4.2.1, using in addition the identity $Q \gg FBUF 1 = Q$, we prove that $Q \gg FBUF 1 = QF$. For this we do not need the additional axiom.

The main trick in the proof is that we introduce yet another 'view of queues': process QF_σ is split in two parts, a read-process QF_σ^r and a write process QF_σ^s . The read-process takes care of reading new data, whereas the write process outputs the data in σ . When the write process is ready, it sends a message *ready* to the read-process and dies. When the read-process, after reading a sequence ρ of messages, receives the *ready*-signal it behaves again like process QF_ρ . The fact that the length of the sequence of data σ in QF_σ^s can only decrease, allows us to use induction.

We extend the alphabet¹ with actions *ready*, *ready*^{*} and $\overline{\text{ready}}$, and define

1. See Note 1 in Section 3.2.1.

communication by $\gamma(\text{ready}, \text{ready}^*) = \overline{\text{ready}}$. For $\sigma \in D^*$ and $d \in D$ we define:

$$\begin{aligned} QF_\epsilon^r &= \sum_{d \in D} \downarrow d \cdot QF_d^r \\ QF_{\sigma*d}^r &= \sum_{e \in D} \downarrow e \cdot QF_{e*\sigma*d}^r + \text{ready}^* \cdot QF_{\sigma*d} \\ QF_\epsilon^s &= \text{ready} \\ QF_{\sigma*d}^s &= (\uparrow d + \tau) \cdot QF_\sigma^s \end{aligned}$$

A short but nontrivial proof gives:

$$QF_\sigma = \tau_{\{\overline{\text{ready}}\}} \circ \partial_{\{\text{ready}, \text{ready}^*\}} (QF_\epsilon^r \parallel QF_\sigma^s)$$

Also in this step we do not use the extra axiom. We claim that:
 $\tau \cdot QF_{\sigma*\sigma'}^s \subseteq QF_{\sigma*d*\sigma'}^s$.

The proof of the claim goes with induction to the length of σ' . If $|\sigma'| = 0$ the claim holds trivially. Now suppose the claim is proved for $|\sigma'| \leq n$. Choose $\bar{\sigma}$ with length n , and $e \in D$. We have that:

$$\tau \cdot QF_{\sigma*\bar{\sigma}*e}^s = \tau \cdot (\uparrow e \cdot QF_{\sigma*\bar{\sigma}}^s + \tau \cdot QF_{\sigma*\bar{\sigma}}^s) =$$

(this is the only step where we use axiom T4)

$$= \uparrow e \cdot QF_{\sigma*\bar{\sigma}}^s + \tau \cdot QF_{\sigma*\bar{\sigma}}^s \subseteq$$

(because on the one hand $\uparrow e \cdot QF_{\sigma*\bar{\sigma}}^s \subseteq \uparrow e \cdot QF_{\sigma*d*\bar{\sigma}}^s$ because of the induction hypothesis and axiom T3, and on the other hand $\tau \cdot QF_{\sigma*\bar{\sigma}}^s \subseteq \tau \cdot QF_{\sigma*d*\bar{\sigma}}^s$ because of the induction hypothesis and axiom T2)

$$\subseteq \uparrow e \cdot QF_{\sigma*d*\bar{\sigma}}^s + \tau \cdot QF_{\sigma*d*\bar{\sigma}}^s = QF_{\sigma*d*\bar{\sigma}*e}^s$$

This finishes the proof of the claim. A corollary is that $\tau \cdot QF_{\sigma*\sigma'} \subseteq QF_{\sigma*d*\sigma'}$:

$$\tau \cdot QF_{\sigma*\sigma'} = \tau \cdot \tau_{\{\overline{\text{ready}}\}} \circ \partial_{\{\text{ready}, \text{ready}^*\}} (QF_\epsilon^r \parallel QF_{\sigma*\sigma'}^s) \subseteq$$

(Use the observation of Section 2.1.4 that $\tau x \subseteq y$ implies $\tau x \parallel z \subseteq y \parallel z$)

$$\subseteq \tau_{\{\overline{\text{ready}}\}} \circ \partial_{\{\text{ready}, \text{ready}^*\}} (QF_\epsilon^r \parallel QF_{\sigma*d*\sigma'}^s) = QF_{\sigma*d*\sigma'}$$

We have shown that process QF_σ is indistinguishable from a process that can lose each datum at every moment. Using the notation of Section 4.4.2 we can write down the following equation for processes $QF_{\sigma*d}$:

$$QF_{\sigma*d} = \sum_{e \in D} \downarrow e \cdot QF_{e*\sigma*d} + \uparrow d \cdot QF_\sigma + \sum_{\rho \in R(\sigma*d)} \tau \cdot QF_\rho$$

Application of RSP gives that the process $FQUEUE$ of Section 4.4.2 equals process QF . But according to Theorem 4.4.3 also $FQUEUE = FQ$. $\square$

4.6. The faulty and damaging queue. In the specification of certain link layer protocols we have to deal with a communication channel that behaves like a FIFO-queue with unbounded capacity (this is of course a simplifying assumption), but has some additional properties: (1) a datum can be damaged at every moment it is in the queue; the environment cannot prevent this event, and (2) a datum can be lost at every moment it is in the queue. We give a process algebra specification of this process in two steps. First we specify the Faulty and Damaging Buffer with capacity one (FDBUF1). We assume that the domain of data D contains a special element er , representing a damaged datum.

$$\begin{aligned} FDBUF1 &= \sum_{d \in D} \downarrow d \cdot FDBUF1^d \\ FDBUF1^d &= \uparrow d \cdot FDBUF1 + \tau \cdot (\uparrow er + \tau) \cdot FDBUF1 \end{aligned}$$

With the help of this process we can now easily define the Faulty and Damaging Queue (FDQ):

$$FDQ = \sum_{d \in D} \downarrow d \cdot (FDQ \gg FDBUF1^d)$$

4.6.1. LEMMA: $FDBUF1^d \gg FDBUF1 = \tau \cdot (FDBUF1 \gg FDBUF1^d)$.

PROOF: $FDBUF1^d \gg FDBUF1 =$

$$\begin{aligned} &= \tau \cdot (FDBUF1 \gg FDBUF1^d) + \tau \cdot ((\uparrow er + \tau) \cdot FDBUF1 \gg FDBUF1) = \\ &= \tau \cdot (FDBUF1 \gg FDBUF1^d) + \tau \cdot (\tau \cdot (FDBUF1 \gg FDBUF1^{er}) + \\ &\quad + \tau \cdot (FDBUF1 \gg FDBUF1)) = \end{aligned}$$

$$\begin{aligned} &\stackrel{note}{=} \tau \cdot (FDBUF1 \gg FDBUF1^d) + \tau \cdot (\tau \cdot (FDBUF1 \gg (\tau \cdot (\uparrow er + \tau) \cdot FDBUF1)) + \\ &\quad + \tau \cdot (FDBUF1 \gg FDBUF1)) = \end{aligned}$$

$$\begin{aligned} &= \tau \cdot (FDBUF1 \gg FDBUF1^d) + \tau \cdot (\tau \cdot (FDBUF1 \gg (\uparrow er + \tau) \cdot FDBUF1) + \\ &\quad + \tau \cdot (FDBUF1 \gg FDBUF1)) = \end{aligned}$$

$$\begin{aligned} &= \tau \cdot (FDBUF1 \gg FDBUF1^d) + \tau \cdot (FDBUF1 \gg (\uparrow er + \tau) \cdot FDBUF1) = \\ &= \tau \cdot (FDBUF1 \gg FDBUF1^d) \end{aligned}$$

$$Note. \quad FDBUF1^{er} = \uparrow er \cdot FDBUF1 + \tau \cdot (\uparrow er + \tau) \cdot FDBUF1 =$$

$$\stackrel{T2}{=} \tau \cdot (\uparrow er + \tau) \cdot FDBUF1. \quad \square$$

Once we have Lemma 4.6.1, it is standard to prove that process FDQ is guardedly specifiable. It is moreover easy to derive an analogous version of Proposition 4.4.4 for FDQ .

4.6.2. *Remark.* One might ask if there is not a τ too many in the specification of process $FDBUF1$. Why not specify the faulty and damaging buffer simply as follows?

$$\begin{aligned} FDB1 &= \sum_{d \in D} \downarrow \cdot FDB1^d \\ FDB1^d &= (\uparrow d + \uparrow er + \tau) \cdot FDB1 \end{aligned}$$

A first observation we make is that if $D \neq \{er\}$:

$$SACP_{\tau}^{\#} \not\sim FDBUF1 = FDB1$$

This is because the two processes are different in bisimulation semantics. Process $FDBUF1$ can input a datum d different from er , and then get into a state where either an output action $\uparrow er$ will be performed or no output action at all. This means that it is possible that a datum is first damaged and then lost. Process $FDB1$ does not have such a state.

For similar reasons we also have the following fact:

$$SACP_{\tau}^{\#} \not\sim FDB1^d \gg FDB1 = \tau \cdot (FDB1 \gg FDB1^d)$$

This means that if we work with a queue defined with the help of $FDB1$, our standard technique to prove facts about queues is not applicable. Note that processes $FDB1$ and $FDBUF1$ are trivially equal if we work in a setting where the law T4 ($\tau(\tau x + y) = \tau x + y$) is valid.

4.7. *The faulty and stuttering queue.* This section is about a very curious queue: a FIFO-queue that can lose or duplicate any element contained in it at every moment. An infinite specification of this process can be found in LARSEN & MILNER [127]. The basic component we use in the specification of the Faulty and Stuttering Queue is a Faulty and Stuttering Buffer with capacity 1:

$$\begin{aligned} FSBUF1 &= \sum_{d \in D} \downarrow d \cdot FSBUF1^d \\ FSBUF1^d &= \uparrow d \cdot FSBUF1^d + \tau \cdot FSBUF1 \end{aligned}$$

$$FSQ = \sum_{d \in D} \downarrow d \cdot (FSQ \gg FSBUF1^d)$$

When we place two faulty and stuttering buffers in a chain, then we have the possibility of an infinite number of internal actions (the first buffer stutters and the second one loses all its input). This implies that, in the specification of the faulty and stuttering queue, we have to guard against unguarded recursion. We need a fairness assumption if we want to exclude the possibility of infinite stuttering.

First we prove a simple lemma:

4.7.1. LEMMA: $FSBUF 1^d \gg FSBUF 1^d = \tau \cdot (FSBUF 1^d \gg FSBUF 1^d) =$
 $= FSBUF 1^d \gg FSBUF 1 = \tau \cdot (FSBUF 1^d \gg FSBUF 1).$

PROOF:

$$\begin{aligned} FSBUF 1^d \gg FSBUF 1^d &\supseteq \tau \cdot (FSBUF 1^d \gg FSBUF 1) \supseteq FSBUF 1^d \gg FSBUF 1 \supseteq \\ &\supseteq \tau \cdot (FSBUF 1^d \gg FSBUF 1^d) \supseteq FSBUF 1^d \gg FSBUF 1^d \quad \square \end{aligned}$$

The proof of the next lemma is more involved.

4.7.2. LEMMA:

$$SACP_{\tau}^{\#} + KFAR \cdot FSBUF 1^d \gg FSBUF 1 = \tau \cdot (FSBUF 1 \gg FSBUF 1^d).$$

PROOF:

$$\begin{aligned} FSBUF 1^d \gg FSBUF 1 &= \\ &= c(d) \cdot (FSBUF 1^d \gg FSBUF 1^d) + \tau \cdot (FSBUF 1 \gg FSBUF 1) \\ FSBUF 1^d \gg FSBUF 1^d &= \\ &= \tau \cdot (FSBUF 1^d \gg FSBUF 1) + \tau \cdot (FSBUF 1 \gg FSBUF 1^d) + \\ &\quad + \uparrow d \cdot (FSBUF 1^d \gg FSBUF 1^d) \end{aligned}$$

Application of CFAR gives ($I = \{c(d) \mid d \in D\}$):

$$\begin{aligned} FSBUF 1^d \gg FSBUF 1 &= \tau_I (FSBUF 1^d \gg FSBUF 1) = \\ &= \tau \cdot \tau_I (\tau \cdot (FSBUF 1 \gg FSBUF 1) + \tau \cdot (FSBUF 1 \gg FSBUF 1^d) + \\ &\quad + \uparrow d \cdot (FSBUF 1^d \gg FSBUF 1^d)) = \\ &\stackrel{4.7.1}{=} \tau \cdot (\tau \cdot (FSBUF 1 \gg FSBUF 1) + \tau \cdot (FSBUF 1 \gg FSBUF 1^d) + \\ &\quad + \uparrow d \cdot (FSBUF 1^d \gg FSBUF 1)) \end{aligned}$$

In addition we derive:

$$\begin{aligned} FSBUF 1 \gg FSBUF 1^d &= \sum_{e \in D} \downarrow e \cdot (FSBUF 1^e \gg FSBUF 1^d) + \\ &\quad + \tau \cdot (FSBUF 1 \gg FSBUF 1) + \uparrow d \cdot (FSBUF 1 \gg FSBUF 1^d) \\ FSBUF 1^e \gg FSBUF 1^d &= \tau \cdot (FSBUF 1 \gg FSBUF 1^d) + \\ &\quad + \tau \cdot (FSBUF 1^e \gg FSBUF 1) + \uparrow d \cdot (FSBUF 1^e \gg FSBUF 1^d) \\ FSBUF 1 \gg FSBUF 1 &= \sum_{d \in D} \downarrow d \cdot (FSBUF 1^d \gg FSBUF 1) \end{aligned}$$

Let E be the following guarded system of recursion equations:

$$\begin{aligned} Y^d &= \tau \cdot (\tau \cdot Y + \tau \cdot Y^d + \uparrow d \cdot Y^d) \\ Y^d &= \sum_{e \in D} \downarrow e \cdot Y^{ed} + \tau \cdot Y + \uparrow d \cdot Y^d \end{aligned}$$

$$Y^{ed} = \tau \cdot Y^d + \tau \cdot Y^e + \uparrow d \cdot Y^{ed}$$

$$Y = \sum_{d \in D} \downarrow d \cdot Y^d$$

RSP gives that $FSBUF1^d \gg FSBUF1 = Y^d$ and $FSBUF1 \gg FSBUF1^d = Y^d$. Thus it suffices to prove that $Y^d = \tau \cdot Y^d$. Let F be the following guarded system of recursion equations:

$$Z^d = \sum_{e \in D} \downarrow e \cdot Z^{ed} + \tau \cdot Z + \uparrow d \cdot Z^d$$

$$Z^{ed} = \tau \cdot Z^d + \tau \cdot Z^e + \uparrow d \cdot Z^{ed}$$

$$Z = \sum_{d \in D} \downarrow d \cdot Z^d$$

We derive:

$$\tau \cdot Z^d = \tau \cdot \tau \cdot Z^d = \tau \cdot (\tau \cdot Z + \uparrow d \cdot Z^d + \tau \cdot Z^d)$$

If we substitute $\tau \cdot Z^d$ for Y^d , Z^d for Y^d , Z^{ed} for Y^{ed} and Z for Y , then RSP gives: $\tau \cdot Z^d = Y^d$ and $Z^d = Y^d$. Consequently $Y^d = \tau \cdot Y^d$. $\square$

From Lemma 4.7.2 all the rest follows: process FSQ is guardedly specifiable and we can derive an analogous version of Proposition 4.4.4.

5. A PROTOCOL VERIFICATION

In this section we present the specification and verification of a variant of the Alternating Bit Protocol, resembling the ones discussed in KOYMANS & MULDER [123] and LARSEN & MILNER [127]. The aim of this exercise is to illustrate the usefulness of the proof technique developed in the previous section. The architecture of the *Concurrent Alternating Bit Protocol (CABP)* is as follows:




FIGURE 5.

Elements of a finite set of data are to be transmitted by the CABP from port 1 to port 2. Verification of the CABP amounts to a proof that (1) the protocol

will eventually send at port 2 all and only data it has read at port 1, and (2) the protocol will send the data at port 2 in the same order as it has read them at port 1.

In the CABP sender and receiver send frames continuously. Since sender and receiver will have a different clock in general, the number of data that can be in the channels at a certain moment is in principle unlimited. In this section we assume that the channels behave like the process FQ as described in Section 4.4: a FIFO-queue with unbounded capacity that can either lose frames or pass them on correctly.

In the protocol, the sender consists of two components A and D , whereas the receiver consists of components B and C . One might propose to collapse A and D into a sender process, and B plus C into a receiver process. The resulting processes would be more complicated and in the correctness proof we would have to decompose them again.

5.1. Specification. Let D be a finite set of data which have to be sent by the CABP from port 1 to port 2. Let $B = \{0, 1\}$. $\mathcal{D} = (D \times B) \cup B$ is the set of data which occur as parameter in the actions of the chaining operators. The set of ports is $\mathbb{P} = \{1, 2, 3, 4\}$, the set of data that can be communicated at these ports is $\mathbb{D} = D \cup \{next\}$. Alphabet $\mathcal{A}$ and communication function γ are now defined by the standard scheme for the chaining operators, augmented with actions $ri(d)$, $si(d)$ and $ci(d)$, for which we have communications $\gamma(ri(d), si(d)) = ci(d)$ ($i \in \mathbb{P}$ and $d \in \mathbb{D}$).

We now give the specifications of processes A , B , C and D . Here b ranges over $B = \{0, 1\}$ and d over D (the overloading of names B and D should cause no confusion). The specifications are standard and need no further comment.

$A = A^0$ $A^b = \sum_{d \in D} r1(d) \cdot A^{db}$ $A^{db} = \uparrow db \cdot A^{db} + r3(next) \cdot A^{1-b}$	$B = B^0$ $B^b = \sum_{d \in D} \downarrow(d, 1-b) \cdot B^b + \sum_{d \in D} \downarrow db \cdot B^{db}$ $B^{db} = s2(d) \cdot s4(next) \cdot B^{1-b}$
$D = D^0$ $D^b = \downarrow(1-b) \cdot D^b + \downarrow b \cdot s3(next) \cdot D^{1-b}$	$C = C^1$ (not C^0 !) $C^b = \uparrow b \cdot C^b + r4(next) \cdot C^{1-b}$

Let H and I be the following sets of actions:

$$H = \{r3(next), s3(next), r4(next), s4(next)\}$$

$$I = \{c3(next), c4(next)\}$$

The Concurrent Alternating Bit Protocol is defined by:

$$CABP = \tau_1 \circ \partial_H((A \gg FQ \gg B) \parallel (C \gg FQ \gg D))$$

5.2. Verification. If we do not abstract from the internal actions of the protocol, then the number of states is infinite. This means that a straightforward calculation of the state graph is not possible. A strategy which is often applied in cases like this is that one substitutes a buffer with capacity 1 for the communication channels. As a result the system is finite and can be verified automatically. Next a buffer with capacity 2 is substituted, followed by another automatic verification, etc.. The verification for the case of buffers with capacity 155 takes 23 hours CPU time. Thereafter it is decided that ‘the protocol is correct’.

Of course it is not so difficult to specify a protocol that is correct for buffers with capacity less or equal than 155, but fails when the capacity is 156. The conclusion that the protocol is correct for arbitrary buffer size because it works in the cases where the buffer size is less than 156, is therefore influenced by other observations. It is for example intuitively not very plausible that the CABP works for buffer size 155, but not for buffer size 156, because the specification is so short and the only numbers which occur in it are 0 and 1.

Because intuitions can be wrong people look for formal techniques which tell in which situations induction over certain protocol parameters is allowed.

The basic merit of the results of Section 4 is that they make it possible to use inductive arguments when dealing with the length of queues in protocol systems. In the verification below we show that the protocol is correct if the channels behave as faulty FIFO-queues with unbounded capacity. However, a minor change in the proof is enough to show that the protocol also works if the channels behave as n -buffers, faulty n -buffers, perfect queues, faulty and stuttering queues, etc.

The following two lemmas will be used to show that, after abstraction, the number of states of the protocol is finite. The first lemma says that if, at the head of the queue, there is a datum that will be thrown away by the receiver because it is of the wrong type, this datum can be thrown away immediately.

5.2.1. LEMMA:

- i) $FBUF\ 1^{db} \gg B^{1-b} = \tau \cdot (FBUF\ 1 \gg B^{1-b});$
- ii) $FBUF\ 1^{db} \gg s\ 4(next) \cdot B^{1-b} = \tau \cdot (FBUF\ 1 \gg s\ 4(next) \cdot B^{1-b});$
- iii) $FBUF\ 1^{db} \gg B^{db} = \tau \cdot (FBUF\ 1 \gg B^{db}).$

PROOF: The proof of (i) is trivial. Part (ii) goes as follows:

$$\begin{aligned}
 FBUF\ 1^{db} \gg s\ 4(next) \cdot B^{1-b} &= \\
 &= \tau \cdot (FBUF\ 1 \gg s\ 4(next) \cdot B^{1-b}) + s\ 4(next) \cdot (FBUF\ 1^{db} \gg B^{1-b}) = \\
 &\stackrel{(i)}{=} \tau \cdot (FBUF\ 1 \gg s\ 4(next) \cdot B^{1-b}) + s\ 4(next) \cdot (FBUF\ 1 \gg B^{1-b}) = \\
 &= \tau \cdot (FBUF\ 1 \gg s\ 4(next) \cdot B^{1-b}) \quad (\text{summand inclusion})
 \end{aligned}$$

The proof of part (iii) is similar:

$$\begin{aligned}
 FBUF1^{db} \gg B^{db} &= \tau \cdot (FBUF1 \gg B^{db}) + s2(d) \cdot (FBUF1^{db} \gg s4(next) \cdot B^{1-b}) = \\
 &\stackrel{(ii)}{=} \tau \cdot (FBUF1 \gg B^{db}) + s2(d) \cdot (FBUF1 \gg s4(next) \cdot B^{1-b}) = \\
 &= \tau \cdot (FBUF1 \gg B^{db}) \quad \square
 \end{aligned}$$

The next lemma says that if two frames, of a type that the receiver is willing to accept, are at the head of the queue, one of these can be deleted without changing the process (modulo an initial τ).

5.2.2. LEMMA: $FBUF1^{db} \gg FBUF1^{db} \gg B^b = \tau \cdot (FBUF1 \gg FBUF1^{db} \gg B^b)$.

PROOF: $FBUF1^{db} \gg FBUF1^{db} \gg B^b =$

$$\begin{aligned}
 &= \tau \cdot (FBUF1 \gg FBUF1^{db} \gg B^b) + \tau \cdot (FBUF1^{db} \gg FBUF1 \gg B^b) + \\
 &\quad + \tau \cdot (FBUF1^{db} \gg FBUF1 \gg B^{db}) = \\
 &\stackrel{4.4.1}{=} \tau \cdot (FBUF1 \gg FBUF1^{db} \gg B^b) + \tau \cdot (FBUF1 \gg FBUF1^{db} \gg B^b) + \\
 &\quad + \tau \cdot (FBUF1 \gg FBUF1^{db} \gg B^{db}) = \\
 &\stackrel{5.2.1}{=} \tau \cdot (FBUF1 \gg FBUF1^{db} \gg B^b) + \tau \cdot (FBUF1 \gg FBUF1 \gg B^{db}) = \\
 &= \tau \cdot (FBUF1 \gg FBUF1^{db} \gg B^b) \quad \square
 \end{aligned}$$

5.2.3. We will now derive a transition diagram for process $A \ggg FQ \gg B$. In the derivation we use Lemmas 5.2.1 and 5.2.2 to keep the diagram finite. Furthermore we stop the derivation at those places where an action is performed that corresponds to the acknowledgement of a frame that has not yet arrived. The result of the calculations is presented in Figure 6. The grey arcs correspond to places where we stopped the derivation.


$$A \ggg FQ \gg B = A^0 \ggg FQ \gg B^0 \quad (0)$$

$$A^b \ggg FQ \gg B^b = \sum_{d \in D} r1(d) \cdot (A^{db} \ggg FQ \gg B^b) \quad (1)$$

$$\begin{aligned}
 A^{db} \ggg FQ \gg B^b &= c(db) \cdot (A^{db} \ggg FQ \gg FBUF1^{db} \gg B^b) + \\
 &\quad + r3(next) \cdot (A^{1-b} \ggg FQ \gg B^b) \quad (2)
 \end{aligned}$$

$$A^{db} \ggg FQ \gg FBUF1^{db} \gg B^b = \quad (3)$$

$$\begin{aligned}
 &c(db) \cdot (A^{db} \ggg FQ \gg FBUF1^{db} \gg FBUF1^{db} \gg B^b) + \\
 &\quad + \tau \cdot (A^{db} \ggg FQ \gg FBUF1 \gg B^b) + \\
 &\quad + \tau \cdot (A^{db} \ggg FQ \gg FBUF1 \gg B^{db}) + \\
 &\quad + r3(next) \cdot (A^{1-b} \ggg FQ \gg FBUF1^{db} \gg B^b) =
 \end{aligned}$$

FIGURE 6. Transition diagram of process $A \gg FQ \gg B$

(Lemmas 5.2.2 and 4.4.4(i))

$$\begin{aligned}
 &= c(db) \cdot (A^{db} \gg FQ \gg FBUF 1^{db} \gg B^b) + \\
 &\quad + \tau \cdot (A^{db} \gg FQ \gg B^b) + \tau \cdot (A^{db} \gg FQ \gg B^{db}) + \\
 &\quad + r3(next) \cdot (A^{1-b} \gg FQ \gg FBUF 1^{db} \gg B^b) \\
 A^{db} \gg FQ \gg B^{db} &= c(db) \cdot (A^{db} \gg FQ \gg FBUF 1^{db} \gg B^{db}) + \\
 &\quad + s2(d) \cdot (A^{db} \gg FQ \gg s4(next) \cdot B^{1-b}) + \\
 &\quad + r3(next) \cdot (A^{1-b} \gg FQ \gg B^{db}) =
 \end{aligned} \tag{4}$$

(Lemmas 5.2.1(iii) and 4.4.4(i))

$$\begin{aligned}
 &= c(db) \cdot (A^{db} \gg FQ \gg B^{db}) + \\
 &\quad + s2(d) \cdot (A^{db} \gg FQ \gg s4(next) \cdot B^{1-b}) + \\
 &\quad + r3(next) \cdot (A^{1-b} \gg FQ \gg B^{db}) \\
 A^{db} \gg FQ \gg s4(next) \cdot B^{1-b} &= \\
 &\quad c(db) \cdot (A^{db} \gg FQ \gg FBUF 1^{db} \gg s4(next) \cdot B^{1-b}) + \\
 &\quad + s4(next) \cdot (A^{db} \gg FQ \gg B^{1-b}) + \\
 &\quad + r3(next) \cdot (A^{1-b} \gg FQ \gg s4(next) \cdot B^{1-b}) =
 \end{aligned} \tag{5}$$

(Lemmas 5.2.1(ii) and 4.4.4(i))

$$= c(db) \cdot (A^{db} \gg FQ \gg s4(next) \cdot B^{1-b}) +$$

$$\begin{aligned}
& + s4(next) \cdot (A^{db} \ggg FQ \gg B^{1-b}) + \\
& + r3(next) \cdot (A^{1-b} \ggg FQ \gg s4(next) \cdot B^{1-b}) \\
A^{db} \ggg FQ \gg B^{1-b} & = c(db) \cdot (A^{db} \ggg FQ \gg FBUF1^{db} \gg B^{1-b}) + \\
& + r3(next) \cdot (A^{1-b} \ggg FQ \gg B^{1-b})
\end{aligned} \tag{6}$$

(Lemmas 5.2.1(i) and 4.4.4(i))

$$= c(db) \cdot (A^{db} \ggg FQ \gg B^{1-b}) + r3(next) \cdot (A^{1-b} \ggg FQ \gg B^{1-b})$$

5.2.4. Summarising, we have shown that $A \ggg FQ \gg B$ satisfies the following system of recursion equations.

$$X = X_1^0$$

$$X_1^b = \sum_{d \in D} r1(d) \cdot X_2^{db}$$

$$X_2^{db} = c(db) \cdot X_3^{db} + Y_2^b$$

$$Y_2^b = r3(next) \cdot (A^{1-b} \ggg FQ \gg B^b)$$

$$X_3^{db} = \tau \cdot X_2^{db} + c(db) \cdot X_3^{db} + \tau \cdot X_4^{db} + Y_3^{db}$$

$$Y_3^{db} = r3(next) \cdot (A^{1-b} \ggg FQ \gg FBUF1^{db} \gg B^b)$$

$$X_4^{db} = c(db) \cdot X_4^{db} + s2(d) \cdot X_5^{db} + Y_4^{db}$$

$$Y_4^{db} = r3(next) \cdot (A^{1-b} \ggg FQ \gg B^{db})$$

$$X_5^{db} = c(db) \cdot X_5^{db} + s4(next) \cdot X_6^{db} + Y_5^{db}$$

$$Y_5^{db} = r3(next) \cdot (A^{1-b} \ggg FQ \gg s4(next) \cdot B^{1-b})$$

$$X_6^{db} = r3(next) \cdot X_1^{1-b} + c(db) \cdot X_6^{db}$$

Using CFAR immediately gives the next lemma.

5.2.5. LEMMA: Let U be specified by:

$U = U_1^0$	
$U_1^b = \sum_{d \in D} r\,1(d) \cdot U_2^{db}$	
$U_2^{db} = \tau \cdot U_4^{db} + V_2^b + V_3^{db}$	$V_2^b = r\,3(next) \cdot (A^{1-b} \gg FQ \gg B^b)$
	$V_3^{db} = r\,3(next) \cdot (A^{1-b} \gg FQ \gg FBUF\,1^{db} \gg B^b)$
$U_4^{db} = s\,2(d) \cdot U_5^{db} + V_4^{db}$	$V_4^{db} = r\,3(next) \cdot (A^{1-b} \gg FQ \gg B^{db})$
$U_5^{db} = s\,4(next) \cdot U_6^{db} + V_5^{db}$	$V_5^{db} = r\,3(next) \cdot (A^{1-b} \gg FQ \gg s\,4(next) \cdot B^{1-b})$
$U_6^{db} = r\,3(next) \cdot U_1^{1-b}$	

Then: $SACP_{\tau}^{\#} + KFAR \vdash U = A \gg FQ \gg B$.

In the same way we can derive similar lemmas for ‘the other side’ of the protocol.

5.2.6. LEMMA:

- i) $FBUF\,1^b \gg D^{1-b} = \tau \cdot (FBUF\,1 \gg D^{1-b});$
- ii) $FBUF\,1^b \gg s\,3(next) \cdot D^{1-b} = \tau \cdot (FBUF\,1 \gg s\,3(next) \cdot D^{1-b});$
- iii) $FBUF\,1^b \gg FBUF\,1^b \gg D^b = \tau \cdot (FBUF\,1 \gg FBUF\,1^b \gg D^b).$

5.2.7. LEMMA: Let W be specified by:

$W = W_1^1$	
$W_1^b = \tau \cdot r\,4(next) \cdot W_2^{1-b}$	$Z_1^b = r\,4(next) \cdot (C^{1-b} \gg FQ \gg D^b)$
$W_2^b = \tau \cdot W_3^b + Z_1^b + Z_2^b$	$Z_2^b = r\,4(next) \cdot (C^{1-b} \gg FQ \gg FBUF\,1^b \gg D^b)$
$W_3^b = s\,3(next) \cdot W_1^b + Z_3^b$	$Z_3^b = r\,4(next) \cdot (C^{1-b} \gg FQ \gg s\,3(next) \cdot D^{1-b})$

Then: $SACP_{\tau}^{\#} + KFAR \vdash C \gg FQ \gg D = W$.

The fact that CABP is a correct protocol is asserted by

5.2.8. THEOREM: $SACP_{\tau}^{\#} + KFAR \vdash CABP = \tau \cdot (\sum_{d \in D} r\,1(d) \cdot s\,2(d)) \cdot CABP$.

PROOF: Lemmas 5.2.5 and 5.2.7 together give that we can write CABP as:

$$CABP = \tau_I \circ \partial_H(U \parallel W)$$

A straightforward expansion gives:

$$\tau_I \circ \partial_H(U \parallel W) = \tau \cdot \left(\sum_{d \in D} r \, 1(d) \cdot s \, 2(d) \right) \cdot \left(\sum_{e \in D} r \, 1(e) \cdot s \, 2(e) \right) \cdot \tau_I \circ \partial_H(U \parallel W)$$

The variables V and Z vanish in the expansion, due to the fact that they only occur in situations where a receiver component sends a premature acknowledgement. An application of RSP concludes the proof of the theorem. $\square$

5.2.9. Remark. A serious problem that has to be faced in the context of algebraic protocol verification is the fairness issue. In the verifications of this chapter we used KFAR to deal with fairness. KFAR is the algebraic equivalent of the statement: ‘if anything can go well infinitely often, it will go well infinitely often’. In most applications a more subtle treatment of fairness is desirable. Moreover KFAR is incompatible with lots of semantics between bisimulation and trace semantics. In [29] it is proved that failure semantics is inconsistent with the rule KFAR. In the same paper a restricted version KFAR^- of KFAR is presented which is consistent with the axioms of failure semantics, but this version is not powerful enough to allow for a verification of the CABP. The argument for this is simple: KFAR^- allows for the fair abstraction of *unstable divergence*. This means that a process will never stay forever in a conservative cluster of internal τ -steps if it can be exited by another internal τ -step. Since in the CABP component C can always perform an internal step, and since the protocol is finite state (after suitable abstraction), there must be a conservative cluster of internal steps which can only be exited by performing an observable action. Thus the CABP contains stable divergence.

CONCLUSIONS AND OPEN PROBLEMS

In this chapter we presented a language making it possible to give modular specifications of process algebras. The language contains operations $+$ and $\square$, which are standard in the theory of structured algebraic specifications, and moreover two new operators H and S . Two applications have been presented of the new operators: we showed how the left-merge operator can be hidden if this is needed and we described how the chaining operator can be defined in a clean way in terms of more elementary operators. It is clear that there are much more applications of our approach. Numerous other process combinators can be defined in terms of more elementary operators in the same way as we did with the chaining operators. Maybe also other model theoretic operations can be used in a process algebra setting (cartesian products?).

Strictly speaking we have not introduced a ‘module algebra’ as in [23]: we do not interpret module expressions in an algebra. However, this can be done without any problem. An interesting topic of research is to look for axioms to manipulate module expressions. Due to the presence of the operators H and S , an elimination theorem for module expressions as in [23] will probably not be achievable.

An important open problem for us is the question whether the proof system of Table 1 is complete for first order logic.

In this chapter the modules are parametrised by a set of actions. These actions themselves do not have any structure. The most natural way to look towards actions like $s\ 1(d_0)$ however, is to see them as actions parametrised by data. We would like to include the notion of a parametrised action in our framework but it turns out that this is not trivial. Related work in this area has been done by MAUW [129] and MAUW & VELTINK [130].

In order to prove the associativity of the chaining operators, we needed auxiliary actions $\bar{s}(d)$, $\bar{r}(d)$, etc. Also in other situations it often turns out to be useful to introduce auxiliary actions in verifications. At present we have to introduce these actions right at the beginning of a specification. This is embarrassing for a reader who does not know about the future use of these actions in the verification. But of course also the authors don't like to rewrite their specification all the time when they work on the verification. Therefore we would like to have a proof principle saying that it is allowed to use 'fresh' atomic actions in proofs. We think that it is possible to add a 'Fresh Atom Principle' (FAP) to our formal setting, but some work still has to be done.

In our view Section 4 convincingly shows that chaining operators are useful in dealing with FIFO-queues. We think that in general it will be often the case that a new application requires new operators and laws.

In Section 4.5 we presented a simple example of a realistic situation where bisimulation semantics does not work: a FIFO-queue which can loose data at every place is different from a FIFO-queue which can only loose data at the end. Adding the law T4, which holds in ready trace semantics (and hence in failure semantics), made it possible to prove the two queues equal.

For the correctness of protocols which involve faulty queues one normally needs some fairness assumption. Koomen's Fair Abstraction Rule (KFAR) often forms an adequate, although not optimal, way to model fairness. An interesting open problem is therefore the question whether the module $SACP_7^\# + T4 + KFAR$ is consistent (conjecture: yes).

The verification of the Concurrent Alternating Bit Protocol as presented here takes 4 pages (or 5 if the proofs of the standard facts about the queues are included). Our proof is considerably shorter than the proof of similar protocols in [123] and [127] (15 and 11 pages respectively). But maybe this comparison is not altogether fair because the proofs in these papers were meant as an illustration of new modular proof techniques. Our proof shows that the axioms of bisimulation semantics with fair abstraction are sufficient for the modular verification of simple protocols like this. The axioms of bisimulation semantics will turn out to be not sufficient for more substantial modular verifications because bisimulation semantics is not fully abstract. We could give a shorter and simpler proof of the protocol by using the notion of redundancy in context of [169]: the grey arcs in Figure 6 all correspond to summands which are redundant in the context in which they occur. Additional proof techniques will certainly be needed for the modular verification of more complex protocols.

ACKNOWLEDGEMENTS

Our thanks to Jan Bergstra for his help in the development of the H -operator and to Kees Middelburg for helpful comments on an earlier version.

APPENDIX: LOGICS

In this appendix equational, conditional equational and first order logic are defined. Since all these logics share the concepts of variables and terms, these will be treated first.

1. Variables and terms. Let σ be a signature. A σ -variable is an expression x_S with $x \in \text{NAMES}$ and $(S:S) \in \sigma$. A *valuation* of the σ -variables in a σ -algebra $\mathcal{A}$ is a function ξ that takes every σ -variable x_S into an element of $S^{\mathcal{A}}$.

For any $(S:S) \in \sigma$ the set T_S^{σ} of σ -terms of sort S is defined inductively by:

- $x_S \in T_S^{\sigma}$ for any σ -variable x_S .
- If $f: S_1 \times \dots \times S_n \rightarrow S$ is in σ and $t_i \in T_{S_i}^{\sigma}$ for $i = 1, \dots, n$ then $f_{S_1 \times \dots \times S_n \rightarrow S}(t_1, \dots, t_n) \in T_S^{\sigma}$.

The ξ -evaluation $\llbracket t \rrbracket^{\xi} \in S^{\mathcal{A}}$ of a σ -term $t \in T_S^{\sigma}$ in a σ -algebra $\mathcal{A}$ (with ξ a valuation) is defined by:

- $\llbracket x_S \rrbracket^{\xi} = \xi(x_S) \in S^{\mathcal{A}}$.
- $\llbracket f_{S_1 \times \dots \times S_n \rightarrow S}(t_1, \dots, t_n) \rrbracket^{\xi} = f_{S_1 \times \dots \times S_n \rightarrow S}^{\mathcal{A}}(\llbracket t_1 \rrbracket^{\xi}, \dots, \llbracket t_n \rrbracket^{\xi})$.

2. Equational logic. The set F_{σ}^{eq} of equations or equational formulas over σ is defined by:

- If $t_i \in T_{S_i}^{\sigma}$ for $i = 1, 2$ and certain $S:S$ in σ , then $(t_1 = t_2) \in F_{\sigma}^{eq}$.

An equation $(t_1 = t_2) \in F_{\sigma}^{eq}$ is ξ -true in a σ -algebra $\mathcal{A}$, notation $\mathcal{A}, \xi \models_{\sigma}^{eq} t_1 = t_2$, if $\llbracket t_1 \rrbracket^{\xi} = \llbracket t_2 \rrbracket^{\xi}$.

Such an equation $\phi \in F_{\sigma}^{eq}$ is true in $\mathcal{A}$, notation $\mathcal{A} \models_{\sigma}^{eq} \phi$, if $\mathcal{A}, \xi \models_{\sigma}^{eq} \phi$ for all valuations ξ .

An inference system I_{σ}^{eq} for equational logic is displayed in Table 12 below. There t, u and v are terms over σ and x is a variable. Furthermore $t[u/x]$ is the result of substituting u for all occurrences of x in t . Of course u and x should be of the same sort. Finally an inference rule $\frac{H}{\phi}$ with $H = \emptyset$ is called an *axiom* and denoted simply by ϕ .

$t = t$	$\frac{u = v}{v = u}$	$\frac{t = u, u = v}{t = v}$	$\frac{u = v}{t[u/x] = t[v/x]}$	$\frac{u = v}{u[t/x] = v[t/x]}$
---------	-----------------------	------------------------------	---------------------------------	---------------------------------

TABLE 12

3. *Conditional equational logic.* The set F_σ^{at} of *atomic formulas* over σ is defined by:

- If $t_i \in T_S^\sigma$ for $i = 1, 2$ and certain $S:S$ in σ , then $(t_1 = t_2) \in F_\sigma^{at}$.
- If $R:p \subseteq S_1 \times \dots \times S_n$ is in σ and $t_i \in T_{S_i}^\sigma$ for $i = 1, \dots, n$ then $p_{S_1 \times \dots \times S_n}(t_1, \dots, t_n) \in F_\sigma^{at}$.

The set F_σ^{ceql} of *conditional equational formulas* over σ is defined by:

- If $C \subseteq F_\sigma^{at}$ and $\alpha \in F_\sigma^{at}$ then $(C \Rightarrow \alpha) \in F_\sigma^{ceql}$.

The ξ -*truth* of formulas $\phi \in F_\sigma^{at} \cup F_\sigma^{ceql}$ in a σ -algebra $\mathcal{Q}$ is defined by:

- $\mathcal{Q}, \xi \models_\sigma^{ceql} t_1 = t_2$ if $\llbracket t_1 \rrbracket^\xi = \llbracket t_2 \rrbracket^\xi$.
- $\mathcal{Q}, \xi \models_\sigma^{ceql} p_{S_1 \times \dots \times S_n}(t_1, \dots, t_n)$ if $p_{S_1 \times \dots \times S_n}^\mathcal{Q}(\llbracket t_1 \rrbracket^\xi, \dots, \llbracket t_n \rrbracket^\xi)$.
- $\mathcal{Q}, \xi \models_\sigma^{ceql} C \Rightarrow \alpha$ if $\mathcal{Q}, \xi \not\models_\sigma^{ceql} \beta$ for some $\beta \in C$ or $\mathcal{Q}, \xi \models_\sigma^{ceql} \alpha$.

ϕ is *true* in $\mathcal{Q}$, notation $\mathcal{Q} \models_\sigma^{ceql} \phi$, if $\mathcal{Q}, \xi \models_\sigma^{ceql} \phi$ for all valuations ξ .

An inference system I_σ^{ceql} for conditional equational logic is displayed in Table 13 below. There α and α_i are atomic formulas, C is a set of atomic formulas, ϕ is a conditional equational formula, t_i, t, u and v are terms over σ and x_i and x are variables. Furthermore $\alpha[u/x]$ is the result of substituting u for all occurrences of x in α . Of course u and x should be of the same sort. Likewise $\phi[t_i/x_i (i \in I)]$ is the result of simultaneous substitution for $i \in I$ of t_i for all occurrences of x_i in ϕ . An inference rule $\frac{\phi}{\phi}$ is again denoted by ϕ and a conditional equational formula $\emptyset \Rightarrow \alpha$ by α .

$C \Rightarrow \alpha$ if $\alpha \in C$	$\frac{C \Rightarrow \alpha_i (i \in I), \{\alpha_i \mid i \in I\} \Rightarrow \alpha}{C \Rightarrow \alpha}$	$\frac{\phi}{\phi[t_i/x_i (i \in I)]}$
$t = t$	$\{u = v\} \Rightarrow (v = u)$	$\{t = u, u = v\} \Rightarrow (t = u)$
	$\{u = v, \alpha[u/x]\} \Rightarrow (\alpha[v/x])$	

TABLE 13

The logic described above is *infinitary conditional equational logic*. *Finitary conditional equational logic* is obtained by the extra requirement that in conditional equational formulas $C \Rightarrow \alpha$ the set of conditions C should be finite. In that case the inference rule

$$\frac{\phi}{\phi[t_i/x_i (i \in I)]} \quad \text{can be replaced by} \quad \frac{\phi}{\phi[t/x]}.$$

Furthermore *(in)finitary conditional logic* is obtained by omitting all reference to the equality predicate $=$.

4. *First order logic.* The set F_{σ}^{foleq} of *first order formulas with equality* over σ is defined by:

- If $t_i \in T_{S_i}^{\sigma}$ for $i = 1, 2$ and certain $S : S$ in σ , then $(t_1 = t_2) \in F_{\sigma}^{foleq}$.
- If $R : p \subseteq S_1 \times \dots \times S_n$ is in σ and $t_i \in T_{S_i}^{\sigma}$ for $i = 1, \dots, n$ then $p_{S_1 \times \dots \times S_n}(t_1, \dots, t_n) \in F_{\sigma}^{foleq}$.
- If $\phi \in F_{\sigma}^{foleq}$ then $\neg \phi \in F_{\sigma}^{foleq}$.
- If ϕ and $\psi \in F_{\sigma}^{foleq}$ then $(\phi \rightarrow \psi) \in F_{\sigma}^{foleq}$.
- If ϕ and $\psi \in F_{\sigma}^{foleq}$ then $(\phi \wedge \psi) \in F_{\sigma}^{foleq}$.
- If ϕ and $\psi \in F_{\sigma}^{foleq}$ then $(\phi \vee \psi) \in F_{\sigma}^{foleq}$.
- If ϕ and $\psi \in F_{\sigma}^{foleq}$ then $(\phi \leftrightarrow \psi) \in F_{\sigma}^{foleq}$.
- If x_S is a σ -variable and $\phi \in F_{\sigma}^{foleq}$ then $\forall x_S(\phi) \in F_{\sigma}^{foleq}$.
- If x_S is a σ -variable and $\phi \in F_{\sigma}^{foleq}$ then $\exists x_S(\phi) \in F_{\sigma}^{foleq}$.

The ξ -*truth* of a formula $\phi \in F_{\sigma}^{foleq}$ in a σ -algebra $\mathcal{Q}$ is defined inductively by:

- $\mathcal{Q}, \xi \models_{\sigma}^{foleq} t_1 = t_2$ if $\llbracket t_1 \rrbracket^{\xi} = \llbracket t_2 \rrbracket^{\xi}$.
- $\mathcal{Q}, \xi \models_{\sigma}^{foleq} p_{S_1 \times \dots \times S_n}(t_1, \dots, t_n)$ if $p_{S_1 \times \dots \times S_n}(\llbracket t_1 \rrbracket^{\xi}, \dots, \llbracket t_n \rrbracket^{\xi})$.
- $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \neg \phi$ if $\mathcal{Q}, \xi \not\models_{\sigma}^{foleq} \phi$.
- $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \phi \rightarrow \psi$ if $\mathcal{Q}, \xi \not\models_{\sigma}^{foleq} \phi$ or $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \psi$.
- $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \phi \wedge \psi$ if $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \phi$ and $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \psi$.
- $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \phi \vee \psi$ if $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \phi$ or $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \psi$.
- $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \phi \leftrightarrow \psi$ if $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \phi$ if and only if $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \psi$.
- $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \forall x_S(\phi)$ if $\mathcal{Q}, \xi' \models_{\sigma}^{foleq} \phi$ for all valuations ξ' with $\xi'(y_{S'}) = \xi(y_{S'})$ for all variables $y_{S'} \neq x_S$.
- $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \exists x_S(\phi)$ if $\mathcal{Q}, \xi' \models_{\sigma}^{foleq} \phi$ for some valuation ξ' with $\xi'(y_{S'}) = \xi(y_{S'})$ for all variables $y_{S'} \neq x_S$.

ϕ is *true* is $\mathcal{Q}$, notation $\mathcal{Q} \models_{\sigma}^{foleq} \phi$, if $\mathcal{Q}, \xi \models_{\sigma}^{foleq} \phi$ for all valuations ξ .

An inference system I_{σ}^{foleq} for first order logic with equality is displayed in Table 14 on the next page. There ϕ, ψ and ρ are elements of F_{σ}^{foleq} , α is an atomic formula (constructed by means of the first two clauses in the definition of F_{σ}^{foleq} only), t, u and v are terms over σ and x is a variable. An occurrence of a variable x in a formula ϕ is *bound* if it occurs in a subformula $\forall x(\psi)$ or $\exists x(\psi)$ of ϕ . Otherwise it is *free*. A variable is *free* in a formula ϕ if all its occurrences in ϕ are free. $\phi[t/x]$ denotes the result of substituting u for all free occurrences of x in t . Of course u and x should be of the same sort. Now t is *free for x in ϕ* if all free occurrences of variables in t remain free in $\phi[t/x]$. As before an inference rule $\frac{H}{\phi}$ with $H = \emptyset$ is called an *axiom* and denoted simply by ϕ .

$\frac{\phi, \phi \rightarrow \psi}{\psi}$ <i>modus ponens</i>	$\frac{\phi}{\forall x(\phi)}$ <i>generalisation</i>
$\phi \rightarrow (\psi \rightarrow \phi)$ $\{\phi \rightarrow (\psi \rightarrow \rho)\} \rightarrow \{(\phi \rightarrow \psi) \rightarrow (\phi \rightarrow \rho)\}$ $\{\forall x(\phi \rightarrow \psi)\} \rightarrow \{\phi \rightarrow \forall x(\psi)\}$, if x is not free in ϕ $(\neg \phi \rightarrow \phi) \rightarrow \phi$ $\neg \phi \rightarrow (\phi \rightarrow \psi)$ $\forall x(\phi) \rightarrow \phi[t/x]$, if t is free for x in ϕ	<i>deduction axioms</i> <i>axiom of the excluded middle</i> <i>axiom of contradiction</i> <i>axiom of specialisation</i>
$(\phi \wedge \psi) \rightarrow \phi$ $(\phi \wedge \psi) \rightarrow \psi$ $\phi \rightarrow \{\psi \rightarrow (\phi \wedge \psi)\}$	$\phi \rightarrow (\phi \vee \psi)$ $\psi \rightarrow (\phi \vee \psi)$ $(\phi \vee \psi) \rightarrow (\neg \phi \rightarrow \psi)$
$t = t$	$(\phi \leftrightarrow \psi) \rightarrow \{(\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)\}$ $\{(\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)\} \rightarrow (\phi \leftrightarrow \psi)$ $\exists x(\phi) \leftrightarrow \neg \forall x(\neg \phi)$ $(u = v) \rightarrow (v = u)$ $\{(t = u) \wedge (u = v)\} \rightarrow (t = v)$ $(u = v) \rightarrow (\alpha[u/x] \leftrightarrow \alpha[v/x])$

TABLE 14

First order logic is obtained from first order logic with equality by omitting all reference to $=$. It is also possible to present first order logic without the connectives $\wedge$, $\vee$ and $\leftrightarrow$ and the quantifier $\exists$, and introduce them as notational abbreviations. In that case the third block of Table 14 can be omitted.

5. *Expressiveness*. One can translate an equation $\alpha \in F_{\sigma}^{eq}$ by a (finitary) conditional equational formula $\emptyset \Rightarrow \alpha$ and a finitary conditional equational formula $\{\alpha_1, \dots, \alpha_n\} \Rightarrow \alpha$ into a first order formula $(\alpha_1 \wedge \dots \wedge \alpha_n) \rightarrow \alpha$. Using this translation we have $F_{\sigma}^{eq} \subset F_{\sigma}^{ceq} \subset F_{\sigma}^{foeq}$ and furthermore $\mathcal{Q} \models_{\sigma}^{eq} \phi \Leftrightarrow \mathcal{Q} \models_{\sigma}^{ceq} \phi$ for $\phi \in F_{\sigma}^{eq}$ and $\mathcal{Q} \models_{\sigma}^{ceq} \phi \Leftrightarrow \mathcal{Q} \models_{\sigma}^{foeq} \phi$ for $\phi \in F_{\sigma}^{ceq}$. This means that first order logic with equality is more expressive than equational logic and finitary conditional equational logic is somewhere in between. However first order logic with equality and infinitary conditional equational logic have incomparable expressive power.

6. *Completeness*. For all logics mentioned above the following completeness result is known to hold: $Alg(\sigma, T) \models_{\sigma}^c \phi \Rightarrow T \vdash_{\sigma}^c \phi$. The reverse direction also holds, since all these logics are obviously sound. As a corollary we have

$$T \vdash_{\sigma}^{eq} \phi \Leftrightarrow T \vdash_{\sigma}^{ceq} \phi \quad \text{for } \phi \in F_{\sigma}^{eq} \text{ and}$$

$$T \vdash_{\sigma}^{ceq} \phi \Leftrightarrow T \vdash_{\sigma}^{foeq} \phi \quad \text{for } \phi \in F_{\sigma}^{ceq}.$$

For this reason in most process algebra papers it is not made explicit which logic is used in verifications: the space needed for stating this could be saved, since the resulting notion of provability would be the same anyway. However, the situation changes when formulas are proved from modules. Equational logic and conditional equational logic are not complete anymore and for first

order logic with equality this is still an open problem (as far as we know). Here a logic $\mathcal{L}$ is complete if $M \models^{\mathcal{L}} \phi \Rightarrow M \vdash^{\mathcal{L}} \phi$. It is easily shown that

$$\begin{aligned} M \vdash^{eq} \phi &\Rightarrow M \vdash^{ceq} \phi && \text{for } \phi \in F_{\Sigma(M)}^{eq} \text{ and} \\ M \vdash^{ceq} \phi &\Rightarrow M \vdash^{foeq} \phi && \text{for } \phi \in F_{\Sigma(M)}^{ceq}, \end{aligned}$$

but the reverse directions do not hold. Thus we should state exactly in which logic our results are proved.

7. Notation. This chapter employs infinitary conditional equational logic. However, no proof trees are constructed; proofs are given in a slightly informal way, that allows a straightforward translation into formal proofs by the reader. Furthermore all type information given in the subscripts of variables, function and predicate symbols is omitted, since confusion about the correct types is almost impossible. Outside Section 1 and this appendix inference rules $\frac{H}{\phi}$ do not occur, but all conditional equational formulas $C \Rightarrow \alpha$ are written $\frac{C}{\alpha}$, as is usual. However, the suggested similarity between inference rules and conditional equational formulas is misleading: $\frac{H}{\phi}$ holds in an algebra $\mathcal{Q}$ if $(\mathcal{Q}, \xi \models \psi$ for all $\psi \in H$ and all valuations ξ) implies $(\mathcal{Q}, \xi \models \phi$ for all valuations ξ), while $\frac{C}{\alpha}$ holds in $\mathcal{Q}$ if for all valuations ξ : $(\mathcal{Q}, \xi \models \beta$ for all $\beta \in C$ implies $\mathcal{Q}, \xi \models \alpha$).

8. Positive and universal formulas. In equational logic all formulas are both positive and universal. In conditional equational logic all formulas are universal and the positive formulas are the atomic ones. In first order logic with equality the positive formulas are the ones without the connectives $\neg$, $\rightarrow$ and $\leftrightarrow$ and the universal ones are the formulas without quantifiers. Model theory (see for instance [140]) teaches us that a formula ϕ is preserved under homomorphisms (respectively subalgebras) iff there is a positive (respectively universal) formula ψ with $\vdash^{foeq} \psi \leftrightarrow \phi$.

Chapter III

Branching time and abstraction in bisimulation semantics

Rob van Glabbeek & Peter Weijland

In comparative concurrency semantics one usually distinguishes between *linear time* and *branching time* semantic equivalences. Milner's notion of *observation equivalence* is often mentioned as the standard example of a branching time equivalence. In this chapter we investigate whether observation equivalence really does respect the branching structure of processes, and find that in the presence of the unobservable action τ of CCS this is not the case.

Therefore the notion of *branching bisimulation equivalence* is introduced which strongly preserves the branching structure of processes, in the sense that it preserves computations together with the potentials in all intermediate states that are passed through, even if silent moves are involved. On closed CCS-terms branching bisimulation congruence can be completely axiomatized by the single axiom scheme:

$$a.(\tau.(y + z) + y) = a.(y + z)$$

(where a ranges over all actions) and the usual laws for strong congruence.

We also establish that for sequential processes observation equivalence is not preserved under refinement of actions, whereas branching bisimulation is.

For a large class of processes it turns out that branching bisimulation and observation equivalence are the same. As far as we know, all protocols that have been verified in the setting of observation equivalence happen to fit in this class, and hence are also valid in the stronger setting of branching bisimulation equivalence.

TABLE OF CONTENTS

Introduction	118
1. Branching and abstraction	120
2. Axioms	130
3. Branches and traces	143
4. Completeness proofs	149
5. Correspondence	158
6. Refinement of actions	159
7. Divergence	163
8. Modal Characterizations	164
Conclusion	167

INTRODUCTION

When comparing semantic equivalences for concurrency, it is common practice to distinguish between *linear time* and *branching time* equivalences (see for instance DE BAKKER, BERGSTRA, KLOP & MEYER [18], PNUELI [152]). In the former, a process is determined by its possible executions, whereas in the latter also the branching structure of processes is taken into account. The standard example of a linear time equivalence is *trace equivalence* as employed in HOARE [115]; the standard example of a branching time equivalence is *observation equivalence* or *bisimulation equivalence* as defined by MILNER [134] and PARK [145] (cf. MILNER [136-138]). Furthermore, there are several *decorated trace equivalences* in between (cf. Chapter I), preserving part of the branching structure of processes but for the rest resembling trace equivalence.

Originally, the most popular argument for employing branching time semantics was the fact that it allows a proper modelling of *deadlock behaviour*, whereas linear time semantics does not. However, this advantage is shared with the decorated trace semantics which have the additional advantage of only distinguishing between processes that can be told apart by some notion of *observation* or *testing*. The main criticism on observation equivalence - and branching time equivalences in general - is that it is not an observational equivalence in that sense: distinctions between processes are made that cannot be observed or tested, unless observers are equipped with extraordinary abilities like that of a copying facility together with the capability of global testing as in MILNER [134,135] or ABRAMSKY [1].

Nevertheless, branching time semantics is of fundamental importance in concurrency, exactly because it is independent of the precise nature of observability. Which one of the decorated trace equivalences provides a suitable modelling of observable behaviour depends in great extent on the tools an observer has to test processes. And in general a protocol verification in a particular decorated trace semantics, does not carry over to a setting in which observers are a bit more powerful. On the other hand, branching time semantics preserves the internal branching structure of processes and thus certainly their observable behaviour as far as it can be captured by decorated traces. A protocol verified in branching time semantics is automatically valid in each of the decorated trace semantics.

Probably one of the most important features in process algebra is that of abstraction, since it provides us with a mechanism to *hide* actions that are not observable, or not interesting for any other reason. By abstraction, some of the actions in a process are made *invisible* or *silent*. Consequently, any consecutive execution of hidden steps cannot be recognized since we simply do not 'see' anything happen.

Algebraically, in ACP_τ of BERGSTRA & KLOP [26] abstraction has the form of a renaming operator which renames actions into a *silent move* called τ . In MILNER's CCS [134] these silent moves result from synchronization. This new constant τ is introduced in the algebraic models as well: for instance in the *graph models* (cf. [26,134]) we find the existence of τ -edges, and so the question was how to find a satisfactory extension of the original definition of *bisimulation equivalence* that we had on process graphs without τ .

It turns out that there exist many possibilities for extending bisimulation equivalence to process graphs with τ -steps. One such possible extension is incorporated in Milner's notion of *observation equivalence*—called τ -bisimulation equivalence in [26]—, which resembles ordinary bisimulation, but permits arbitrary sequences of τ -steps to precede or follow corresponding atomic actions. A different notion, the so-called η -bisimulation, was suggested by BAETEN & VAN GLABBEEK [15], yielding a weaker set of abstraction axioms. In MILNER [135] another notion of observational equivalence was introduced which in this chapter is referred to as *delay bisimulation equivalence*. As we will show, the treatments of Milner and Baeten & Van Glabbeek fit into a natural structure of four possible variations of bisimulation equivalence involving silent steps. The structure is completed by defining *branching bisimulation equivalence*. As it turns out, observation equivalence is the coarsest equivalence of the four, in the sense of identifying more processes. η - and delay bisimulation equivalence are two incomparable finer notions whereas branching bisimulation equivalence is the finest of all.

In a certain sense the usual notion of observation equivalence does not preserve the branching structure of a process. For instance, the processes $a(\tau b + c)$ and $a(\tau b + c) + ab$ are observation equivalent. However, in the first term, in each computation the choice between b and c is made after the a -step, whereas the second term has a computation in which b is already chosen when the a -step occurs. For this reason one may wonder whether or not we should accept the so-called third τ -law— $a(\tau x + y) = a(\tau x + y) + ax$ —that is responsible for the former equivalence. Similarly, the processes $\tau a + b$ and $\tau a + a + b$ are observation equivalent. However, only in the first term every computation in which a occurs passes through a state where a did not yet happen but the possibility to do b instead is already discarded. Hence we argue that also the second τ -law— $\tau x = \tau x + x$ —(responsible for the latter equivalence) does not respect branching time.

These examples shows us that while preserving observation equivalence, we can introduce new paths in a graph that were not there before. To be precise: the *traces* are the same, but the sequences of intermediate nodes are different (modulo observation equivalence), since in the definition of observation equivalence there is no restriction whatsoever on the nature of the nodes that are passed through during the execution of a sequence of τ -steps, preceding or following corresponding atomic actions. This is the key point in our definition of branching bisimulation equivalence: in two bisimilar processes every computation in the one process corresponds to a computation in the other, in such a way that all intermediate states of these computations correspond as well, due to the bisimulation relation. It turns out that it can be defined by a small change in the definition of observation equivalence.

The fact that observation equivalence can be too coarse in its identifications is illustrated even more strongly by the problems that it may cause in practical applications and analysis. As an example, it can be shown (cf. GRAF & SIFAKIS [102]) that there is no modal logic with eventually operator $\diamond$ which is adequate for observation equivalence. Here $\diamond\phi$ means that all paths will sooner or later reach a state where ϕ holds. Indeed, suppose that such a logic *would* exist, then this means that two

processes are observation equivalent iff they satisfy the same modal formulas. For instance, there exists a formula ϕ such that: $(\tau b + c) \models \phi$ and $b \not\models \phi$ since obviously these processes are not observation equivalent. However, from $(\tau b + c) \models \phi$ it follows that we have $a(\tau b + c) \models \blacklozenge \phi$ whereas from $b \not\models \phi$ we find $a(\tau b + c) + ab \not\models \blacklozenge \phi$ although both processes are observation equivalent. Obviously, this inconsistency is due to the third τ -law. Similarly there must be a formula ψ such that $a \models \psi$ and $\tau a + a + ba \not\models \psi$. Thus $\tau a + ba \models \blacklozenge \psi$, whereas $\tau a + a + ba \not\models \blacklozenge \psi$, although both processes are observation equivalent. This time the inconsistency is due to the second τ -law.

A paper by JONSSON & PARROW [119] on deciding bisimulation equivalence shows a different kind of struggle with the third τ -law. In this paper, infinite data flow is turned into a finite state representation by considering symbolic transitions. This provides us with a method to decide on the equivalence of infinite data flow programs. It turns out to work easily for strong equivalence, but for observation equivalence there is no straightforward generalization of the former results and a less intuitive transition system is needed to fix this problem. It is easy to see that using branching bisimulation would serve as a key to a more natural solution of this problem.

Having at least four options for the definition of bisimulation congruence involving τ -steps, in any particular application it becomes important to have a clear intuition about which kind of abstraction is preferable. In an important class of problems one can prove however, that all four notions of bisimulation yield the same equivalence. In particular this is the case if one of the two bisimulating graphs does not have any τ -steps. It is interesting to observe that, as far as we know, all case studies on protocol verification performed so far fit into this class of problems, hence all of their proofs that have been given in the setting of observation equivalence still hold in branching bisimulation semantics.

1. BRANCHING AND ABSTRACTION

In this section we define the semantic equivalences that we want to discuss on a domain of *process graphs*. Since we focus on branching and abstraction, we have chosen to abstain from a proper modelling of divergence, concurrency, real-time behaviour and stochastic aspects of processes. Moreover, we will disregard the nature of the actions that our processes may perform: they will be modelled as uninterpreted symbols $a, b, c, \dots$ from a given set Act . We have chosen process graphs (or labelled transition systems) to represent processes, since they clearly visualize the aspects of the modelled systems' behaviour we are interested in. The nodes in our graphs (or states in our transition systems) remain anonymous. A common alternative is to use closed expressions in a system description language like CCS or ACP as nodes in process graphs, but here we prefer to separate the semantic issues from the treatment of a particular language. In the next section, however, we will give an interpretation of certain subsets of CCS and ACP in (parts of) the graph model and discuss the algebraic aspects of our equivalences.

DEFINITION 1.1 A *process graph* is a connected, rooted, edge-labelled, directed graph.

In an edge-labelled directed graph, edges go from one node to another (or the same) node and are labelled with elements from a certain set Act . One can have more than one edge between two nodes as long as they carry different labels. A rooted graph has one special node which is indicated as the root node. We require process graphs to be connected: they need not be finite, but one must be able to reach every node from the root node by following a finite path. If r and s are nodes in a graph, then $r \rightarrow^a s$ denotes an edge from r to s with label a or it will be used as a *proposition* saying that such an edge exists. Process graphs represent concurrent systems in the following way: the elements of Act are *actions* a system may perform; the nodes of a process graph represent the states of a concurrent system; the root is the initial state and if $r \rightarrow^a s$, then the system can evolve from state r to state s by performing an action a . The domain of process graphs will be denoted by G .

On G we consider the notion of *bisimulation equivalence*, which originally was due to PARK [145] and used in MILNER [136-138] and in a different formulation already in MILNER [134]. On the domain of process graphs, a bisimulation usually is defined as a relation $R \subseteq \text{nodes}(g) \times \text{nodes}(h)$ on the nodes of graphs g and h satisfying:

- i. The roots of g and h are related by R
- ii. If $R(r, s)$ and $r \rightarrow^a r'$, then there exists a node s' such that $s \rightarrow^a s'$ and $R(r', s')$
- iii. If $R(r, s)$ and $s \rightarrow^a s'$, then there exists a node r' such that $r \rightarrow^a r'$ and $R(r', s')$.

Equivalently—as is done in this chapter—one can obtain bisimulation equivalence from a *symmetric* relation R between nodes of g and h , only satisfying (i) and (ii). Such a symmetric relation can be defined as a relation $R \subseteq \text{nodes}(g) \times \text{nodes}(h) \cup \text{nodes}(h) \times \text{nodes}(g)$ such that $R(r, s) \Leftrightarrow R(s, r)$, or alternatively, as a set of unordered pairs of nodes $R \subseteq \{\{r, s\} : r \in \text{nodes}(g), s \in \text{nodes}(h)\}$. In the latter case $R(r, s)$ abbreviates $\{r, s\} \in R$. Note that this restriction to symmetric relations does not cause any loss of generality.

DEFINITION 1.2 Two graphs g and h in G are *bisimilar*—notation: $g \approx h$ —if there exists a symmetric relation R (called a *bisimulation*) between their nodes such that:

- i. The roots of g and h are related by R
- ii. If $R(r, s)$ and $r \rightarrow^a r'$, then there exists a node s' such that $s \rightarrow^a s'$ and $R(r', s')$

Bisimilarity turns out to be an equivalence relation on G which is called *bisimulation equivalence*. Depending on the context we will sometimes use Milner's terminology and refer to bisimulation equivalence as *strong equivalence* or *strong congruence*.

Now let us postulate the existence of a special action $\tau \in \text{Act}$, that represents an unobservable, internal move of a process. We write $r \Rightarrow s$ for a path from r to s consisting of an arbitrary number (≥ 0) of τ -steps.

The definition of strong congruence was the starting point of MILNER [134] when he considered abstraction in CCS. Having in mind that τ -steps are not observable, he suggested to simply require that for g and h to be equivalent, (i) every possible a -step ($a \neq \tau$) in the one graph should correspond with an a -step in the other (as for usual bisimulation equivalence), apart from some arbitrary long sequences of τ -steps that are allowed to precede or follow, and (ii) every τ -step should correspond to an arbitrary long (≥ 0) τ -sequence. This way he obtained his notion of *observation equivalence* (cf.

MILNER [134,137,138])—or τ -bisimulation equivalence—which can be defined as follows:

DEFINITION 1.3 Two graphs g and h are τ -bisimilar—notation: $g \equiv_{\tau} h$ —if there exists a symmetric relation R (called a τ -bisimulation) between their nodes such that:

- i. The roots are related by R
- ii. If $R(r,s)$ and $r \rightarrow^a r'$, then either $a=\tau$ and $R(r',s)$, or there exists a path $s \Rightarrow s_1 \rightarrow^a s_2 \Rightarrow s'$ such that $R(r',s')$.

Again, $\equiv_{\tau}$ is an equivalence on G which is called τ -bisimulation equivalence, also known as *observation equivalence* or *weak (bisimulation) equivalence*.


For $\sigma=a_1a_2\cdots a_n$ a sequence of visible actions let $p \Rightarrow^{\sigma} q$ denote $p \Rightarrow \rightarrow^{a_1} \Rightarrow \rightarrow^{a_2} \Rightarrow \cdots \Rightarrow \rightarrow^{a_1} \Rightarrow q$, i.e. a path from p to q passing through a sequence of actions that reduces to σ after leaving out the internal ones. Then a τ -bisimulation can equivalently be defined as a symmetric relation between the nodes of graphs g and h , such that:

- i. The roots are related by R
- ii. If $R(r,s)$ and $r \Rightarrow^{\sigma} r'$, then there exists a path $s \Rightarrow^{\sigma} s'$ such that $R(r',s')$.


Although less useful for purposes of verification, in this (original) form the τ -bisimulation shows a clear parallel with the τ -less version.

Still, to some extent, the notion of τ -bisimulation cannot be regarded as the natural generalization of ordinary bisimulation to an abstract setting with hidden steps. The reason for this is that an important feature of a bisimulation is missing for τ -bisimulation, namely the property that any computation in the one process corresponds to a computation in the other, in such a way that all intermediate states of these computations correspond as well, due to the bisimulation relation. When HENNESSY & MILNER [111] introduced the first version of observation equivalence, they also insisted on relating the intermediate states of computations, as they tell us: "... any satisfactory comparison of the behaviour of concurrent programs must take into account their intermediate states as they progress through a computation, because differing intermediate states can be exploited in different program contexts to produce different overall behaviour ..." and: "If we consider a computation as a sequence of experiments (or communications), then the above remarks show that the intermediate states are compared. In fact, if p is to be equivalent to q , there must be a strong relationship between their respective intermediate states. At each intermediate stage in the computations, the respective "potentials" must also be the same". However, in Milner's observation equivalence, when satisfying the second requirement of Definition 1.3 one may execute arbitrary many τ -steps in a graph without worrying about the status of the nodes that are passed through in the meantime.¹


1. The quoted text was actually intended to motivate a non-functional semantics for reactive processes. We also note that observational equivalence *does* take the intermediate states of related processes into account, although not on the level of matching computations. Here we merely take the same ideas that motivated Hennessy and Milner to an extreme. In the conclusion we will show how the different intermediate states of corresponding computations of observationally equivalent processes can be exploited in different program contexts to produce different overall behaviour.



(a)



(b)



(c)

FIGURE 1. Observation equivalence.

As an illustration, in Figure 1 we consider a path $a \cdot \tau \cdot b \cdot \tau \cdot c$ with outgoing edges $d_1, \dots, d_4$, and it follows easily that all three graphs are observation equivalent. Note that one may add extra b -edges as in (b) and (c) without disturbing equivalence. However, in both (b) and (c) a new computation path is introduced—in which the outgoing edge d_2 (or d_3 respectively) is missing—and such a path did not occur in (a). Or—to put it differently—in the path introduced in (b) the options d_1 and d_2 are discarded simultaneously, whereas in (a) it corresponds to a path containing a state

where the option d_1 is already discarded but d_2 is still possible. Also in the path introduced in (c) the choice not to perform d_3 is already made with the execution of the b-step, whereas in (a) it corresponds to a path in which this choice is made only after the b-step. Thus we argue that observation equivalence does not preserve the branching structure of processes and hence lacks one of the main characteristics of bisimulation semantics.

Consider the following alternative definition of bisimulation in order to see how we can overcome this deficit.

DEFINITION 1.4 Two graphs g and h are *branching bisimilar*—notation: $g \approx_b h$ —if there exists a symmetric relation R (called a *branching bisimulation*) between the nodes of g and h such that:

- i. The roots are related by R
- ii. If $R(r, s)$ and $r \xrightarrow{a} r'$, then either $a = \tau$ and $R(r', s)$, or there exists a path $s \Rightarrow s_1 \xrightarrow{a} s_2 \Rightarrow s'$ such that $R(r, s_1)$, $R(r', s_2)$ and $R(r', s')$.

In a picture, the difference between branching and τ -bisimulation can be characterized as follows:




FIGURE 2. Bisimulations with τ .

The double arrow corresponds to the symbol $\Rightarrow$. Ordinary τ -bisimulation (Definition 1.3) says that every a -step $r \xrightarrow{a} r'$ corresponds with a path $s \Rightarrow s_1 \xrightarrow{a} s_2 \Rightarrow s'$ and so we obtain Figure 2 *without* the lines marked with (1) and (2). Branching bisimulation moreover requires relations between r and s_1 and between r' and s_2 and thus we obtain Figure 2 *with* (1) and (2). Note that if $g \approx_b h$ then there exists a *largest* branching bisimulation between g and h , since the set of branching bisimulations is closed under arbitrary union. One can easily check that branching bisimilarity is an equivalence on G , referred to as *branching bisimulation equivalence* or *branching equivalence* for short.

Obviously, branching equivalence more strongly preserves the branching structure of a graph since the starting and endnodes of the τ -paths $s \Rightarrow s_1$ and $s_2 \Rightarrow s$ are related to the same nodes. Observe that in Figure 1 there are no branching bisimulations

between any of the graphs (a), (b) and (c). In particular, adding extra edges as in (b) and (c) no longer preserves branching equivalence. Equivalently, we could have strengthened Definition 1.3 (ii) by requiring *all* intermediate nodes in $s \Rightarrow s_1$ and $s_2 \Rightarrow s$ to be related with r and r' respectively. In fact, this would yield the notion we really want to define. That this alternative definition yields the same equivalence relation can be seen by use of the following lemma:

LEMMA 1.1 (stuttering lemma)

Let R be the largest branching bisimulation between g and h .

If $r \xrightarrow{\tau} r_1 \xrightarrow{\tau} \dots \xrightarrow{\tau} r_m \xrightarrow{\tau} r'$ ($m \geq 0$) is a path such that $R(r,s)$ and $R(r',s)$ then $\forall 1 \leq i \leq m: R(r_i,s)$.

PROOF First we prove Lemma 1.1 for a slightly different kind of bisimulation, defined as follows:

DEFINITION A *semi branching bisimulation* between two graphs g and h is a symmetric relation R between the nodes of g and h such that:

- i. The roots are related by R
- ii. If $R(v,w)$ and $v \xrightarrow{a} v'$ then either
 - (a) $a = \tau$ and there exists a path $w \Rightarrow w'$ such that $R(v,w')$ and $R(v',w')$, or:
 - (b) there exists a path $w \Rightarrow w_1 \xrightarrow{a} w_2 \Rightarrow w'$ such that $R(v,w_1)$, $R(v',w_2)$ and $R(v',w')$.

The difference with branching bisimulation is in case (a), which can be illustrated by:



FIGURE 3. Semi branching (left) and branching bisimulation.

Now let $(*)$ denote the property, mentioned in the lemma. Observe that (a) any branching bisimulation is a semi branching bisimulation and (b) any semi branching bisimulation satisfying $(*)$ is a branching bisimulation.

CLAIM The largest semi branching bisimulation between g and h satisfies $(*)$.

Let R be the largest semi branching bisimulation between g and h , let s be a node and let $r \xrightarrow{\tau} r_1 \xrightarrow{\tau} \dots \xrightarrow{\tau} r_m \xrightarrow{\tau} r'$ ($m \geq 0$) be a path such that $R(r,s)$ and $R(r',s)$. Then we prove that $R' = R \cup \{ \{r_i, s\} : 1 \leq i \leq m \}$ is a semi branching bisimulation. We check the conditions:

- (i) Clearly, the root nodes of g and h are related by R' (since by R).

(ii) Suppose $R'(v, w)$ and $v \rightarrow^a v'$. If $R(v, w)$ then it follows that either (a) $a=\tau$ and there exists a path $w \Rightarrow w'$ such that $R(v, w')$ and $R(v', w')$, or (b) there exists a path $w \Rightarrow w_1 \rightarrow^a w_2 \Rightarrow w'$ such that $R(v, w_1)$, $R(v', w_2)$ and $R(v', w')$. Hence, from $R \subseteq R'$ we find that R' satisfies the requirements in the definition above.

So assume *not* $R(v, w)$, then we find that either (1) $v=s$ and $w=r_i$ or (2) $v=r_i$ and $w=s$.

(1) If $s \rightarrow^a s'$ then it follows from $R(r', s)$ that

either: $a=\tau$ and there is a path $r' \Rightarrow r''$ such that $R(r'', s)$ and $R(r'', s')$. Hence there is a path $r_i \Rightarrow r' \Rightarrow r''$ such that $R'(r'', s)$ and $R'(r'', s')$ as required.

or: there is a path $r' \Rightarrow t_1 \rightarrow^a t_2 \Rightarrow r''$ such that $R(t_1, s)$, $R(t_2, s')$ and $R(r'', s')$ and hence $r_i \Rightarrow r' \Rightarrow t_1 \rightarrow^a t_2 \Rightarrow r''$ with $R'(t_1, s)$, $R'(t_2, s')$ and $R'(r'', s')$.

(2) If $r_i \rightarrow^a r''$ then $r \rightarrow^\tau r_1 \rightarrow^\tau \dots \rightarrow^\tau r_i \rightarrow^a r''$ and since $R(r, s)$ we find that there exists a sequence $s \Rightarrow s_1 \Rightarrow \dots \Rightarrow s_i$ such that $R(r_1, s_1), \dots, R(r_i, s_i)$. It follows from $R(r_i, s_i)$ that

either: $a=\tau$ and there exists a path $s_i \Rightarrow s'$ such that $R(r_i, s')$ and $R(r'', s')$. Hence $s \Rightarrow s'$ with $R'(r_i, s')$ and $R'(r'', s')$ as required.

or: there exists a path $s_i \Rightarrow t_1 \rightarrow^a t_2 \Rightarrow s''$ such that $R(r_i, t_1)$, $R(r'', t_2)$ and $R(r'', s'')$, and hence $s \Rightarrow s_i \Rightarrow t_1 \rightarrow^a t_2 \Rightarrow s''$ with $R'(r_i, t_1)$, $R'(r'', t_2)$ and $R'(r'', s'')$.

This proves that R' is a semi branching bisimulation. Since R is the largest we find $R=R'$.

So we proved the claim. Finally, conclude that the largest semi branching bisimulation is equal to the largest branching bisimulation, and thus we proved the lemma. $\square$

The stuttering lemma will play a crucial role in some of the results we will present later. An alternative proof, presented in a somewhat different setting, can be found in DE NICOLA, MONTANARI & VAANDRAGER [66].

It follows from Figure 2 that we can find two more kinds of bisimulation with τ , since we can leave out (1) while still having (2) and vice versa. Consider the following two definitions:

DEFINITION 1.5 Two graphs g and h are η -bisimilar—notation $g \rightleftharpoons_\eta h$ —if there exists a symmetric relation R (called an η -bisimulation) between their nodes such that:

- i. The roots are related by R
- ii. If $R(r, s)$ and $r \rightarrow^a r'$, then either $a=\tau$ and $R(r', s)$, or there exists a path $s' \Rightarrow s_1 \rightarrow^a s_2 \Rightarrow s'$ such that $R(r, s_1)$ and $R(r', s')$.

DEFINITION 1.6 Two graphs g and h are *delay bisimilar*— $g \rightleftharpoons_d h$ —if there exists a symmetric relation R (called a *delay bisimulation*) between their nodes such that:

- i. The roots are related by R
- ii. If $R(r, s)$ and $r \rightarrow^a r'$, then either $a=\tau$ and $R(r', s)$, or there exists a path $s \Rightarrow s_1 \rightarrow^a s_2 \Rightarrow s'$ such that $R(r', s_2)$ and $R(r', s')$.

Notice the subtle differences between both definitions (and Definition 1.4). In Definition 1.5 the notion of η -bisimulation corresponds to Figure 2 without the relation (2) but with (1). Similarly, with delay bisimulation we have (2) but not (1). It is easy to see that in the definition of both branching and delay bisimulation the existence requirement of a node s' such that $s_2 \Rightarrow s'$ and $R(r', s')$ is redundant.

From the definitions we find immediately that $g \simeq_b h \Rightarrow g \simeq_\eta h \Rightarrow g \simeq_\tau h$ and similarly $g \simeq_b h \Rightarrow g \simeq_d h \Rightarrow g \simeq_\tau h$. Observe that in Figure 1 we find an η -bisimulation between (a) and (c) and a delay bisimulation between (a) and (b). Conversely, there is no η -bisimulation between (a) and (b) and no delay bisimulation between (a) and (c), so all implications are strict.

The notion of η -bisimulation was first introduced by BAETEN & VAN GLABBEK [15] as a finer version of observation equivalence. A variant of delay bisimulation—only differing in the treatment of divergence—first appeared in MILNER [135], also under the name observational equivalence.

HISTORICAL NOTE:

The first semantic equivalence preserving the branching structure of processes was defined in HENNESSY & MILNER [111] and MILNER [134]. In [134] it was called *strong equivalence* or *strong congruence*. It was defined in terms of a decreasing sequence $\sim_0, \sim_1, \dots, \sim_k, \dots$ of equivalence relations. Originally, these relations were defined on CCS expressions that figured as states in transition systems, but one can also define them on nodes of (possibly different) process graphs.

DEFINITION 1.7 Let r and s be nodes of process graphs. Then:

$r \sim_0 s$ is always true

$r \sim_{k+1} s$ iff for all $a \in \text{Act}$

- (i) if $r \xrightarrow{a} r'$ then there exists a node s' such that $s \xrightarrow{a} s'$ and $r' \sim_k s'$
- (ii) if $s \xrightarrow{a} s'$ then there exists a node r' such that $r \xrightarrow{a} r'$ and $r' \sim_k s'$

$r \sim s$ iff for all $k \in \mathbb{N}$: $r \sim_k s$.

Two graphs g and h are *strongly equivalent*, notation $g \sim h$, if $\text{root}(g) \sim \text{root}(h)$.

A process graph is *finitely branching* if each node has only finitely many outgoing edges. In HENNESSY & MILNER [111] and MILNER [134] strong congruence was defined only on CCS expressions corresponding with finitely branching graphs. On this domain, as was shown in [111] and [134], strong congruence 'satisfies its definition' in the following sense:

PROPOSITION 1.2 Let r and s be nodes of finitely branching process graphs.

Then $r \sim s$ iff for all $a \in \text{Act}$:

- i. if $r \xrightarrow{a} r'$ then there exists a node s' such that $s \xrightarrow{a} s'$ and $r' \sim s'$
- ii. if $s \xrightarrow{a} s'$ then there exists a node r' such that $r \xrightarrow{a} r'$ and $r' \sim s'$.

Strong equivalence is closely related to the notion of *bisimulation*, introduced by PARK [145] (cf. Definition 1.2). It is easy to verify that any bisimulation is included in each of the relations $\sim_k$ for $k \in \mathbb{N}$. Hence bisimulation equivalence is at least as discriminating as strong equivalence. On the other hand, from the former proposition it follows that with respect to finitely branching process graphs strong equivalence is a bisimulation, and hence the two notions coincide. With respect to infinitely branching graphs, $\sim$ is strictly coarser than bisimulation equivalence as can be seen from the following example. Consider the graphs




FIGURE 4. 'Strongly equivalent' graphs that are not bisimilar.

One can easily verify that these graphs are strongly equivalent in the sense of definition 1.7, but not bisimilar.

PROPOSITION 1.3 (SANDERSON [160])

- i. With respect to finitely branching process graphs the notions $\sim$ and $\cong$ coincide;
- ii. With respect to infinitely branching process graphs $\cong$ is strictly contained in $\sim$.

Starting from this observation, there are two different ways in which the notion of strong equivalence (in HENNESSY & MILNER [111] and MILNER [134] defined for finitely branching processes only) can be extended to infinitely branching process graphs. In MILNER [136] strong equivalence is chosen to be the relation of Definition 1.2, so "strong equivalence" and "bisimulation equivalence" are synonyms.

In the presence of a special action τ , representing an unobservable move of a process, one looks for a semantic equivalence that abstracts from internal moves in a process and for the rest resembles bisimulation equivalence. Such an abstract equivalence has to satisfy requirements such as:

- it is coarser than (strong) bisimulation equivalence
- it is equal to bisimulation equivalence with respect to processes not containing τ -edges
- it does not discriminate between the graphs





FIGURE 5. Contraction of internal moves.

The definition of strong congruence ($\sim$) was the starting point of HENNESSY & MILNER [111] when they introduced abstraction in CCS. Having in mind that τ -steps are not observable, they proposed that two process graph g and h are equivalent if every visible step in the one graph corresponds with a similar step in the other, apart from some arbitrary long sequences of τ -steps that are allowed to precede or follow. This way they obtained a notion of *observational equivalence*. Originally, this relation was defined in the style of Definition 1.7, but in order to facilitate comparison with the other equivalences, we will present it in bisimulation style.

DEFINITION 1.8 Two graphs g and h are *observational equivalent* in the sense of HENNESSY & MILNER if there exists a symmetric relation R between the nodes of g and h such that:

- The roots are related by R
- If $R(r, s)$ and $r \xrightarrow{a} r'$ ($a \neq \tau$), then there is a path $s \Rightarrow s_1 \xrightarrow{a} s_2 \Rightarrow s'$ with $R(r', s')$.

This equivalence does not coincide with the one from Definition 1.3. As an example, consider the following two process graphs:

FIGURE 6. Observational equivalent in the sense of HENNESSY & MILNER [111], but not τ -bisimilar.

These graphs are observational equivalent in the sense of HENNESSY & MILNER [111]; the relation R has been indicated in the figure above. But they are distinguished by MILNER [134] (Definition 1.3).

Unfortunately, this type of observational equivalence turned out not to be a congruence for the CCS parallel composition operator, the free merge, or any other operator representing concurrent activity (cf. HENNESSY & MILNER [111]). Hence Milner's version of observation equivalence [134], which we call τ -bisimulation equivalence, can be regarded as an improvement. Both notions satisfy the requirements mentioned above, but additionally τ -bisimulation equivalence is a congruence for the CCS parallel composition operator. Since observational equivalence in the sense of HENNESSY & MILNER [111] is coarser than τ -bisimulation equivalence, the criticism that τ -bisimulation equivalence does not preserve the branching structure of processes also applies to the variant of [111].

2. AXIOMS

In this section we will turn several parts of our graph domain G into algebras, by defining some operations on them. This will enable us to give equational characterizations of the equivalences studied in the previous section. In the first subsection we use the operators of the axiom system BPA_τ (cf. BERGSTRA & KLOP [26]): action constants, alternative and sequential composition. In the second subsection we take the operators inaction, prefixing and alternative composition of CCS (cf. MILNER [134]). Finally, in the third subsection we combine the features of the previous two approaches, thereby obtaining the kernel of the extended algebra ACP_τ (cf. BERGSTRA & KLOP [26]). We will not consider parallel composition, restriction (or encapsulation), hiding and relabelling. However, we claim that these can be added without problem.

2.1. BASIC PROCESS ALGEBRA

For sake of convenience, in this subsection we will only consider *root unwound* process graphs, i.e. process graphs with no incoming edges at the root. Since each bisimulation equivalence class of process graphs contains a root unwound graph, this does not cause any loss of generality. Furthermore, we restrict ourselves to *non-trivial graphs*—having at least one edge—and we assume our graphs to be *divergence free*, meaning that they do not contain infinite τ -paths. The latter restriction will be cancelled later, but for the time being it suits us since having it we can stay closer to CCS in our presentation, as will be explained in Section 2.3. (NOTE: apart from arguments about presentation, one may argue that there is still discussion about the role of divergence in bisimulation equivalence on processes, such as the dichotomy between explicit divergence (MILNER [135], WALKER [181]) and fair abstraction (MILNER [134], BAETEN, BERGSTRA & KLOP [13]), see also Section 7). The domain of root unwound, non-trivial and divergence free process graphs will be denoted by G_{BPA} . Clearly $G_{BPA} \subseteq G$.

In order to equip G_{BPA} with some structure, we introduce two binary infix written operators $+$ and $\cdot$ and constants for every element in Act .

DEFINITION 2.1 The constants $a \in Act$ and the operators $+$ and $\cdot$ are defined on G_{BPA} as follows:

- (i) Constants $a \in Act$ are interpreted by one-edge graphs labelled by a
- (ii) $(g + h)$ can be constructed by identifying the root nodes of g and h
- (iii) $(g \cdot h)$ is constructed by identifying all endnodes (leaves) in g with the root of h .
If g is without endnodes, then the result is just g .

Formally, identifying is defined as dividing out the equivalence relation 'to be identified' on the disjoint union of the nodes of g and h . The convention applies that any pointer (such as a name) to a node of g or h is also used to refer to its equivalence class, this being a node of $g+h$ or $g \cdot h$. In particular, if "R" is the name of a relation between the nodes of g and g' , then "R" is also the name of a relation between the nodes of $g \cdot h$ and g' , or $g \cdot h$ and $g' \cdot h'$. Which of these relations is meant by a given occurrence of "R" should be determined by the context.

As in regular algebra we will often leave out brackets and $\cdot$, assuming that $\cdot$ will always bind stronger than $+$.

The operators $+$ and $\cdot$ are well-defined, even after dividing out bisimulation equivalence on G_{BPA} , as follows from the following proposition, the proof of which is straightforward and omitted.

PROPOSITION 2.1 *Bisimulation equivalence is a congruence with respect to the operators $+$ and $\cdot$.*

Hence the structure $(G_{BPA}/\approx, +, \cdot, Act)$ is a well-defined algebra. Considering its first order equational theory we find the axiom system BPA (cf. BERGSTRA & KLOP [25]) which stands for *Basic Process Algebra*.

$x + y = y + x$	A1
$(x + y) + z = x + (y + z)$	A2
$x + x = x$	A3
$(x + y)z = xz + yz$	A4
$(xy)z = x(yz)$	A5

Table 1. BPA.

Now let us say that a theory Γ is a *complete axiomatization* of a model M if for every pair of closed terms p and q we have: $\Gamma \vdash p=q$ if and only if $M \models p=q$. This definition deviates from the standard one, since usually also open terms are considered. Then the following theorem is due to BERGSTRA & KLOP [26]:

THEOREM 2.2 *BPA is a complete axiomatization of $(G_{BPA}/\equiv, +, \cdot, \text{Act})$.*

Observe that in the presence of the trivial graph, BPA is not sound with respect to $(G_{BPA}/\equiv, +, \cdot, \text{Act})$: Axiom A4 no longer holds, with the trivial graph substituted for the variable y . For this reason it was excluded from G_{BPA} from the beginning.

In the same way one may wish to find axiomatizations for algebras resulting from dividing out the other equivalences of Section 1. However, as it turns out these equivalences are not congruences with respect to the operator $+$. In the case of observation equivalence this problem was solved by MILNER [134] by simply taking the closure of $\equiv_\tau$ with respect to all contexts in CCS, thereby obtaining *observation congruence*. Similarly in HENNESSY & MILNER [111] *observational congruence* was defined as the CCS-closure of their variant of observational equivalence (Definition 1.8) and this congruence coincides with the one of MILNER [134]. BERGSTRA & KLOP [26] formulated an additional condition, yielding an immediate definition of observation congruence by means of bisimulation relations.

DEFINITION 2.2 (root condition) A relation R between nodes of process graphs is called *rooted* if root nodes are related with root nodes only.

Observe that every bisimulation as in Definition 1.2 is rooted, but this is not necessarily the case for the relations defined in Definitions 1.3-1.6. For any two process graphs g and h and $* \in \{\tau, b, \eta, d\}$ we write $R: g \equiv_{r*} h$ if R is a rooted $*$ -bisimulation between g and h , and $g \equiv_{r*} h$ if such a relation exists.

THEOREM 2.3 For $* \in \{\tau, b, \eta, d\}$, $\equiv_{r*}$ is a congruence on G_{BPA} with respect to $+$ and $\cdot$.

PROOF We prove Theorem 2.3 for $\equiv_{rb}$. The other proofs proceed in the same way.

$\equiv_{rb}$ is reflexive since the identity relation is a rooted branching bisimulation between any graph and itself, and it is symmetric by definition. Furthermore, assume that $R: g \equiv_{rb} g'$ and $S: g' \equiv_{rb} g''$ and define: $T(r, r'') : \Leftrightarrow$ for some r' in g' : $R(r, r')$ and $S(r', r'')$. Now one can easily prove that $T: g \equiv_{rb} g''$, and so $\equiv_{rb}$ is transitive. Thus we proved that $\equiv_{rb}$ is an equivalence. So it is left to prove that $\equiv_{rb}$ respects the operators. Suppose that $R: g \equiv_{rb} g'$ and $S: h \equiv_{rb} h'$.

$\pm$: We prove that $(R \cup S): (g + h) \equiv_{rb} (g' + h')$.

(i) Obviously the roots of $(g + h)$ and $(g' + h')$ are related.

(ii) Assume that in $(g + h)$ we have an edge $r \xrightarrow{a} r'$ and suppose we have $(R \cup S)(r, s)$ then from the construction of $(g + h)$ it follows that this edge either originates from g or from h ; let us say it is g . It follows from $(R \cup S)(r, s)$ that we have either $R(r, s)$ or $S(r, s)$, so we have two distinct cases:

Firstly, suppose that $R(r, s)$. Then either $a = \tau$ and $R(r', s)$ - hence $(R \cup S)(r', s)$ and $(R \cup S)$ satisfies Definition 1.4 - or there exists a path $s \Rightarrow s_1 \xrightarrow{a} s_2 \Rightarrow s'$ in g' such that $R(r, s_1)$, $R(r', s_2)$ and $R(r', s')$. Obviously, we can find the same path in $(g' + h')$ and we have that $(R \cup S)(r, s_1)$, $(R \cup S)(r', s_2)$ and $(R \cup S)(r', s')$ as required.

Secondly, suppose that we do *not* have $R(r,s)$. Then we have $S(r,s)$, and since we assumed that the edge $r \rightarrow^a r'$ came from g , we find that r has to be the (joint) root node of g and h . However, in S root nodes are related with root nodes only (the root condition), and so s must be the joint root node of g' and h' . Hence $R(r,s)$, which is a contradiction.

(iii) Obviously, the root nodes of $(g + h)$ and $(g' + h')$ are uniquely related by $(R \cup S)$.

$\therefore$ we prove $R \cup S: (g \cdot h) \cong_{rb} (g' \cdot h')$.

(i) Clearly, the roots of both graphs are related by R , hence by $R \cup S$.

(ii) Assume that in $(g \cdot h)$ we have an edge $r \rightarrow^a r'$ and suppose we have $(R \cup S)(r,s)$ then from the construction of $(g \cdot h)$ it follows that this edge either originates from g or from h .

(1) Firstly, let us say it is from g . From $(R \cup S)(r,s)$ we find that either $R(r,s)$ or $S(r,s)$. Since r cannot be an endnode in g we have $R(r,s)$. It follows from the fact that R is a rooted branching bisimulation that either $a=\tau$ and $R(r',s)$ - hence $(R \cup S)(r',s)$ as required - or there is a path $s \Rightarrow s_1 \rightarrow^a s_2 \Rightarrow s'$ in g' such that $R(r,s_1)$, $R(r',s_2)$ and $R(r',s')$ and thus we find the same path in $(g' \cdot h')$ such that $(R \cup S)(r,s_1)$, $(R \cup S)(r',s_2)$ and $(R \cup S)(r',s')$, as is required.

(2) Secondly, assume $r \rightarrow^a r'$ is from h .

- In case $R(r,s)$, we find that r is an endnode in g (since those are the only nodes of g that are identified with nodes from h). Suppose s is an endnode in g' , then it is identified with the root node of h' , and since S is a rooted branching bisimulation we find:

either $a=\tau$ and $S(r',s)$, hence $(R \cup S)(r',s)$;

or there exists a path $s \Rightarrow s_1 \rightarrow s_2 \Rightarrow s'$ such that $S(r,s_1)$, $S(r',s_2)$ and $S(r',s')$ and hence $(R \cup S)(r,s_1)$, $(R \cup S)(r',s_2)$ and $(R \cup S)(r',s')$ as required.

So let us assume that s is *not* an endnode in g' , then it has at least one outgoing edge $s \rightarrow^b s_1$. Since R is a rooted branching bisimulation and $R(r,s)$, we find that $b=\tau$ and $R(r,s_1)$. The same argument holds for s_1 and thus we find a path $s=s_0 \rightarrow^\tau s_1 \rightarrow^\tau s_2 \rightarrow^\tau \dots$ such that $R(r,s_i)$. Since all graphs in G_{BPA} are divergence free we have that all nodes s_i are distinct and furthermore the sequence $s=s_0 \rightarrow^\tau s_1 \rightarrow^\tau s_2 \rightarrow^\tau \dots$ has bounded length. Hence there exists a path $s \Rightarrow s'$ to an endnode s' in g' , such that $R(r,s')$ (and hence $(R \cup S)(r,s')$). Note that s' is identified with the root node of h' . Combining this result with the former part, we find that the conditions of Definition 1.4 are satisfied as required.

- In case *not* $R(r,s)$, then $S(r,s)$ and both r and s are from h and h' respectively. Now the requirement follows immediately from the fact that S is a branching bisimulation.

(iii) Clearly, the root nodes are uniquely related by $(R \cup S)$. □

THEOREM 2.4 *Provided that there exists at least one action $a \in \text{Act}$ with $a \neq \tau$, $\cong_{r*}$ is the coarsest congruence on G_{BPA} with respect to $+$ that is contained in $\cong_*$, for $* \in \{\tau, b, \eta, d\}$. Hence $\cong_{r\tau}$ coincides with observation congruence.*

PROOF The idea for this proof is due to J.W. Klop (personal communication). Let g and $h \in G_{BPA}$ and suppose $g+k \approx_* h+k$ for any graph $k \in G_{BPA}$. Suppose there is an action $a \in \text{Act}$ ($a \neq \tau$) that does not occur in g and h . Then $g+a \approx_* h+a$. Let R be a $*$ -bisimulation between $g+a$ and $h+a$, then R must be rooted. Therefore the restriction of R to the nodes of g and h is a rooted $*$ -bisimulation between g and h .

If no 'fresh atom' $a \in \text{Act}$ can be found a variant of this method still works. First note that for each infinite cardinal κ there are at least κ $*$ -bisimulation equivalence classes of graphs with less than κ nodes. (Choose an action $a \in \text{Act}$ ($a \neq \tau$) and define for each ordinal $\lambda > 0$ the graphs g_λ as follows: $g_1 = a$, $g_{\lambda+1} = g_\lambda + a g_\lambda$ and for λ a limit ordinal g_λ is constructed from all graphs g_μ for $\mu < \lambda$ by identifying their roots. Then with transfinite induction it follows that no two different g_λ 's are $*$ -bisimilar. Furthermore, for infinite λ , the cardinality of the nodes of g_λ is the cardinality of λ .) Thus for any two graphs g and h there must be a graph $k \in G_{BPA}$ with the same cardinality such that k is not bisimilar with any subgraph corresponding with a node in g or h . Now take a $*$ -bisimulation between $g+\tau k$ and $h+\tau k$. $\square$

The equivalence relations $\approx_{\tau*}$ are called *rooted $*$ -bisimulation equivalence* or *$*$ -bisimulation congruence*. As a consequence of Theorem 2.3, we find that all structures $(G_{BPA}/\approx_{\tau*}, +, \cdot, \text{Act})$ are well-defined algebras, every one of which may satisfy a different equational theory. In a slightly different setting, MILNER [134] found that the algebra $(G_{BPA}/\approx_{\tau*}, +, \cdot, \text{Act})$ can be completely axiomatized by BPA together with the following three equations:

$x\tau = x$	T1
$\tau x = \tau x + x$	T2
$a(\tau x + y) = a(\tau x + y) + ax$	T3

Table 2. τ -laws ($a \in \text{Act}$).

THEOREM 2.5 *BPA + T1-T3 is a complete axiomatization of $(G_{BPA}/\approx_{\tau*}, +, \cdot, \text{Act})$.*

In the setting of BPA and process graphs, this theorem was first established in BERGSTRA & KLOP [26]. Its proof will be given in Section 4, together with the proofs of the Theorems 2.6-2.8.

From Figure 1 one can observe that the constructions (b) and (c) are highly fundamental for the behaviour of τ in the graph model. For instance, by simplifying Figure 1 (b) one finds the second τ -law T2, whereas T3 can be easily found from Figure 1 (c). This shows us that the extra τ -laws T2 and T3 originate from the fact that observation equivalence does not preserve branching structures.

Since branching bisimulation equivalence distinguishes between all three graphs in Figure 1, we expect that the laws T2 and T3 will no longer hold in $(G_{BPA}/\equiv_{rb}, +, \cdot, \text{Act})$. As it turns out, Axiom T3 is completely dropped and T2 is considerably weakened to Axiom H2 from the following table:

$x\tau = x$	H1 (T1)
$x(\tau(y + z) + y) = x(y + z)$	H2

Table 3. τ -laws for branching bisimulation.

H1 is the same axiom as T1 whereas H2 is derivable from T1 and T2 as one can check easily. Both axioms refer to the axiomatization of η , a constant for abstraction from BAETEN & VAN GLABBEEK [15] similar to τ . In fact they are a variation on the first two η -laws in the sense that in [15] the second law H2 was only introduced for atomic actions x , instead of taking x as a general variable ranging over all processes. On the domain of closed terms the two variants are equally powerful. In Section 4 we will establish that

THEOREM 2.6 $BPA + H1\text{-}H2$ is a complete axiomatization of $(G_{BPA}/\equiv_{rb}, +, \cdot, \text{Act})$.

Obviously, $\equiv_{\eta}$ is a coarser notion than $\equiv_{rb}$ and it respects the Axioms H1-H2. As it turns out we have the additional Axiom H3 which was introduced earlier as T3.

$x\tau = x$	H1 (T1)
$x(\tau(y + z) + y) = x(y + z)$	H2
$a(\tau x + y) = a(\tau x + y) + ax$	H3 (T3)

Table 4. η -laws ($a \in \text{Act}$).


BAETEN & VAN GLABBEEK [15] established a completeness theorem for rooted η -bisimulation:

THEOREM 2.7 $BPA + H1\text{-}H3$ is a complete axiomatization of $(G_{BPA}/\equiv_{\eta}, +, \cdot, \text{Act})$.

So, on closed terms, the difference between H2 and T2 is precisely all the difference there is between the usual τ -laws and η . Finally a completeness theorem for delay bisimulation was (in the setting of CCS) established by WALKER [181].

THEOREM 2.8 $BPA + T1\text{-}T2$ is a complete axiomatization of $(G_{BPA}/\equiv_{\tau\Delta}, +, \cdot, \text{Act})$.

Resuming we have the following diagram (see Figure 7):



T1	$x\tau = x$	H1
	$x(\tau(y + z) + y) = x(y + z)$	H2
T2	$\tau x = \tau x + x$	
T3	$a(\tau x + y) = a(\tau x + y) + ax$	H3

FIGURE 7. Four notions of bisimulation with $\tau (a \in \text{Act})$.

2.2 CCS

In the setting of CCS we extend the graph domain G_{BPA} to G_{CCS} consisting of the root unwound process graphs, thus no longer excluding the trivial graph (the one-node graph without edges) nor any of the graphs with divergences (i.e. infinite τ -paths). We obtain: $G_{\text{BPA}} \subseteq G_{\text{CCS}} \subseteq G$.

We introduce a constant 0 for inaction, a binary infix written operator $+$ for alternative composition, and unary operators $a.$ for prefixing ($a \in \text{Act}$).

DEFINITION 2.3 The constant 0 and the operators $+$ and $a.$ are defined on G_{CCS} as follows:

- (i) The constant 0 is interpreted as the trivial graph
- (ii) $(g + h)$ can be constructed by identifying the root nodes of g and h
- (iii) $(a.g)$ is constructed from g by adding a new node which will be the root of $a.g$, and a new a -labelled edge from the root of $a.g$ to the root of g .

We will often leave out brackets, assuming that $+$ will be the weakest operator symbol. For agents p we will often write ap instead of $a.p$ in order to avoid non-essential distinctions between CCS and ACP. Similarly, we write Act for the set of prefix operators $\{a.: a \in \text{Act}\}$. MILNER [134] proved that the operators from Act and $+$ all are well-defined on $G_{\text{CCS}}/\equiv$:

PROPOSITION 2.9 *Bisimulation equivalence is a congruence with respect to the operators from Act and +.*

So again, the structure $(G_{CCS}/\equiv, 0, +, \text{Act})$ is a well-defined algebra, and as in the case of $(G_{BPA}/\equiv, +, \cdot, \text{Act})$ we can find a complete axiomatization of its equalities with respect to closed terms:

$x + y = y + x$	A1
$(x + y) + z = x + (y + z)$	A2
$x + x = x$	A3
$x + 0 = x$	A6

Table 5. Basic CCS.

Let us call the theory from Table 5 *Basic CCS*, and write $BCCS := A1-A3, A6$. Then the following theorem is due to HENNESSY & MILNER [111] and MILNER [134].

THEOREM 2.10 *BCCS is a complete axiomatization of $(G_{CCS}/\equiv, 0, +, \text{Act})$.*

As before, we have four other equivalences $\equiv_{r*}$ for $* \in \{\tau, b, \eta, d\}$ on G_{CCS} which can be considered. First we establish that they are congruences.

THEOREM 2.11 *For $* \in \{\tau, b, \eta, d\}$, $\equiv_{r*}$ is a congruence on G_{CCS} w.r.t. + and Act.*

PROOF We prove it for $\equiv_{rb}$. The other proofs proceed in the same way.

The proof that $\equiv_{rb}$ is an equivalence and respects + can be copied from the proof of Theorem 2.3. So it is left to prove that it respects the operators in Act. So suppose that $R: g \equiv_{rb} g'$ and p, p' are the root nodes of $a.g$ and $a.g'$. Put $R^* := R \cup \{p, p'\}$. Then we prove $R^*: (a.g) \equiv_{rb} (a.g')$.

(i) Clearly, the roots of both graphs are related by R^* .

(ii) Assume that in $(a.g)$ we have an edge $r \rightarrow^b r'$ and suppose $R^*(r, s)$ then from the construction of $(a.g)$ it follows that either $r=p$ or this edge originates from g .

If $r=p$ then by the definition of R^* we have $s=p'$. Furthermore, $b=a$ and r' is the root node of g and by the construction of prefixing we find that in g' there exists an edge $s \rightarrow^a s'$ to the root node s' of g' . Since R is a branching bisimulation we find $R(r', s')$ and hence $R^*(r', s')$.

If $r \rightarrow^b r'$ originates from g then it follows from the definition of R^* that $R(r, s)$, from which the requirement follows immediately.

(iii) Clearly, the root nodes are uniquely related by R^* . □

It follows from Theorem 2.4 that $\equiv_{r*}$ is moreover the coarsest BCCS-congruence contained in $\equiv_*$.

Now consider the axioms from the following table:

H1'	$a\tau x = ax$	T1'
H2'	$a(\tau(y + z) + y) = a(y + z)$	
	$\tau x = \tau x + x$	T2'
H3'	$a(\tau x + y) = a(\tau x + y) + ax$	T3'

Table 6. τ -laws in CCS ($a \in \text{Act}$).

The only difference between these axioms and the ones introduced in the previous section is the replacement of sequential composition by prefixing in the Axioms T1 (H1) and H2. The prime accents (') refer to this replacement. Note that H1' is derivable from H2. We find the following completeness results:

THEOREM 2.12

- (i) BCCS is a complete axiomatization of $(G_{\text{CCS}}/\equiv, 0, +, \text{Act})$
- (ii) BCCS + T1'-T3' is a complete axiomatization of $(G_{\text{CCS}}/\equiv_{\tau}, 0, +, \text{Act})$.
- (iii) BCCS + H2' is a complete axiomatization of $(G_{\text{CCS}}/\equiv_{\tau b}, 0, +, \text{Act})$.
- (iv) BCCS + H2'-H3' is a complete axiomatization of $(G_{\text{CCS}}/\equiv_{\tau m}, 0, +, \text{Act})$.
- (v) BCCS + T1'-T2' is a complete axiomatization of $(G_{\text{CCS}}/\equiv_{\tau d}, 0, +, \text{Act})$.

For the proof of Theorem 2.12, we refer to Section 4.

2.3. TERMINATION

In the previous two subsections, we presented two models: the model $(G_{\text{BPA}}/\equiv_{\tau^*}, +, \cdot, \text{Act})$ for BPA with sequential composition, and $(G_{\text{CCS}}/\equiv_{\tau^*}, 0, +, \text{Act})$ for BCCS with prefixing. As we argued before, including the trivial graph in $G_{\text{BPA}}/\equiv_{\tau^*}$ would destroy the soundness of BPA in the corresponding model, namely of the Axiom A4. Furthermore, from $G_{\text{BPA}}/\equiv_{\tau^*}$ we have to exclude graphs containing infinite τ -paths since otherwise sequential composition no longer respects the equivalences— i.e. the equivalences $\equiv_{\tau^*}$ are no longer congruences with respect to $\cdot$. For consider the following example:




FIGURE 8. Equivalent graphs with and without divergence.

In Figure 8, we find two equivalent graphs, one with and one without divergence, which we informally denote by $a \cdot \tau^\omega$ and a . So: $a \cdot \tau^\omega \equiv_{\tau^*} a$, for $* \in \{\tau, b, \eta, d\}$. However, since $a \cdot \tau^\omega$ does not contain any endnodes we find that $(a \cdot \tau^\omega) \cdot b = a \cdot \tau^\omega$ and $a \cdot \tau^\omega \not\equiv_{\tau^*} ab$. So in the presence of divergence $\equiv_{\tau^*}$ no longer is a congruence with respect to $\cdot$.

The question arises whether the virtues of $(G_{BPA}/\equiv_{\tau^*}, +, \cdot, \text{Act})$ and $(G_{CCS}/\equiv_{\tau^*}, 0, +, \text{Act})$ can be combined, i.e. whether it is possible to define inaction and general sequential composition in one model (without destroying the intuitively plausible Axiom A4) as well as to define general sequential composition on graphs with possible divergence paths, while respecting the equivalences. We will give a positive answer to this question by once again extending G_{CCS} to a larger domain G_{ACP} (so: $G_{BPA} \subseteq G_{CCS} \subseteq G_{ACP}$).

Let us extend the set Act with an additional label, written as $\checkmark$. Then, in G_{ACP} we will distinguish between successful and unsuccessful termination of a process by adding a *termination edge* to the endnodes which are considered to terminate successfully. Such termination edges consist of an outgoing edge labelled with $\checkmark$ to a new endnode. Let G_{ACP} consist of all graphs that can be obtained from non-trivial, root unwound graphs from G_{BPA} by adding termination edges to some of their endnodes. Next we add the trivial graph to G_{ACP} but assume that G_{ACP} is without the graph consisting of a single termination edge, i.e. the graph representing instant termination.

Observe that in graphs from G_{ACP} every node has at most one outgoing termination edge and if it has one, then it does not have any other outgoing edges. Furthermore, if a node has an incoming termination edge then it is an endnode and it does not have any other incoming edges. We immediately find that $G_{CCS} \subseteq G_{ACP}$ and $G_{ACP} \subseteq G^\checkmark$, where $G^\checkmark$ is the set of process graphs with $\checkmark$ as a possible edge-label. The difference between G_{CCS} and G_{ACP} is that the latter distinguishes between two kinds of termination.




FIGURE 9. Process graphs with termination edges.

With respect to the algebraic operators, we simply combine the operators from BPA and CCS, but we adapt the definitions of action constants and sequential composition

to the presence of $\sqrt{}$ -labels. This is done in the following definition. The new operator for sequential composition will again be denoted by $\cdot$, and similarly for action constants. It will appear from the context (whether it is about G_{BPA} or G_{ACP}) which one of the Definitions 2.1 and 2.4 presents their current interpretation. In case of doubt we underline the BPA operators.

DEFINITION 2.4 On G_{ACP} the constants 0 and a (for $a \in \text{Act}$) and the operators $+$ and $\cdot$ are defined as follows:

- (i) 0 is the trivial graph
- (ii) Constants $a \in \text{Act}$ are interpreted by the left hand side of Figure 9
- (iii) $(g + h)$ can be constructed by identifying the root nodes of g and h
- (iv) $(g \cdot h)$ is constructed by identifying every node in g with an outgoing termination edge with the root node of h while deleting its termination edge. The root node of $(g \cdot h)$ is that of g . If g is without termination edges, then $(g \cdot h)$ is just g .

The prefixing operator of CCS can now be defined by: $a.g = a \cdot g$. In the subdomain G_{CCS} of G_{ACP} all processes end in deadlock (unsuccessful termination), so $g \cdot h = g$. This explains the absence of sequential composition on G_{CCS} . Let G'_{BPA} be the subdomain of G_{ACP} consisting of all divergence free graphs from G_{ACP} only ending with successful termination. Then $(G'_{BPA}, +, \cdot, \text{Act})$ and $(G_{BPA}, +, \cdot, \underline{\text{Act}})$ are isomorphic algebras and the latter can be interpreted as a notational abbreviation of the former, where all $\sqrt{}$ -labels have been left out.

On the new graph domain G_{ACP} we can define the bisimulation relations from Definition 1.2-1.6 and 2.2, taking into account that $\sqrt{} \in \text{Act}$. That is, termination edges are not treated differently from other edges. The relations on G_{BPA} , inherited through the isomorphism from G'_{BPA} , coincide with the relations considered in Subsection 2.1. However, this is no longer true if divergent graphs would be added to G_{BPA} ; in that case all relations need an additional clause:

- If $R(r, s)$ and r is an endnode then there is a path $s \Rightarrow s'$ to an endnode s' .

In order to prevent this complication in Section 2.1, there we treated divergence free graphs only.

The fact that Definition 2.4 provides us with a proper algebraic structure on G_{ACP} follows from the following theorem:

THEOREM 2.13 All equivalences $\equiv$, $\equiv_{\tau}$, $\equiv_{rb}$, $\equiv_{\tau\eta}$ and $\equiv_{rd}$ are congruences with respect to the operators $+$ and $\cdot$ on G_{ACP} .

PROOF Again we prove the theorem for $\equiv_{rb}$. The fact that on G_{ACP} they are congruences with respect to $+$ can be found from the proof of Theorem 2.3. Considering the proof for $\cdot$, suppose that $R: g \equiv_{rb} g'$ and $S: h \equiv_{rb} h'$. Let R' be the restriction of R to the nodes in g that also appear in $g \cdot h$ (i.e. the nodes without incoming $\sqrt{}$ -edges). We prove that $R' \cup S: (g \cdot h) \equiv_{rb} (g' \cdot h')$.

- (i) Clearly the roots of $(g \cdot h)$ and $(g' \cdot h')$ are related by $R' \cup S$.

(ii) Assume that in $(g \cdot h)$ we have an edge $r \rightarrow^a r'$ and suppose $(R' \cup S)(r, s)$, then from the construction of $(g \cdot h)$ it follows that this edge either originates from g or from h . If it is from g , then the proof proceeds as in the proof of Theorem 2.3. So assume $r \rightarrow^a r'$ is from h .

—In case $R'(r, s)$, we find that in g the node r has an outgoing termination edge $r \rightarrow^\vee r''$ to an endnode r'' (since those are the only nodes of g that are identified with nodes from h). Since R is a branching bisimulation, we find that in g' there exists a path $s \Rightarrow s' \rightarrow^\vee s''$ such that $R(r, s')$ and $R(r'', s'')$. By applying the definition of $\models_{rb}$ we even find that all nodes in $s \Rightarrow s'$ are related with r . Furthermore, by construction of $(g' \cdot h')$ the node s' is identified with the root node of h' , and since S is a rooted branching bisimulation between h and h' , we find:

either $a = \tau$ and $S(r', s')$, hence $(R' \cup S)(r', s')$;

or there exists a path $s' \Rightarrow s_1 \rightarrow s_2 \Rightarrow s_3$ such that $S(r, s_1)$, $S(r', s_2)$ and $S(r', s_3)$ and hence $(R' \cup S)(r, s_1)$, $(R' \cup S)(r', s_2)$ and $(R' \cup S)(r', s_3)$ as required.

—In case *not* $R'(r, s)$, then $S(r, s)$ and both r and s are from h and h' respectively. Now the requirement follows immediately from the fact that S is a branching bisimulation.

(iii) In G_{ACP} the root node cannot have an outgoing termination edge, and hence the root nodes of $(g \cdot h)$ and $(g' \cdot h')$ are only related by R' (they are not identified with nodes from h or h'). Hence $(R' \cup S)$ is rooted since R is. $\square$

As a consequence, we find a well-defined algebra $(G_{ACP}/\equiv, 0, +, \cdot, \text{Act})$, and four others with domain $G_{ACP}/\equiv_{r*}$ ($* \in \{\tau, b, \eta, d\}$). To start with, we find that the following basic theory is valid in all five algebras (see Table 7):

$x + y = y + x$	A1
$(x + y) + z = x + (y + z)$	A2
$x + x = x$	A3
$(x + y)z = xz + yz$	A4
$(xy)z = x(yz)$	A5
$x + 0 = x$	A6
$0 \cdot x = 0$	A7

Table 7. BPA_0 .

The theory BPA_0 is the kernel of the axiom system ACP, introduced in BERGSTRA & KLOP [25], where 0 was called δ . As before, we have the following completeness theorem for the five respective algebras:

THEOREM 2.14

- (i) BPA_0 is a complete axiomatization of $(G_{ACP/\equiv, 0, +, \cdot, \text{Act}})$
- (ii) $BPA_0 + T1-T3$ is a complete axiomatization of $(G_{ACP/\equiv_{\tau}, 0, +, \cdot, \text{Act}})$
- (iii) $BPA_0 + H1-H2$ is a complete axiomatization of $(G_{ACP/\equiv_{rb}, 0, +, \cdot, \text{Act}})$
- (iv) $BPA_0 + H1-H3$ is a complete axiomatization of $(G_{ACP/\equiv_{\eta}, 0, +, \cdot, \text{Act}})$
- (v) $BPA_0 + T1-T2$ is a complete axiomatization of $(G_{ACP/\equiv_{rd}, 0, +, \cdot, \text{Act}})$.

Again for the proof of this completeness theorem we refer to Section 4.

2.4. TURNING EQUIVALENCES INTO CONGRUENCES

τ -bisimulation equivalence is not a congruence for $+$, and therefore τ -bisimulation congruence is defined as the closure of τ -bisimulation equivalence under contexts, or by means of the root condition. We established (in Theorem 2.4) that (if $\text{Act} - \{\tau\}$ is nonempty) both methods yield the same result. In this respect η -, delay and branching bisimulation behave exactly the same. However, each τ -bisimulation equivalence class consists of at most two τ -bisimulation congruence classes (this follows from Exercise 7.6 of HENNESSY in MILNER [134]), as is the case for delay bisimulation, whereas η - and branching bisimulation equivalence classes may contain many congruence classes. Nevertheless, for all four bisimulations there exists a close relationship between rooted and non-rooted bisimulation, since the root condition (Definition 2.2) only works on the root nodes:

THEOREM 2.15 *For all root unwound graphs g and h and $* \in \{\tau, b, \eta, d\}$ we have:*

$$g \equiv_* h \text{ if and only if } \tau \cdot g \equiv_{\tau*} \tau \cdot h.$$

PROOF If R is a $*$ -bisimulation between g and h and r, s are the roots of $\tau \cdot g$ and $\tau \cdot h$ then $R \cup \{r, s\}$ is a rooted $*$ -bisimulation between $\tau \cdot g$ and $\tau \cdot h$. On the other hand, if R is a rooted $*$ -bisimulation between $\tau \cdot g$ and $\tau \cdot h$, then the roots of g and h are related by R , so R restricted to the nodes of g and h is a $*$ -bisimulation between g and h . $\square$

This theorem provides us with a tool to decide upon $*$ -bisimulation equivalence, using the axiom systems of $*$ -bisimulation congruence.

In defining rooted bisimulations we restricted ourselves to root-unwound process graphs. The reason for doing so is that the definitions given do not work for graphs with incoming edges in the root. As remarked in Section 2.1 already, this does not cause any loss of generality. It is straightforward to define an operator ρ for root-unwinding, such that for any graph g , $\rho(g)$ is root-unwound and strongly bisimilar with g (and moreover if g is root-unwound already, $\rho(g)$ and g are isomorphic). Now the definitions of rooted bisimulation can be extended to all process graphs by letting $g \equiv_{\tau*} h$ iff $\rho(g) \equiv_{\tau*} \rho(h)$. In case we do not restrict to root unwound process graphs,

the definitions of the operators $+$ and $\cdot$ on process graphs have to be upgraded as well. An easy way to do so is by defining $g + h$ as $\rho(g) + \rho(h)$ and $g \cdot h$ as $g \cdot \rho(h)$.

A rooted branching bisimulation between two graphs g and h can alternatively be defined as a branching bisimulation R such that, for a visible or invisible:

- iii. If $\text{root}(g) \rightarrow^a r$, then there is an edge $\text{root}(h) \rightarrow^a s$ with $R(r, s)$.
- iv. If $\text{root}(h) \rightarrow^a s$, then there is an edge $\text{root}(g) \rightarrow^a r$ with $R(r, s)$.

This means that a rooted branching bisimulation behaves like a strong bisimulation in the first step, and like a branching bisimulation further on. It is easy to see that for root-unwound graphs this definition agrees with Definition 2.2. However, it has the advantage that it can directly be used for graphs that are not root-unwound. Also delay, η - and τ -bisimulation congruence can be characterized in a similar fashion (cf. MILNER [138] and BAETEN & VAN GLABBEK [15]) (but giving up the analogy with a strong bisimulation in the first step). However, the characterizations of this type of the various bisimulation congruences have not exactly the same appearance, whereas the root condition gives a uniform way for turning weak, delay, η - and branching bisimulation into a congruence.

3. BRANCHES AND TRACES

As we saw in Figure 1, while preserving observation equivalence we are able to introduce new 'paths' in a graph. To be more precise: in these new paths alternative options may branch off in places different from any in the old paths. So far, we claimed to have solved this problem by defining a new kind of bisimulation, but as of yet we still have to prove that our solution solves the problem in a fundamental way. In this section we will establish an alternative characterization of branching bisimulation. In fact, we will show how branching bisimulation preserves the branching structure of graphs. Let us first consider ordinary bisimulation.

DEFINITION 3.1 A *concrete trace* of a process graph is a finite sequence $(a_1, a_2, a_3, \dots, a_k)$ of actions from Act, such that there exists a path $r_0 \rightarrow^{a_1} r_1 \rightarrow^{a_2} r_2 \rightarrow \dots \rightarrow^{a_k} r_k$ from the root node r_0 .

Two graphs g and h are said to be *concrete trace equivalent*, notation $g \equiv_t h$, if their *concrete trace sets* (i.e. the sets of their concrete traces) are equal. It is easily checked that $\equiv_t$ is a congruence on G_{BPA} and $g \rightleftharpoons h \Rightarrow g \equiv_t h$. Consequently we find that $G_{BPA}/\equiv_t$ is a model for BPA. Compared to bisimulation, concrete trace equivalence makes many more identifications. For example, we find that $G_{BPA}/\equiv_t$ satisfies the equation $x(y + z) = xy + xz$ which cannot be proved from BPA.

The main reason for this is that a concrete trace does not provide us with information about the branching potentials in the intermediate nodes. Therefore we cannot distinguish between processes $a(b + c)$ and $(ab + ac)$. In the following we will use *colours* at the nodes to indicate these potentials.

DEFINITION 3.2 A *coloured graph* is a process graph with colours $C \in \mathcal{C}$ as labels at the nodes.

Obviously, in a coloured graph we have traces which have colours in the nodes:

DEFINITION 3.3 A *concrete coloured trace* of a coloured graph g is a sequence of the form $(C_0, a_1, C_1, a_2, C_2, \dots, a_k, C_k)$ for which there exists a path $r_0 \xrightarrow{a_1} r_1 \xrightarrow{a_2} r_2 \rightarrow \dots \rightarrow^{a_k} r_k$ in g , starting from the root node r_0 , such that r_i has colour C_i .

The concrete coloured traces of a node r in a graph g are the concrete coloured traces of the subgraph $(g)_r$ of g that has r as its root node. This graph is obtained from g by deleting all nodes and edges which are inaccessible from r .

The question remains how to detect the colour differences of the nodes, or—to put it differently—how to define the concept of 'branching potential in a node' properly. There are several ways to do this. A natural definition is the following:

DEFINITION 3.4 A *concrete consistent colouring* of a set of graphs is a colouring of their nodes with the property that two nodes have the same colour only if they have the same concrete coloured trace set.

Obviously, the *trivial* colouring—in which every node has a different colour—is consistent on any set of graphs. Note that—even apart from the choice of the colours—a set of graphs can have more than one consistent colouring. For instance, consider a set containing only an infinite graph representing a^ω or $a \cdot a \cdot a \dots$ then obviously the *homogeneous* colouring—in which every node has the same colour—is a consistent one, as well as the *alternating* or the trivial colouring.

Let us say two graphs g and h are *concrete coloured trace equivalent*—notation: $g \equiv_{cc} h$ —if for some concrete consistent colouring on $\{g, h\}$ they have the same concrete coloured trace set, or equivalently, if for some concrete consistent colouring on $\{g, h\}$ the root nodes have the same colour. Then we have the following important characterization:

THEOREM 3.1 $g \equiv h$ if and only if $g \equiv_{cc} h$.

PROOF $\Rightarrow$: Suppose R is the largest bisimulation relation between g and h . Let $\underline{R}$ be the transitive closure of R , then $\underline{R}$ is an equivalence relation on the set of nodes from g and h . Let $\mathcal{C}$ be the set of equivalence classes induced by $\underline{R}$ and label every node with its own equivalence class. Then this colouring is consistent on g and h .

To see this let r_0 be a node in g say, and $(C_0, a_1, C_1, a_2, C_2, \dots, a_k, C_k)$ be a concrete coloured trace which corresponds to a path $r_0 \xrightarrow{a_1} r_1 \xrightarrow{a_2} r_2 \rightarrow \dots \rightarrow^{a_k} r_k$ starting from r_0 . Now suppose for some node s_0 in h we have $R(r_0, s_0)$, then we find from Definition 1.2 that $s_0 \xrightarrow{a_1} s_1$ for some s_1 such that $R(r_1, s_1)$. Thus r_1 and s_1 have the same colour C_1 . By induction we find that s_0 has the same concrete

coloured trace $(C_0, a_1, C_1, a_2, C_2, \dots, a_k, C_k)$. So R preserves concrete coloured trace sets, hence so does $\underline{R}$.

Since the roots of g and h are related we find $g \equiv_{cc} h$.

$\Leftarrow$: Suppose that g and h have the same concrete coloured trace sets. Then consider the relation R which relates two nodes of g and h iff they are labelled with the same colour. It is easy to prove that R is a bisimulation between g and h . $\square$

So far we did not have any notion of abstraction in the definition of coloured traces, so if a coloured graph has τ -labels then these are treated as if they were ordinary actions. In the following definition we find how to abstract from these τ -steps. The idea is simple: τ -steps can only be left out if they are *inert*, meaning that they are between two nodes that have the same colour. Thus it is not only that inert steps are not observable, but even more, they do not cause any change in the overall state of the machine.

DEFINITION 3.5 A *coloured trace* of a coloured graph is a sequence of the form $(C_0, a_1, C_1, a_2, C_2, \dots, a_k, C_k)$ which is obtained from a concrete coloured trace of this graph by replacing all subsequences $(C, \tau, C, \tau, \dots, \tau, C)$ by C .

DEFINITION 3.6 A *consistent colouring* of a set of graphs is a colouring of their nodes with the property that two nodes have the same colour only if they have the same coloured trace set. Furthermore such a colouring is *rooted* if no root-node has the same colour as a non-root node.

For two root unwound graphs g and h let us write $g \equiv_c h$ if for some consistent colouring on $\{g, h\}$ they have the same coloured trace set, and $g \equiv_{rc} h$ if moreover this colouring is rooted. Then we find the following characterization for (rooted) branching bisimulation:

THEOREM 3.2

- i. $g \rightleftharpoons_b h$ if and only if $g \equiv_c h$
- ii. $g \rightleftharpoons_{rb} h$ if and only if $g \equiv_{rc} h$.

PROOF $\Rightarrow$: Suppose R is the largest (rooted) branching bisimulation between g and h . Let $\underline{R}$ be its transitive closure and C the set of equivalence classes induced by $\underline{R}$. Then the colouring in which every node is labelled with its own equivalence class is consistent (and rooted) on g and h .

To see this, let us write $C(r)$ for the colour of the node r and assume that, for certain nodes r_0 and s_0 , $R(r_0, s_0)$ and r_0 has an coloured trace $(C_0, a_1, C_1, a_2, C_2, \dots, a_k, C_k)$. Then there exists a path of the form $r_0 \xrightarrow{\tau} u_1 \xrightarrow{\tau} \dots \xrightarrow{\tau} u_m \xrightarrow{a_1} r_1$ ($m \geq 0$) such that $C(r_1) = C_1$ and for all i : $C(u_i) = C(r_0) = C_0$. For every edge $u_i \xrightarrow{\tau} u_{i+1}$ ($0 \leq i < m$, $u_0 = r_0$) there exists a path $v_i \Rightarrow v_{i+1}$ ($v_0 = s_0$) such that $R(u_i, v_i)$, and all intermediate nodes are related to either u_i or u_{i+1} (by Lemma 1.1), hence all v_i have the same colour C_0 . So we find a path $s_0 \Rightarrow v_m$ with only one colour in the nodes such that $R(u_m, v_m)$.

Next, since $u_m \rightarrow^{a_1} r_1$ and $R(u_m, v_m)$ we find that either $a_1 = \tau$ and $R(r_1, v_m)$ —in which case $C_1 = C_0$ in contradiction with $(C_0, a_1, C_1, a_2, C_2, \dots, a_k, C_k)$ being a coloured trace—or there is a path $v_m \Rightarrow t_1 \rightarrow^{a_1} s_1$ such that $R(u_m, t_1)$ and $R(r_1, s_1)$. Again by Lemma 1.1 we find that t_1 and all the intermediate nodes in $\Rightarrow$ have the same colour as v_m and so we find a coloured trace (C_0, a_1, C_1) of s_0 . By repeating this argument k times, we find that s_0 has a coloured trace $(C_0, a_1, C_1, a_2, C_2, \dots, a_k, C_k)$ and so R preserves coloured trace sets. Thus $\underline{R}$ induces a consistent colouring and since the roots are related we find $g \equiv_c h$. If moreover R is rooted, then so is the induced colouring.

$\Leftarrow$: Consider a (rooted) consistent colouring such that the coloured trace sets of g and h are equal with respect to that colouring. Let R be the relation between nodes of g and h relating two nodes iff they have the same colour, then it is easy to see that R is a (rooted) branching bisimulation. $\square$

This characterization provides us with a clear intuition about what branching bisimulation actually is, since the difference between *inert* steps—not changing the state of the machine—and relevant τ -steps—that behave as common actions—is visualized immediately by the (change of) colours at the nodes. It follows that branching bisimulation equivalence preserves computations together with the potentials in all intermediate states that are passed through.

Another way of looking at the colouring of a graph is the following. Since trace-equivalence is too weak to characterize branching bisimilarity we can add more information to traces in order to distinguish between processes. Consider the following definition:

DEFINITION 3.7 For ordinals α the α -trace set of a graph g is defined as follows:

1. The α -trace set of a node r of g is the set of all γ -traces of r , for $\gamma < \alpha$.
2. An α -trace of r is made of a sequence $(T_0, a_1, T_1, a_2, \dots, a_k, T_k)$, where a_i are actions from Act and T_i are α -trace sets such that g has a path of the form $r_0 \rightarrow^{a_1} r_1 \rightarrow^{a_2} \dots \rightarrow^{a_k} r_k$ and r_i has α -trace set T_i , by replacing all subsequences $(T, \tau, T, \tau, \dots, \tau, T)$ by T .
3. The α -trace set of g is the α -trace set of its root.

Note that all 0-trace sets are empty, and the 1-trace set of g is just the set of its concrete traces from which τ 's have been left out. Two graphs g and h are α -trace equivalent—notation $g \approx_\alpha h$ —if they have the same α -trace set. Let us say that they are *hypertrace equivalent*—notation $g \approx h$ —if $g \approx_\alpha h$ for all ordinals α . Note that if $\lambda < \alpha$ then $g \approx_\alpha h$ implies $g \approx_\lambda h$. From this it immediately follows that if $G' \subseteq G$ is a set of process graphs then on G' the notion of α -trace equivalence stabilizes for some ordinal—i.e. there exists a *closure ordinal* α such that, for $g, h \in G'$, $g \approx h$ iff $g \approx_\alpha h$. It will follow from the proof of Theorem 3.3 that the smallest ordinal with $(g \approx_\alpha h \Leftrightarrow g \approx_{\alpha+1} h)$ is a closure ordinal. Furthermore if G has cardinality β then β

must be a closure ordinal. Next we prove that hypertrace equivalence coincides with coloured trace equivalence:

THEOREM 3.3 $g \approx h$ if and only if $g \equiv_c h$.

PROOF $\Rightarrow$: Let G' be a set of process graphs containing g, h and all their subgraphs and let α be the smallest ordinal such that, for $g', h' \in G'$, $g' \approx_\alpha h'$ iff $g' \approx_{\alpha+1} h'$. If $g' \approx_{\alpha+1} h'$ then by definition g' and h' have the same set of γ -traces ($\gamma \leq \alpha$). Since α -traces are recognizable from their form, this implies that g' and h' must have the same α -traces. Consider the colouring on g and h in which every node is coloured with its own α -trace set. Now a coloured trace $(C_0, a_1, C_1, a_2, \dots, a_k, C_k)$ of a node r is precisely an α -trace and by definition of α we have that r and r' have the same α -trace set only if they have the same α -traces, i.e. they have the same colour only if they have the same coloured traces. Hence the colouring is consistent.

Now $g \approx h \Rightarrow g \approx_{\alpha+1} h \Rightarrow g$ and h have the same coloured traces $\Rightarrow g \equiv_c h$.

$\Leftarrow$: Take a consistent colouring on g and h such that the roots of g and h have the same colour. Then with transfinite induction on γ it is easy to prove that equally coloured nodes have the same γ -traces for all ordinals γ . $\square$

Hence we find that $\approx$ is equivalent to $\equiv_c$, and hence to $\equiv_b$ (Theorem 3.2). Note that compared to *readiness semantics* (cf. OLDEROG & HOARE [144]), *possible-futures semantics* (cf. ROUNDS & BROOKES [159]) and *ready trace semantics* (cf. BAETEN, BERGSTRA & KLOP [14]) in an α -trace ($\alpha \geq 1$) we keep track of a lot more information. Apart from just all one-step exits from the endstate of a partial execution we are now able to see all traces (and higher traces) that can be chosen at every intermediate state during the execution.

The notion of hypertrace equivalence gives us an indication of the amount of extra information that is needed to turn trace equivalence into branching bisimulation equivalence. Furthermore, it provides us with an idea of how to build a consistent colouring on a set of graphs by distinguishing more and more between nodes.

A construction similar to Definition 3.7 was used by MILNER [137,138] to characterize observation equivalence in the spirit of Definition 1.7.

As a tool for further analysis we have the following proposition:

PROPOSITION 3.4 *It is possible to colour the nodes of a root unwound process graph g in such a way that two nodes have the same colour iff they can be related by a rooted branching autobisimulation on g (relating g with itself). This colouring is rooted and consistent.*

PROOF For every root unwound process graph g the largest rooted branching autobisimulation on g is an equivalence relation on the nodes. It follows from the proof of Theorem 3.2 that every node can be labelled with its equivalence class as a colour, in order to obtain a rooted consistent colouring. $\square$

This colouring of a graph is called its *canonical colouring*. Note that two different nodes r and s of a root unwound process graph g have the same colour with respect to its canonical colouring if and only if $r, s \neq \text{root}(g)$ and $(g)_r \cong_b (g)_s$ (the subgraph $(g)_r$ of g with root r is defined in the beginning of this section; furthermore, remember that in the canonical colouring root nodes have colours different from those of internal nodes). In this case we say that r and s are *rooted branching bisimilar*.

DEFINITION 3.8 A root unwound graph is said to be in *normal form* if each node has a different colour with respect to its canonical colouring and it has no τ -loops $r \rightarrow^\tau r$.

Next we show that each root unwound process graph is rooted branching bisimilar with exactly one normal form (up to isomorphism).

DEFINITION 3.9 Let g be a root unwound process graph and consider its canonical colouring with colour set C . Let $N(g)$ - the *normal form* of g - be the graph which can be found from g by contracting all nodes with the same colour and removing τ -loops. To be precise:

1. $N(g)$ has colours $C \in C$ as its nodes.
2. $N(g)$ has an edge $C \rightarrow^a C'$ ($a \neq \tau$) iff g has an edge $r \rightarrow^a r'$ such that $C(r)=C$ and $C(r')=C'$, where $C(r)$ denotes the colour of the node r .
3. $N(g)$ has an edge $C \rightarrow^\tau C'$ iff $C \neq C'$ and g has an edge $r \rightarrow^\tau r'$ with $C(r)=C$ and $C(r')=C'$.

PROPOSITION 3.5 For all root unwound process graphs g : $g \cong_{rb} N(g)$.

PROOF Consider the canonical colouring on g , and the trivial colouring on $N(g)$ in which each node (being a colour from C) is labelled by itself. Let R be the relation relating nodes from g and $N(g)$ iff they have the same colour. Now it follows directly from the construction above that R is a rooted branching bisimulation between g and $N(g)$. $\square$

So in every rooted branching bisimulation equivalence class of root unwound process graphs there is a normal form. We proceed by showing that up to isomorphism there is only one.

DEFINITION 3.10 A *graph isomorphism* is a bijective relation R between the nodes of two process graphs g and h such that:

1. the roots of g and h are related by R
2. if $R(r, s)$ and $R(r', s')$ then $r \rightarrow^a r'$ is an edge in g iff $s \rightarrow^a s'$ is an edge in h ($a \in \text{Act}$).

Note that a graph isomorphism is just a bijective bisimulation, or a bijective branching bisimulation for that matter. Two graphs are isomorphic—notation $g \cong h$ —iff there exists an isomorphism between them. In that case g and h only differ with respect to the identity of the nodes. Note that $\cong$ is a congruence relation on process graphs.

THEOREM 3.6 (normal form theorem)

Let g and h be root unwound graphs that are in normal form.

Then $g \equiv_{rb} h$ if and only if $g \equiv h$.

PROOF This follows since any bisimulation R : $g \equiv_{rb} h$ must be bijective:

- (i) it is surjective because every node in g or h can be reached from the root; hence by Definition 1.4 every node is related to some node in the other graph;
- (ii) it is injective since every node is related with at most one other node: if two different nodes in g are related to the same node in h , then these two are also related by a branching autobisimulation on g , and so with respect to the canonical colouring they have the same colour. But then by Definition 3.8 the nodes are identical, which is a contradiction. $\square$

Theorem 3.6 says that each equivalence class in $G/\equiv_{rb}$ can be represented by one special element: its normal form. It follows that $g \equiv_{rb} h$ if and only if $N(g) \equiv N(h)$.

4. COMPLETENESS PROOFS

In this section we will present the proofs of the completeness Theorems 2.6, 2.7, 2.8 and 2.5. By means of a rather trivial adaptation of the contents of this section one obtains the completeness theorems for CCS and ACP_τ (Theorems 2.12 and 2.14). The basic idea in these proofs is to establish a *graph rewriting system* on finite process graphs, which is confluent and terminating. We prove that (i) normal forms with respect to the graph rewriting system are normal forms in the sense of Definition 3.8, hence two normal forms are bisimilar iff they are equal (i.e. isomorphic). Furthermore we prove that every rewriting step in the system (ii) preserves bisimulation, and (iii) corresponds to a proof step in the theory. Then we conclude:

- two finite graphs are bisimilar iff they have the same normal form
- if two graphs have the same normal form then the corresponding terms can be proved equal.

To start with, let us consider some definitions.

DEFINITION 4.1 Let $H \subseteq G$ be the set of *finite* process graphs and $H^+ \subseteq G_{BPA}$ the set of finite, non-trivial process graphs. Here a process graph is *finite* if it has only finitely many paths.

Note that finite process graphs are acyclic and thus certainly root-unwound, and contain only finitely many nodes and edges. Later on, we will establish a correspondence between graphs from H^+ and closed terms in BPA_τ , i.e. the signature of BPA together with the extra constant τ . Below we will use the results from the previous section, starting from Proposition 3.4.

DEFINITION 4.2 A pair (r,s) of nodes in a process graph g is called a pair of *double nodes* if $r \neq s$, $r,s \neq \text{root}(g)$ and for all nodes t and labels $a \in \text{Act}$: $r \xrightarrow{a} t \iff s \xrightarrow{a} t$.

DEFINITION 4.3 An edge $r \xrightarrow{\tau} s$ in a process graph g is called *manifestly inert* if $r \neq \text{root}(g)$ and for all nodes t and labels $a \in \text{Act}$ such that $(a,t) \neq (\tau,s)$: $r \xrightarrow{a} t \implies s \xrightarrow{a} t$.




Figure 10. A pair of double nodes (left) and a manifestly inert τ -step.

Note that for finite process graphs g , the requirement $r,s \neq \text{root}(g)$ in Definition 4.2 is redundant. A τ -edge in a root unwound graph g is *inert* if it is between two rooted branching bisimilar nodes (i.e. nodes that have the same colour in the canonical colouring of g). For root unwound graphs it is easily checked that if (r,s) is a pair of double nodes or if $r \xrightarrow{\tau} s$ is manifestly inert, then r and s are rooted branching bisimilar (though it need not be the case that $(g)_r$ and $(g)_s$ are rooted branching bisimilar). As one can see from Figure 10, it is essential in the Definitions 4.2 and 4.3 that this can be found by investigating the outgoing edges only up to one level. For this reason, in Definition 4.3 the τ -step is called *manifestly inert*, since it can be recognized as such. On H , sharing of double nodes and contraction of manifestly inert τ -steps turns out to be strong enough to reduce a graph to its normal form. This means that in the reduction process all rooted branching bisimilar nodes become *manifestly* rooted branching bisimilar.

THEOREM 4.1 A graph $g \in H$ without double nodes or manifestly inert edges is in normal form.

PROOF Let $g \in H$ be a finite graph which is not in normal form. Then with respect to its canonical colouring (Proposition 3.4) it has at least one pair of different nodes with the same colour. Now define the *depth* $d(s)$ of a node s as the number of edges in the longest path starting from s , and the *combined depth* of two nodes as the sum of

their depths. Choose a pair (r,s) of different equally coloured nodes in g with minimal combined depth. Consequently we have:

(*) if r' and s' have the same colour and $d(r') + d(s') < d(r) + d(s)$ then $r'=s'$.

Without loss of generality assume $d(s) \leq d(r)$. Then we prove the following two statements:

1. if $r \rightarrow^a t$ ($a \in \text{Act}$) is an edge in g and $(a,t) \neq (\tau,s)$, then $s \rightarrow^a t$ is an edge in g
2. if $s \rightarrow^a t$ ($a \in \text{Act}$) is an edge in g , then either $r \rightarrow^\tau s$ or $r \rightarrow^a t$ is an edge in g .

From these two statements we find that if $r \rightarrow^\tau s$ is an edge in g then it is manifestly inert, and if $r \rightarrow^\tau s$ is not an edge in g , then (r,s) is a pair of double nodes, which proves our theorem. Note that since r and s are different equally coloured nodes, they both must be different from the root.

ad 1: Let $r \rightarrow^a t$ be an edge in g and $(a,t) \neq (\tau,s)$. Since r and s have the same colour (hence the same coloured traces) we find that either $a=\tau$ and t has the same colour as r and s , or s has the coloured trace $(C(r), a, C(t))$. In the first case it follows from $d(t) < d(r)$ and (*) that $t=s$, which is in contradiction with our assumption $(a,t) \neq (\tau,s)$. So s has a coloured trace $(C(r), a, C(t))$. Suppose that $s \rightarrow^\tau u$ for a node u with colour $C(u)=C(s)=C(r)$, then it follows from $d(u) < d(s)$ and (*) that $u=r$, contradicting $d(u) < d(s) \leq d(r)$. Hence there is a node u such that $s \rightarrow^a u$ and $C(t)=C(u)$, and since $d(t) + d(u) < d(r) + d(s)$ we conclude from (*) that $t=u$. Hence $s \rightarrow^a t$ is an edge in g .

ad 2: Let $s \rightarrow^a t$ be an edge in g . If $C(t)=C(s)=C(r)$ then it follows from (*) and $d(t) < d(s)$ that $r=t$, in contradiction with $d(t) < d(s) \leq d(r)$. So $(C(s), a, C(t))$ is a coloured trace of s , and since r and s have the same colour $(C(s), a, C(t))$ is a coloured trace of r as well. Now if r has an outgoing τ -edge $r \rightarrow^\tau u$ to a node with the same colour $C(r)$, then it follows from $d(u) + d(s) < d(r) + d(s)$ and (*) that $u=s$. If r has no such edge, then it has an edge $r \rightarrow^a u$ with $C(u)=C(t)$, and since $d(u) + d(t) < d(r) + d(s)$ we find that $u=t$. Thus we proved that either $r \rightarrow^\tau s$ or $r \rightarrow^a t$, which proves (2). $\square$

Theorem 4.1 tells us that all we need do in order to turn a finite graph g into its normal form is to repeatedly unify its pairs of double nodes and contract its manifestly inert edges. In the case of finite graphs this can be done in finitely many steps as follows:

DEFINITION 4.4 For any graph $g \in H$ the rewriting relation $\rightarrow_H$ is defined by the following two one-step reductions:

1. *sharing* a pair of double nodes (r,s) : replace all edges $t \rightarrow^a r$ by $t \rightarrow^a s$ (if not already there, otherwise just remove $t \rightarrow^a r$) and remove r together with all its outgoing edges from g ;
2. *contracting* a manifestly inert step $r \rightarrow^\tau s$: replace all edges $t \rightarrow^a r$ by $t \rightarrow^a s$ (if not already there, otherwise just remove $t \rightarrow^a r$) and remove r together with all its outgoing edges from g .

PROPOSITION 4.2 *The rewriting relation $\rightarrow_H$ has the following properties:*

- i. H as well as H^+ are closed under applications of $\rightarrow_H$
- ii. if $g \rightarrow_H h$ then $g \equiv_{rb} h$
- iii. $\rightarrow_H$ is confluent and terminating.

PROOF (i) In applications of $\rightarrow_H$ the root is never removed and in the resulting graph all nodes remain accessible from the root. It is never the case that two edges with the same label appear between the same two nodes. The graph also remains finite (and non-trivial).

(ii) Suppose (r,s) is a pair of double nodes or $r \rightarrow^\tau s$ is a manifestly inert edge in g , and $g \rightarrow_H h$ identifies the nodes r and s ($=$ removes the node r). Let I be the identity relation on the nodes of h then $I \cup \{(r,s)\}$ is a rooted branching bisimulation between g and h . This is easy to prove from the Definitions 4.2 and 4.3.

(iii) $\rightarrow_H$ is terminating since it decreases the number of nodes, and every finite process graph has finitely many nodes. Next, suppose g has two normal forms n and n' , then by the definition of $\rightarrow_H$ n and n' are without pairs of double nodes and without manifestly inert edges. Thus by Theorem 4.1 n and n' are in normal form. By (ii) it follows that $n \equiv_{rb} n'$ and hence by Theorem 3.6 (normal form theorem) we have $n \equiv n'$. $\square$

Next we will establish a correspondence between finite non-trivial graphs and closed BPA_τ -terms, such that the graph reductions of Definition 4.4 correspond to proof steps in $BPA + H1,2$.

Write $s \equiv_\Gamma t$ for $\Gamma \vdash s=t$ saying that s and t are equal *modulo* applications of axioms from Γ and the standard axioms for equality (reflexivity, commutativity, transitivity and replacement). It is quite easy to turn finite non-trivial graphs into BPA_τ -terms as follows. Let $T(BPA_\tau)$ be the set of closed BPA_τ terms.

DEFINITION 4.5 Let $\langle \cdot \rangle: H^+ \rightarrow T(BPA_\tau)$ be a mapping that satisfies

$$\langle g \rangle = \sum_{\substack{r(g) \rightarrow^a s \text{ is an edge in } g; \\ s \text{ not an endnode}}} a \cdot \langle (g)_s \rangle + \sum_{\substack{r(g) \rightarrow^b s \text{ is an edge in } g; \\ s \text{ is an endnode}}} b.$$

Here $r(g)$ denotes the root node of $g \in H^+$ and if p_i is a BPA_τ -term for $i \in I$, with $I = \{i_1, \dots, i_n\}$ a finite non-empty set of indices, then $\sum_{i \in I} p_i$ denotes a BPA_τ -term $p_{i_1} + \dots + p_{i_n}$. Note that the notation $\sum_{i \in I} p_i$ does not determine the order and association of the terms p_i .

If $g \in H^+$, $r(g) \rightarrow^a s$ is an edge in g , and s is not an endnode, then $(g)_s \in H^+$. Furthermore, since $g \in H^+$ is finite, $r(g)$ has only finitely many outgoing edges, so the requirement of Definition 4.5 is well-defined. Moreover, with induction on the depth of its arguments, one easily proves that a mapping that meets this requirement exists. However, for $g \in H^+$, this requirement determines $\langle g \rangle$ is only modulo A1-A2.

PROPOSITION 4.3 *If $g, h \in H^+$ and $g \cong h$, then $A1-A2 \vdash \langle g \rangle = \langle h \rangle$.*

PROOF Trivial. $\square$

DEFINITION 4.6 The denotation $[p]$ of a BPA_τ -term p in the graph domain G , is defined by:

$$\begin{aligned} [a] &= a_G && \text{for } a \in A_\tau \\ [x + y] &= [x] +_G [y] \\ [x \cdot y] &= [x] \cdot_G [y] \end{aligned}$$

where a_G , $+_G$ and $\cdot_G$ are the interpretations in G , of the constants and operators a , $+$ and $\cdot$ of BPA_τ , as defined in Definition 2.1.

Now it turns out that terms of the form $\langle g \rangle$ (for $g \in H^+$) are all of a specific shape, and for terms of this shape, $\langle \cdot \rangle$ is a left-inverse of $[\cdot]$, modulo A1-A2. Consider the following definition:

DEFINITION 4.7 The set BT of *basic terms* over BPA_τ is inductively defined as follows:

1. For all $a \in \text{Act}$ we have $a \in \text{BT}$;
2. If $p, q \in \text{BT}$ then $(p + q) \in \text{BT}$ and for all $a \in \text{Act}$: $a \cdot p \in \text{BT}$.

LEMMA 4.4 *For $g \in H^+$, $\langle g \rangle$ is a basic term and if $p \in \text{BT}$, then $\langle [p] \rangle \equiv_{A1,2} p$.*

PROOF With induction to the structure of terms:

- If $p = a$ ($a \in \text{Act}$) then $[p]$ is the one-edge graph labelled with a , and so $\langle [p] \rangle = p$.
- If $p = a \cdot u$ for some basic term u , then $[p]$ is the graph with an edge labelled a followed by $[u]$.
Then, $\langle [p] \rangle = a \cdot \langle [u] \rangle$ and so by induction we find that $\langle [p] \rangle \equiv_{A1,2} a \cdot u$.
- Suppose $p = u + v$. One can easily see that for graphs g and h : $\langle g +_G h \rangle \equiv_{A1,2} \langle g \rangle + \langle h \rangle$.
Then: $A1-A2 \vdash \langle [u + v] \rangle = \langle [u] \rangle + \langle [v] \rangle = u + v$ (by induction). $\square$

LEMMA 4.5 (elimination)

For every closed BPA_τ -term p there exists a basic term q such that $A4-A5 \vdash p = q$.

PROOF By induction on the structure of p .

- If $p = a$ ($a \in \text{Act}$) then p is a basic term.
- If $p = u \cdot v$ and Lemma 4.5 can be proved for all terms smaller than p , then there exist basic terms u' and v' such that $A4-A5 \models u = u'$, $v = v'$. Now suppose u' has the form $(\sum_i a_i \cdot w_i + \sum_j b_j)$, then we find:

$$\begin{aligned} A4-A5 \models p &= u' \cdot v' = (\sum_i a_i \cdot w_i + \sum_j b_j) \cdot v' = \\ &= \sum_i (a_i \cdot w_i) \cdot v' + \sum_j b_j \cdot v' \quad (\text{by Axiom A4}) \\ &= \sum_i a_i \cdot (w_i \cdot v') + \sum_j b_j \cdot v' \quad (\text{by Axiom A5}) \\ &= \sum_i a_i \cdot q_i + \sum_j b_j \cdot v' \quad \text{for some basic terms } q_i \quad (\text{by induction}) \end{aligned}$$

which is a basic term again.

- If $p = u + v$ then $A4-A5 \models p = u' + v'$ for basic terms u' and v' , and the sum of two basic terms is again a basic term. $\square$

PROPOSITION 4.6 For all closed BPA_{τ} -terms p we have: $A1-A2+A4-A5 \vdash \langle [p] \rangle = p$.

PROOF If ' $p=q$ ' is an instantiation of axioms A4 or A5 (possibly in a context) then $\langle [p] \rangle \equiv_{A1,2} \langle [q] \rangle$. Now the proposition follows immediately from the Lemmas 4.4 and 4.5. $\square$

This concludes the establishment of a correspondence between H^+ and $T(BPA_{\tau})$. Next we will show that every rewriting step on H^+ corresponds to a proof step in $BPA + H1-H2$.

LEMMA 4.7 Let (r,s) be a pair of double nodes or $r \rightarrow^{\tau} s$ be a manifestly inert τ -step in a process graph g , such that neither r nor s are endnodes, and let $a \in \text{Act}$. Then we have: $BPA + H1-H2 \vdash a \cdot \langle (g)_r \rangle = a \cdot \langle (g)_s \rangle$.

PROOF In case (r,s) is a pair of double nodes r has an edge $r \rightarrow^a t$ iff s has an edge $s \rightarrow^a t$ and so $\langle (g)_r \rangle \equiv_{A1,2} \langle (g)_s \rangle$, hence $a \cdot \langle (g)_r \rangle = a \cdot \langle (g)_s \rangle$.

In case $r \rightarrow^{\tau} s$ is a manifestly inert τ -step we distinguish two subcases: First assume that r has more outgoing edges than only $r \rightarrow^{\tau} s$. Then there must be basic terms p and q such that

- (1) $\langle (g)_r \rangle \equiv_{A1,2} \tau \cdot \langle (g)_s \rangle + p$
- (2) $\langle (g)_s \rangle \equiv_{A1,2} p + q$.

So we derive:

$$\begin{aligned} A1,2 + H2 \vdash a \cdot \langle (g)_r \rangle &= a \cdot (\tau \cdot \langle (g)_s \rangle + p) \text{ (by (1))} \\ &= a \cdot (\tau \cdot (p + q) + p) \text{ (by (2))} \\ &= a \cdot (p + q) \text{ (by applying H2)} \\ &= a \cdot \langle (g)_s \rangle \text{ (by (2)).} \end{aligned}$$

In case r has no more outgoing edges than $r \rightarrow^{\tau} s$ we have $\langle (g)_r \rangle = \tau \cdot \langle (g)_s \rangle$, hence

$$A5 + H1 \vdash a \cdot \langle (g)_r \rangle = a \cdot (\tau \cdot \langle (g)_s \rangle) = (a \cdot \tau) \cdot \langle (g)_s \rangle = a \cdot \langle (g)_s \rangle. \quad \square$$

PROPOSITION 4.8 If $g \rightarrow_H h$ then $BPA + H1-H2 \vdash \langle g \rangle = \langle h \rangle$.

PROOF On H the rewriting relation $\rightarrow_H$ can be decomposed in the following elementary reductions:

Take a pair of double nodes (r,s) or a manifestly inert τ -step $r \rightarrow^{\tau} s$ and replace one edge $t \rightarrow^a r$ by $t \rightarrow^a s$ (if not already there, otherwise just remove $t \rightarrow^a r$) and if r has no more incoming edges remove r together with all its outgoing edges from g . So it suffices to prove that if h is obtained from g by means of such an elementary reduction, we have $\langle g \rangle \equiv_{\Gamma} \langle h \rangle$, where $\Gamma = BPA + H1-H2$. From Definition 4.5 it follows that it even suffices to prove $\langle (g)_t \rangle \equiv_{\Gamma} \langle (h)_t \rangle$.

- First consider the case that neither r nor s are endnodes and there is no edge $t \rightarrow^a s$ in g . Then $\langle (g)_t \rangle \equiv_{A1,2} a \cdot \langle (g)_r \rangle + p$ for certain basic term p . Lemma 4.7 says $a \cdot \langle (g)_r \rangle \equiv_{\Gamma} a \cdot \langle (g)_s \rangle$, hence $\langle (g)_t \rangle \equiv_{\Gamma} a \cdot \langle (g)_s \rangle + p \equiv_{A1,2} \langle (h)_t \rangle$.

- In case $t \rightarrow^a s$ is an edge in g , and r, s are still assumed not to be endnodes we have $\langle (g)_t \rangle \equiv_{A1,2} a \cdot \langle (g)_r \rangle + a \cdot \langle (g)_s \rangle + p \equiv_{\Gamma} a \cdot \langle (g)_s \rangle + a \cdot \langle (g)_s \rangle + p \equiv_{A2,3} a \cdot \langle (g)_s \rangle + p \equiv_{A1,2} \langle (h)_t \rangle$.
- If (r, s) is a pair of double nodes than r is an endnode iff s is. In this case we have $\langle (g)_t \rangle \equiv_{A1,2} a + p \equiv_{A1,2} \langle (h)_t \rangle$ if $t \rightarrow^a s$ is not an edge in g and $\langle (g)_t \rangle \equiv_{A1,2} a + a + p \equiv_{A2,3} a + p \equiv_{A1,2} \langle (h)_t \rangle$ otherwise.
- Finally if $t \rightarrow^a s$ is a manifestly inert τ -edge and s is an endnode in g , we have $\langle (g)_t \rangle \equiv_{A1,2} a \cdot \tau + p \equiv_{H1} a + p \equiv_{A1,2} \langle (h)_t \rangle$ if $t \rightarrow^a s$ is not an edge in g and $\langle (g)_t \rangle \equiv_{A1,2} a \cdot \tau + a + p \equiv_{H1} a + a + p \equiv_{A2,3} a + p \equiv_{A1,2} \langle (h)_t \rangle$ otherwise. $\square$

Now we are in the position to prove the completeness of $BPA + H1-H2$ with respect to $G_{BPA/\equiv_{rb}}$:

PROOF OF THEOREM 2.6: (soundness) The fact that $(G_{BPA/\equiv_{rb}}, +, \cdot, \text{Act})$ is a model for $BPA + H1-H2$ follows easily by inspection of the axioms of $BPA + H1-H2$.

(completeness) Let $(G_{BPA/\equiv_{rb}}, +, \cdot, \text{Act}) \models p = q$ for two closed BPA_{τ} -terms p, q , then by definition $[p] \equiv_{rb} [q]$. Let g and h be the unique normal forms of $[p]$ and $[q]$ with respect to $\rightarrow_H$. By Proposition 4.2 we find $g \equiv_{rb} [p] \equiv_{rb} [q] \equiv_{rb} h$. From Theorem 4.1 it follows that g and h must be in normal form in the sense of Section 3 and by the normal form theorem (Theorem 3.6) it then follows that $g \equiv h$. Thus we find $BPA + H1-H2 \vdash p = \langle [p] \rangle = \langle g \rangle = \langle h \rangle = \langle [q] \rangle = q$ using Propositions 4.3, 4.6 and 4.8. So $BPA + H1-H2$ is a complete axiomatization of $G_{BPA/\equiv_{rb}}$. $\square$

Next we will prove the other completeness theorems, using the earlier results in this section. In fact we will extend the graph rewriting system to one which is 'typical' for the corresponding bisimulation relation. The rewrite rules which are added to the system are derived from Figure 1: in case of η -bisimulation we will *saturate* the graph by exhaustively adding edges of the kind of Figure 1 (c), whereas in the case of delay bisimulation we add edges as in Figure 1 (b). For τ -bisimulation we do both. This way we obtain normal forms which are saturated and which turn out to be unique modulo rooted branching bisimulation. From there we establish the completeness result precisely in the same way as before.

DEFINITION 4.8 Let $a \in \text{Act}$, then:

1. The rewriting relation $\rightarrow_{\eta}$ is defined on H by the rule:
if a graph has a path $s \rightarrow^a s_1 \rightarrow^{\tau} s'$ without an edge $s \rightarrow^a s'$ then add $s \rightarrow^a s'$.
2. The rewriting relation $\rightarrow_d$ is defined on H by the rule:
if a graph has a path $s \rightarrow^{\tau} s_1 \rightarrow^a s'$ without an edge $s \rightarrow^a s'$ then add $s \rightarrow^a s'$.
3. Furthermore, we set: $\rightarrow_{\tau} = \rightarrow_{\eta} \cup \rightarrow_d$.

Applications of $\rightarrow_{\eta}$, $\rightarrow_d$ or $\rightarrow_{\tau}$ are referred to as *saturation steps* (cf. BERGSTRA & KLOP [27]).

PROPOSITION 4.9 *The relations $\rightarrow_\eta$, $\rightarrow_d$ and $\rightarrow_\tau$ satisfy the following properties:*

- i. H as well as H^+ are closed under applications of $\rightarrow_\eta$, $\rightarrow_d$ and $\rightarrow_\tau$
- ii. $\rightarrow_\eta$, $\rightarrow_d$ and $\rightarrow_\tau$ are confluent and terminating.

PROOF (i) Directly from Definition 4.8.

(ii) (termination) Let $g \in H$. Let $n(g)$ be the (finite) number of nodes in g , $l(g)$ be the number of labels and $e(g)$ be the number of edges in g . Note that $n(g)$ and $l(g)$ are not changed by $\rightarrow_\eta$, $\rightarrow_d$ and $\rightarrow_\tau$ whereas $e(g)$ increases with every saturation step. Since g is finite we find that $e(g) < n(g) \times l(g) \times n(g)$ and so $n(g) \times l(g) \times n(g) - e(g)$ is positive and decreasing with the number of saturation steps.

(confluence) $\rightarrow_\eta$, $\rightarrow_d$ and $\rightarrow_\tau$ do not eliminate redexes. $\square$

So from Proposition 4.9 we find that any graph $g \in H$ has unique normal forms with respect $\rightarrow_\eta$, $\rightarrow_d$ and $\rightarrow_\tau$. These are written as $H(g)$, $D(g)$ and $T(g)$ and (in that order) are called η -, d - and τ -saturated. The latter is also often referred to as the *transitive closure* of τ -steps. Furthermore, saturation preserves the corresponding bisimulation:

PROPOSITION 4.10 *For all $g, h \in H$:*

- i. if $g \rightarrow_\eta h$ then $g \equiv_{\tau\eta} h$
- ii. if $g \rightarrow_d h$ then $g \equiv_{\tau d} h$
- iii. if $g \rightarrow_\tau h$ then $g \equiv_{\tau\tau} h$.

The proof of the Proposition 4.10 is straightforward.

THEOREM 4.11 (normal form theorem) *Let $g, h \in H$, then*

- i. if g and h are η -saturated process graphs, then $g \equiv_{\tau\eta} h$ if and only if $g \equiv_{\tau b} h$
- ii. if g and h are d -saturated process graphs, then $g \equiv_{\tau d} h$ if and only if $g \equiv_{\tau b} h$
- iii. if g and h are τ -saturated process graphs, then $g \equiv_{\tau\tau} h$ if and only if $g \equiv_{\tau b} h$.

PROOF We will only prove (i). The other cases proceed in the same way.

Suppose that $R: g \equiv_{\tau\eta} h$ then it is sufficient to prove that R is a rooted branching bisimulation:

- (i) The roots of $H(g)$ and $H(h)$ are related and (iii) R satisfies the root condition.
- (ii) If $R(r, s)$ and $r \xrightarrow{a} r'$ then either $a = \tau$ and $R(r', s)$, or $s \Rightarrow s_1 \xrightarrow{a} s_2 \Rightarrow s'$ such that $R(r, s_1)$ and $R(r', s')$. Let $t_1, \dots, t_k$ be such that $s_2 = t_0 \xrightarrow{\tau} t_1 \xrightarrow{\tau} \dots \xrightarrow{\tau} t_k = s'$ ($k \geq 0$) then since g and h are η -saturated there are edges $s_1 \xrightarrow{a} t_i$ and so there is a path $s \Rightarrow s_1 \xrightarrow{a} s'$. $\square$

COROLLARY

- i. $g \equiv_{\tau\eta} h$ if and only if $H(g) \equiv_{\tau b} H(h)$ if and only if $N(H(g)) \equiv N(H(h))$
- ii. $g \equiv_{\tau d} h$ if and only if $D(g) \equiv_{\tau b} D(h)$ if and only if $N(D(g)) \equiv N(D(h))$
- iii. $g \equiv_{\tau\tau} h$ if and only if $T(g) \equiv_{\tau b} T(h)$ if and only if $N(T(g)) \equiv N(T(h))$.

PROOF It follows by Proposition 4.10 that $H(g) \equiv_{\tau\eta} g$, $D(g) \equiv_{\tau d} g$ and $T(g) \equiv_{\tau\tau} g$.

Now apply the normal form theorems 4.11 and 3.6. $\square$

So we find that in each r^* -bisimulation equivalence class of finite process graphs for $* \in \{\tau, \eta, d\}$ there is exactly one $*$ -saturated process graph up to rooted branching bisimulation and exactly one $*$ -saturated normal form up to isomorphism. In order to prove the completeness theorems we still need to prove that rewriting steps correspond to proof steps.

PROPOSITION 4.12 *For finite graphs g and h :*

- i. *If $g \rightarrow_\eta h$ then $A1-A3 + H1,3 \vdash \langle g \rangle = \langle h \rangle$*
- ii. *If $g \rightarrow_d h$ then $A1-A3 + T2 \vdash \langle g \rangle = \langle h \rangle$*
- iii. *If $g \rightarrow_\tau h$ then $A1-A3 + T1,3 \vdash \langle g \rangle = \langle h \rangle$.*

PROOF (i) If $r \rightarrow^a r' \rightarrow^\tau r'' \rightarrow$ is a path in g and $r \rightarrow^a r''$ is added in g to obtain h , then we find that

$$\begin{aligned} \langle (g)_r \rangle &\equiv_{A1-3} \langle (g)_r \rangle + a \cdot \langle (g)_r \rangle \text{ and} \\ \langle (g)_r \rangle &\equiv_{A1-3} \tau \cdot \langle (g)_{r''} \rangle + \langle (g)_r \rangle \text{ and hence:} \\ A1-A3 + H3 \vdash \langle (g)_r \rangle &= \langle (g)_r \rangle + a \cdot (\tau \cdot \langle (g)_{r''} \rangle + \langle (g)_r \rangle) = \\ &= \langle (g)_r \rangle + a \cdot (\tau \cdot \langle (g)_{r''} \rangle + \langle (g)_r \rangle) + a \cdot \langle (g)_{r''} \rangle \text{ (by H3)} = \\ &= \langle (g)_r \rangle + a \cdot \langle (g)_{r''} \rangle = \langle (h)_r \rangle. \end{aligned}$$

In case $r \rightarrow^a r' \rightarrow^\tau r''$ and r'' is an endnode we find:

$$\begin{aligned} A1-A3 + H1,3 \vdash \langle (g)_r \rangle &= \langle (g)_r \rangle + a \cdot (\tau + \langle (g)_{r''} \rangle) = \\ &= \langle (g)_r \rangle + a \cdot (\tau \cdot \tau + \langle (g)_{r''} \rangle) \text{ (by H1)} = \\ &= \langle (g)_r \rangle + a \cdot (\tau \cdot \tau + \langle (g)_{r''} \rangle) + a \cdot \tau \text{ (by H3)} = \\ &= \langle (g)_r \rangle + a = \langle (h)_r \rangle. \end{aligned}$$

From $A1-A3 + H3 \vdash \langle (g)_r \rangle = \langle (h)_r \rangle$ it easily follows that

$$A1-A3 + H3 \vdash \langle g \rangle = \langle h \rangle.$$

(ii) If $r \rightarrow^\tau r' \rightarrow^a r'' \rightarrow$ is a path in g and $r \rightarrow^a r''$ is added in g to obtain h , then:

$$\begin{aligned} \langle (g)_r \rangle &\equiv_{A1-3} \langle (g)_r \rangle + \tau \cdot \langle (g)_{r'} \rangle \text{ and} \\ \langle (g)_{r'} \rangle &\equiv_{A1-3} a \cdot \langle (g)_{r''} \rangle + \langle (g)_{r'} \rangle \text{ and hence:} \\ A1-A3 + T2 \vdash \langle (g)_r \rangle &= \langle (g)_r \rangle + \tau \cdot (a \cdot \langle (g)_{r''} \rangle + \langle (g)_{r'} \rangle) = \\ &= \langle (g)_r \rangle + a \cdot \langle (g)_{r''} \rangle \text{ (by T2 and A3)} = \langle (h)_r \rangle. \end{aligned}$$

In case $r \rightarrow^a r' \rightarrow^\tau r''$ and r'' is an endnode we simply leave out $\cdot \langle (g)_{r''} \rangle$ in the argument above.

Hence $A1-A3 + T2 \vdash \langle g \rangle = \langle h \rangle$.

(iii) Immediately from (i) and (ii). Note that $H1 = T1$ and $H3 = T3$. □

PROOFS OF THE THEOREMS 2.5, 2.7 AND 2.8

The soundness theorems follow easily after inspection of the axioms. Of the completeness theorems we only prove Theorem 2.7. The others proceed in the same way.

Let $(G_{BPA/\equiv_{\tau\eta}}, +, \cdot, \text{Act}) \models p=q$ for two closed BPA_τ -terms p, q , then by definition $[p] \equiv_{\tau\eta} [q]$. Let g and h be the unique normal forms of $[p]$ and $[q]$ with respect to $\rightarrow_\eta$. By Proposition 4.10 we find $g \equiv_{\tau\eta} [p] \equiv_{\tau\eta} [q] \equiv_{\tau\eta} h$. The graphs g and h must be η -saturated and by the normal form theorem (Theorem 4.11) it then follows that $g \equiv_{\tau b} h$. Thus we find $BPA + H1-H3 \vdash p = \langle [p] \rangle = \langle g \rangle = \langle h \rangle = \langle [q] \rangle = q$ using Propositions 4.6 and 4.12 and Theorem 2.6. So $BPA + H1-H3$ is a complete axiomatization of $G_{BPA/\equiv_{\tau\eta}}$. □

5. CORRESPONDENCE

Here we present a theorem which tells us that in quite a number of cases observation and branching bisimulation equivalence are the same. For instance, consider the practical applications where implementations are verified by proving them equal to some specification (after having abstracted from a set of unobservable actions of course). In many such cases, the *specification* does not involve any τ -steps at all: in fact all τ -steps that occur in the verification process originate from the abstraction procedure which is carried out on the implementation.

As it turns out, in all such cases there is no difference between observation and branching bisimulation equivalence. For this reason we may expect many verifications involving observation equivalence to be valid in the stronger setting of branching bisimulation as well. In particular this is the case for all protocol verifications in τ -bisimulation semantics known to the authors.

THEOREM 6.2 *Suppose g and h are two graphs, and g is without edges labelled with τ . Then:*

- i. $g \rightleftharpoons_{\tau} h$ if and only if $g \rightleftharpoons_b h$
- ii. $g \rightleftharpoons_{\tau\tau} h$ if and only if $g \rightleftharpoons_{rb} h$.

PROOF Let R be the *largest* (rooted) τ -bisimulation between g and h . We show that R is even a (rooted) branching bisimulation. Assume that $R(r,s)$ and $r \rightarrow^a r'$ is an edge in g , then either $a=\tau$ and $R(r',s)$ —contradicting the absence of τ -edges in g —or in h there is a path $s \Rightarrow s_1 \rightarrow^a s_2 \Rightarrow s'$ and $R(r',s')$. Assume $s \Rightarrow s_1$ has the form $s = v_0 \rightarrow^{\tau} v_1 \rightarrow^{\tau} \dots \rightarrow^{\tau} v_m = s_1$ ($m \geq 0$) then it follows from $s \rightarrow^{\tau} v_1$ and $R(r,s)$ that for some r_1 : $r \Rightarrow r_1$ and $R(r_1, v_1)$. Since g has no τ -edges we find that $r=r_1$. Repeating this argument m times we find that $R(r, v_i)$ and $R(r, s_1)$.

Furthermore, since $R(r, s_1)$ and $s_1 \rightarrow^a s_2$ we find that $r \rightarrow^a r''$ (g has no τ -steps) such that $R(r'', s_2)$. Since $s_2 = w_0 \rightarrow^{\tau} w_1 \rightarrow^{\tau} \dots \rightarrow^{\tau} w_n = s'$ it follows from the same argument as before that $R(r'', w_i)$ and $R(r'', s')$. Thus we find $R(r', s')$, $R(s', r'')$ and $R(r'', s_2)$ and since R is the largest rooted τ -bisimulation we have $R(r', s_2)$.

On the other hand, if $R(r,s)$ and $r \rightarrow^a r'$ is an edge in h , then either $a=\tau$ and $R(r',s)$ or directly $s \rightarrow^a s'$ such that $R(r',s')$, since g contains no τ -edges. $\square$

For η - instead of branching bisimulation equivalence this theorem was already proven in BAETEN & VAN GLABBEK [15]. From Theorem 2.15 we easily find that for graphs g and h :

$$g \text{ is without } \tau\text{-edges} \Rightarrow (\tau \cdot g \rightleftharpoons_{\tau\tau} \tau \cdot h \Rightarrow \tau \cdot g \rightleftharpoons_{rb} \tau \cdot h).$$

BOUALI [36] presents what is claimed to be "a counter example to the feelings in [4] about the fact that Branching Bisimulation and Weak Bisimulation coincide for a large class of processes". Here "[4]" is GROOTE & VAANDRAGER [106], who refer to the results of this section and confirm that they "are not aware of any protocol that can be verified in the setting of observation equivalence but not in the stronger setting of branching bisimulation equivalence". Bouali's counterexample concerns the

minimization of PETERSON's mutual exclusion algorithm [147]. Here *minimization* means finding an equivalent process that is as small as possible. In branching bisimulation semantics this yields a process with 17 states, whereas in weak bisimulation semantics a process with only 14 states is obtained. In our opinion this does not constitute a counterexample, as this is not a case of *verification*. It would be a verification if the minimal representation modulo weak bisimulation would be an acceptable *specification*. However, it has too many τ -transitions for this purpose. It is a process obtained as the *result* of a calculation, not anything one likely would start with. Much more plausible candidates for a specification are the minimizations modulo coupled simulation (9 states) or failure equivalence (4 states). More on this matter in the conclusion.

6. REFINEMENT

Virtually all semantic equivalences employed in theories of concurrency are—as in this thesis—defined in terms of *actions* that concurrent systems may perform. Mostly, these actions are taken to be *atomic*, meaning that they are considered not to be divisible into smaller parts. In this case, the defined equivalences are said to be based on *action atomicity*.

However, in the top-down design of distributed systems it might be fruitful to model processes at different levels of abstraction. The actions on an abstract level then turn out to represent complex processes on a more concrete level. This methodology does not seem compatible with non-divisibility of actions and for this reason PRATT [154], LAMPORT [125] and others plead for the use of semantic equivalences that are not based on action atomicity.

As indicated in CASTELLANO, DE MICHELIS & POMELLO [46], the concept of action atomicity can be formalized by means of the notion of *refinement of actions*. A semantic equivalence is *preserved under action refinement* if two equivalent processes remain equivalent after replacing all occurrences of an action a by a more complex process $r(a)$. In particular, $r(a)$ may be a sequence of two actions a_1 and a_2 . An equivalence is strictly based on action atomicity if it is not preserved under action refinement.

In the previous sections in this chapter we argued that Milner's notion of observation equivalence does not respect the branching structure of processes, and proposed the finer notion of *branching bisimulation equivalence* which does. In this section we moreover find, that observation equivalence is not preserved under action refinement, whereas branching bisimulation equivalence is.

From the Axioms T3 (see Table 2), it is easy to show why the notion of observation congruence is not preserved under refinement of actions: replacing the action a by the term bc , we obtain $bc(\tau x + y) = bc(\tau x + y) + bcx$, which obviously is not valid in $G/\equiv_{\tau}$. Applying T3, we *do* find $bc(\tau x + y) = b(c(\tau x + y) + cx)$, unfortunately denoting a different process however.

In this section we will prove that branching equivalence is preserved under refinement of actions, and so it allows us to look at actions as abstractions of much larger structures. We will present our result in the style of BPA, and indicate afterwards how our construction can be adapted to obtain refinement theorems in the style of CCS and ACP. Put $A = \text{Act} \setminus \{0\}$ (or $A = \text{Act} \setminus \{0, \surd\}$ if there are $\surd$ -labels around). Consider the following definitions.

DEFINITION 6.1 (substitution) Let $r: A \rightarrow G_{\text{BPA}}$ be a mapping from observable actions to graphs, and suppose $g \in G_{\text{BPA}}$. Then, the graph $r(g)$ can be found as follows.

For every edge $r \xrightarrow{a} r'$ ($a \in A$) in g , take a copy $\underline{r(a)}$ of $r(a)$ ($\in G_{\text{BPA}}$). Next, identify r with the root node of $\underline{r(a)}$, and r' with all endnodes of $\underline{r(a)}$, and remove the edge $r \xrightarrow{a} r'$.

Note that in this definition it is never needed to identify r and r' , since $r(a)$ is non-trivial. This way, the mapping r is extended to the domain G_{BPA} . Note that since $\tau \notin A$, τ -edges cannot be substituted by graphs. Finally, observe that every node in g is a node in $r(g)$.

DEFINITION 6.2 (preservation under action refinement) An equivalence $\approx$ on G_{BPA} is said to be *preserved under refinement of actions* if for every mapping $r: A \rightarrow G_{\text{BPA}}$, we have: $g \approx h \Rightarrow r(g) \approx r(h)$.

In other words, an equivalence $\approx$ is preserved under refinement if it is a congruence with respect to every substitution operator r .

Starting from a relation $R: g \rightleftharpoons_{\text{rb}} h$, we construct a branching bisimulation $r(R): r(g) \rightleftharpoons_{\text{rb}} r(h)$, proving that preserving branching congruence, every edge with a label from A can be replaced by a root unwound non-trivial graph.

DEFINITION 6.3 Let $r: A \rightarrow G_{\text{BPA}}$ be a mapping from observable actions to graphs, $g, h \in G_{\text{BPA}}$ and $R: g \rightleftharpoons_{\text{rb}} h$. Now $r(R)$ is the smallest relation between nodes of $r(g)$ and $r(h)$, such that:

1. $R \subseteq r(R)$.
2. If $r \xrightarrow{a} r'$ and $s \xrightarrow{a} s'$ ($a \in A$) are edges in g and h such that $R(r, s)$ and $R(r', s')$, and both edges are replaced by copies $\underline{r(a)}$ and $\underline{r'(a)}$ of $r(a)$ respectively, then nodes from $\underline{r(a)}$ and $\underline{r'(a)}$ are related by $r(R)$ iff they are copies of the same node in $r(a)$.

Edges $r \xrightarrow{a} r'$ and $s \xrightarrow{a} s'$ ($a \in A$) such that $R(r, s)$ and $R(r', s')$, will be called *related* by R , as well as the copies $\underline{r(a)}$ and $\underline{r'(a)}$ that are substituted for them. Observe, that on nodes from g and h the relation $r(R)$ is equal to R . Note that if $r(R)(r, s)$, then r is a node in g iff s is a node in h .

THEOREM 6.1 (refinement) *Branching congruence is preserved under refinement of actions.*

PROOF We prove that $R: g \rightleftharpoons_{rb} h \Rightarrow r(R): r(g) \rightleftharpoons_{rb} r(h)$ by checking the requirements.

For convenience, in the definition of branching equivalence (Definition 1.4), we omit the requirement of the existence of a path $s_2 \Rightarrow s'$, as it is redundant (see the remark just after Definition 1.6). Then we find:

- i. The root nodes of $r(g)$ and $r(h)$ are related by $r(R)$.
- ii. Assume $r(R)(r,s)$ and in $r(g)$ there is an edge $r \rightarrow^a r'$. Then there are two possibilities (similarly in case $r \rightarrow^a r'$ stems from $r(h)$):

(1) The nodes r and s originate from g and h . Then $R(r,s)$, and by the construction of $r(g)$ we find that either $a=\tau$ and $r \rightarrow^\tau r'$ was already an edge in g , or g has an edge $r \rightarrow^b r^*$ and $r \rightarrow^a r'$ is a copy of an initial edge from $r(b)$.

In the first case it follows from $R: g \rightleftharpoons_{rb} h$ that either $R(r',s)$ —hence $r(R)(r',s)$ —or in h there is a path $s \Rightarrow s_1 \rightarrow^\tau s'$ such that $R(r,s_1)$ and $R(r',s')$. By definition of refinement, the same path also exists in $r(h)$, and thus we have $r(R)(r,s_1)$ and $r(R)(r',s')$.

In the second case there must be a corresponding path $s \Rightarrow s_1 \rightarrow^b s^*$ in h such that $R(r,s_1)$ and $R(r',s^*)$. Then, in $r(h)$ we find a path $s \Rightarrow s_1 \rightarrow^a s'$ (by replacing $\rightarrow^b$ by $r(b)$) such that $r(R)(r,s_1)$ and $r(R)(r',s')$.

(2) The nodes r and s originate from related copies $\underline{r(b)}$ and $\overline{r(b)}$ of a substituted graph $r(b)$ (for some $b \in A$), and are no copies of root or endnodes in $r(b)$. Then $r \rightarrow^a r'$ is an edge in $\underline{r(b)}$. From $r(R)(r,s)$ we find that r and s are copies of the same node from $r(b)$. So, there is an edge $s \rightarrow^a s'$ in $\overline{r(b)}$ where s' is a copy of the node in $r(b)$, corresponding with r' . Clearly $r(R)(r',s')$.

- iii. Since for nodes from g and h we have $r(R)(r,s)$ iff $R(r,s)$, the root condition is satisfied. $\square$

With respect to closed BPA_τ -terms, the refinement theorem can be proved much easier by syntactic analysis of proofs, instead of working with equivalences between graphs. For observe that the Axioms A1-A5 + H1-H2, that form a complete axiomatization of branching congruence for closed terms, do *not* contain any occurrences of (atomic) actions from A . Now assume we have a proof of some equality $s=t$ between closed terms, then this proof consists of a sequence of applications of axioms from A1-A5 + H1-H2. Since all these axioms are universal equations without actions from A , the actions from s and t can be replaced by general variables, and the proof will still hold. Hence, every equation is an instance of a universal equation *without* any actions. Immediately we find that we can substitute arbitrary closed terms for these variables, obtaining refinement for closed terms.

Nevertheless, the semantic proof of the refinement theorem is important since it also holds for larger graphs from G_{BPA} that cannot be represented by closed BPA_τ -terms.

In the setting of BCCS, a substitution should be a mapping $r: A \rightarrow G_{CCS} \setminus \{0\}$, where 0 denotes the trivial graph. Then the semantic proof of the refinement theorem goes exactly as in the setting of BPA. However the syntactic proof breaks down on the absence of general sequential composition and on the presence of actions in the axioms

for branching congruence. In the setting of basic ACP, Definition 6.1 should be adapted such that r' is identified not with all endnodes of $r(a)$, but with all nodes of $r(a)$ that have an outgoing termination edge. These termination edges should then be deleted. Furthermore if certain parts in the resulting graph have become disconnected from the root, they should be deleted as well. Now both the semantic and the syntactic proof of the refinement theorem remain valid.

Delay bisimulation is also preserved under action refinement (as pointed out in DEVILLERS [72]), and η -bisimulation is not (same counterexample as for weak bisimulation). Moreover, delay bisimulation is the coarsest equivalence that is preserved under refinement and finer than τ -bisimulation (i.e. *fully abstract* with respect to weak bisimulation and action refinement) (CHERIEF & SCHNOEBELEEN [48]), but branching bisimulation is not quite the coarsest equivalence that is preserved under refinement and finer than η -bisimulation. This was established in CHERIEF [47], who characterized the *quasi-branching bisimulation equivalence* below as the coarsest equivalence that is preserved under refinement and finer than η -bisimulation.

DEFINITION 6.4 Two graphs g and h are *quasi-branching bisimilar*—notation: $g \approx_b h$ —if there exists a symmetric relation R (called a *quasi-branching bisimulation*) between the nodes of g and h such that:

- i. The roots are related by R
- ii. If $R(r, s)$ and $r \rightarrow^a r'$, then either $a = \tau$ and there is a path $s \Rightarrow s'$ such that $R(r', s')$, or there exists a path $s \Rightarrow s_1 \rightarrow^a s'$ such that $R(r, s_1)$ and $R(r', s')$.

It is easy to see that quasi-branching bisimulation is finer than both delay and η -bisimulation, but coarser than branching bisimulation equivalence, i.e. we have $g \approx_b h \Rightarrow g \approx_{qb} h \Rightarrow g \approx_\eta h$ and $g \approx_{qb} h \Rightarrow g \approx_d h$. The strictness of the first implication follows from the example $\tau(\tau+b)+a \approx_{qb} \tau(\tau+b)+a+\tau$, which does not apply to $\approx_b$. The strictness of the other two implications follows from the fact that neither $\approx_\eta$ nor $\approx_d$ implies the other. Cherief's quasi-branching bisimulation coincides with the *simple branching bisimulation* independently proposed by VAN BENTHEM, VAN EIJCK & STEBLETSOVA [21] and discussed in the Section 8.

Finally it should be noted that action refinement as defined in this section is a meaningful notion that can be used in the design of systems *only* if these system are assumed to be sequential (i.e. performing only one action at a time). In the presence of parallel composition, process graphs as presented here are not sufficiently expressive for defining a refinement operator. For this purpose one may better use causality based models of concurrency, such as event structures or Petri nets. This will be the topic of the following chapter.

7. DIVERGENCE

In the literature on bisimulation semantics roughly three ways are suggested for treating divergence (= infinite τ -paths). The original notion of τ -bisimulation equivalence (HENNESSY & MILNER [111], MILNER [134] and PARK [145]) abstracted from all divergencies; the first two graphs of Figure 11 are equivalent, as well as the two graphs of Figure 8.




Figure 11. Three ways of modeling divergence.

These identifications can be justified by an appeal to fairness (MILNER [134] and BAETEN, BERGSTRA & KLOP [13]), and play a crucial role in many protocol verifications. In BERGSTRA, KLOP & OLDEROG [29] the corresponding semantics is referred to as *bisimulation semantics with fair abstraction*. A variant where divergence is taken into account, in the sense that the first two graphs of Figure 11 are distinguished, as well as the two graphs of Figure 8, was proposed in HENNESSY & PLOTKIN [113] for (a variant of) τ -bisimulation and in MILNER [135] for delay bisimulation. In both cases a complete axiomatization is provided in WALKER [181]. In these semantics the basic notion is a *preorder* rather than an equivalence, and divergence is identified with underspecification. The induced equivalences identify the last two graphs of Figure 11, which are distinguished in τ -bisimulation semantics with fair abstraction. Hence the two notions are incomparable. A semantics that refines both notions was proposed in BERGSTRA, KLOP & OLDEROG [29] under the name *bisimulation semantics with explicit divergence*.

η -, delay and branching bisimulation as presented in this chapter are all based on the variant of τ -bisimulation with fair abstraction. However it is completely straightforward to generalize the τ -bisimulation preorder of HENNESSY & PLOTKIN [113] to a η -bisimulation preorder, and the delay bisimulation preorder of MILNER [135] to a branching bisimulation preorder. Also it is not difficult to define η -, delay and branching bisimulation with explicit divergence in the spirit of BERGSTRA, KLOP

& OLDEROG [29]. For branching bisimulation the definition can conveniently be given in terms of coloured traces.

DEFINITION 7 A node in a coloured graph is *divergent* if it is the starting point of an infinite path of which all nodes have the same colour. A colouring *preserves divergence* if no divergent node has the same colour as a non-divergent node. Two graphs g and h are (rooted) *branching bisimulation equivalent with explicit divergence* if there exists a (rooted) consistent divergence preserving colouring on g and h for which they have the same coloured trace set.

8. MODAL CHARACTERIZATIONS

It is well known (cf. HENNESSY & MILNER [112]) that strong and weak equivalence can be characterized by means of a simple modal language, called Hennessy-Milner logic (HML):

DEFINITION 8 The formulas of the *Hennessy-Milner logic (HML)* are given by the syntax

$$\varphi ::= \langle a \rangle \varphi \mid \varphi \wedge \psi \mid \neg \varphi \mid T$$

The (strong) satisfaction relation $\models$ between processes and formulas is defined by

- i. $g \models T$ for all process graphs g ,
- ii. $g \models \langle a \rangle \varphi$ if there is an edge $\text{root}(g) \rightarrow^a s$ and $(g)_s \models \varphi$,
- iii. $g \models \varphi \wedge \psi$ if $g \models \varphi$ and $g \models \psi$,
- iv. $g \models \neg \varphi$ if $g \not\models \varphi$.

We called the satisfaction relation above *strong* because it is tailored to strong bisimulation equivalence. Hennessy and Milner established that for finitely branching process graphs g and h

$$g \approx h \Leftrightarrow \forall \varphi \in \text{HML} (g \models \varphi \Leftrightarrow h \models \varphi).$$

In MILNER [137] this characterization was generalized to infinitely branching processes by replacing the finite conjunction $\wedge$ by an infinite conjunction $\bigwedge_{i \in I}$.

This result was adapted to weak bisimulation equivalence by replacing the modality $\langle \tau \rangle$ by $\langle \varepsilon \rangle$ and changing the satisfaction relation a bit. What we call the *weak* notion of satisfaction, $\models_\tau$, is given by

$$g \models_\tau \langle a \rangle \varphi \text{ if there is a path } \text{root}(g) \Rightarrow \rightarrow^a \Rightarrow s \text{ and } (g)_s \models_\tau \varphi$$

$$g \models_\tau \langle \varepsilon \rangle \varphi \text{ if there is a path } \text{root}(g) \Rightarrow s \text{ and } (g)_s \models_\tau \varphi$$

and the same clauses for T , $\wedge$ (or $\bigwedge$) and $\neg$ as we had before. Now one obtains:

$$g \approx_\tau h \Leftrightarrow \forall \varphi \in \text{HML} (g \models_\tau \varphi \Leftrightarrow h \models_\tau \varphi).$$

Equivalently, one can introduce a modality $\langle \sigma \rangle$ for any sequence σ of visible actions and put $g \models_\tau \langle \sigma \rangle \varphi$ if there is a path $\text{root}(g) \Rightarrow^\sigma s$ and $(g)_s \models_\tau \varphi$.

In fact, Hennessy and Milner considered a different kind of observational equivalence (Definition 1.8) than what is now known as weak or τ -bisimulation equivalence, and this version of observational equivalence is modally characterized just as weak bisimulation above, but without the modality $\langle \varepsilon \rangle$.

In MILNER [135] a modal characterization of delay bisimulation was given. Define the *delay* notion of satisfaction, $\models_d$, by

$$g \models_d \langle a \rangle \varphi \quad \text{if there is a path } \text{root}(g) \Rightarrow \rightarrow^a s \text{ and } (g)_s \models_d \varphi$$

$$g \models_d \langle \varepsilon \rangle \varphi \quad \text{if there is a path } \text{root}(g) \Rightarrow s \text{ and } (g)_s \models_d \varphi$$

and the same clauses for T , $\wedge$ (or Λ) and $\neg$ as before. Now one obtains:

$$g \equiv_d h \Leftrightarrow \forall \varphi \in \text{HML} (g \models_d \varphi \Leftrightarrow h \models_d \varphi).$$

Actually, this is not precisely Milner's characterization, as he simultaneously changed the treatment of divergence (see the previous section).

Now the question arises if a modal characterization can also be obtained for branching equivalence. A positive answer was given in DE NICOLA & VAANDRAGER [67], but it required the addition of a family of *until*-operators to HML. These operators are also denoted $\langle a \rangle$, but can be distinguished from the operators above because they are binary. $\varphi \langle a \rangle \psi$ says that, during a sequence of internal actions, formula φ remains valid until the action a happens, and afterwards ψ holds:

$$\begin{aligned} g \models_b \varphi \langle a \rangle \psi & \quad \text{if there is a path } \text{root}(g) = s_0 \rightarrow^\tau s_1 \rightarrow^\tau \dots \rightarrow^\tau s_n \rightarrow^a s \text{ (} n \geq 0 \text{)} \\ & \quad \text{such that } (g)_{s_i} \models_b \varphi \text{ for } i=0,1,\dots,n \text{ and } (g)_s \models_b \psi \\ g \models_b \varphi \langle \varepsilon \rangle \psi & \quad \text{if there is a path } \text{root}(g) = s_0 \rightarrow^\tau s_1 \rightarrow^\tau \dots \rightarrow^\tau s_n \text{ (} n \geq 0 \text{)} \text{ such} \\ & \quad \text{that } (g)_{s_i} \models_b \varphi \text{ for } i=0,1,\dots,n-1 \text{ and } (g)_{s_n} \models_b \psi. \end{aligned}$$

Now the unary modalities $\langle a \rangle \varphi$ can be regarded as abbreviations for $T \langle a \rangle \varphi$. In this sense is HML with until operators an extension of HML under the delay interpretation. They established that

$$g \equiv_b h \Leftrightarrow \forall \varphi \in \text{HML} + \text{'until'} (g \models_b \varphi \Leftrightarrow h \models_b \varphi).$$

In VAN GLABBEK [81] a variant of this characterization is proposed. There binary modalities a and τ , written without angular brackets, are considered, defined by

$$\begin{aligned} g \models_b \varphi a \psi & \quad \text{if there is a path } \text{root}(g) \Rightarrow^\tau r \rightarrow^a s \text{ such that } (g)_r \models_b \varphi \text{ and } (g)_s \models_b \psi \\ g \models_b \varphi \tau \psi & \quad \text{if there is a path } \text{root}(g) \Rightarrow^\tau r \text{ such that } r=s \text{ or } r \rightarrow^\tau s \\ & \quad \text{with } (g)_r \models_b \varphi \text{ and } (g)_s \models_b \psi. \end{aligned}$$

Also these modalities, together with T , Λ and $\neg$, characterize branching bisimulation equivalence.

In fact this modal language is equally expressive as the one of De Nicola & Vaandrager. Namely

$$g \models_b \varphi a \psi \Leftrightarrow g \models_b T \langle \varepsilon \rangle (\varphi \langle a \rangle \psi) \text{ and } g \models_b \varphi \tau \psi \Leftrightarrow g \models_b T \langle \varepsilon \rangle (\varphi \wedge (\varphi \langle \varepsilon \rangle \psi))$$

which establishes that it can be expressed in HML+'until', and, using the stuttering lemma,

$$g \models_b \varphi \langle a \rangle \psi \Leftrightarrow g \models_b \varphi \wedge (\varphi a \psi) \quad \text{and} \quad g \models_b \varphi \langle \varepsilon \rangle \psi \Leftrightarrow g \models_b (\varphi \wedge (\varphi \tau \psi)) \vee \psi$$

which establishes that it can express HML+'until'.

In VAN GLABBEEK [81] also a modal characterization of η -bisimulation is provided. Besides the HML modalities it has a binary modality a (for $a \neq \tau$), as in the characterization above, but this time the interpretation of the modalities is weak, rather than delay-like.

$g \models_{\eta} \phi a \psi$ if there is a path $\text{root}(g) \Rightarrow^{\tau} r \rightarrow^a \Rightarrow^{\tau} s$ with $(g)_r \models_{\eta} \phi$ and $(g)_s \models_{\eta} \psi$. Again $\langle a \rangle \phi$ can be regarded as an abbreviation for $\text{Ta}\phi$. However, any attempt to generalize the modality $\langle \varepsilon \rangle \phi$ to a binary version yields, under the weak interpretation, a modality that is no more expressive than $\langle \varepsilon \rangle \phi$. And indeed, HML together with the binary modalities a (for $a \neq \tau$), yields a modal characterization of η -bisimulation:

$$g \cong_{\eta} h \Leftrightarrow \forall \phi \in \text{HML}^{\eta} (g \models_{\eta} \phi \Leftrightarrow h \models_{\eta} \phi).$$

Considering the absence of a binary modality for silent activity in the η -characterization, it appears that one variant is still unexplored, namely HML with the binary modalities a (for $a \neq \tau$), or equivalently $\langle a \rangle$ (for $a \neq \tau$), without a binary modality for internal activity, but employing the delay (= branching) interpretation. This possibility is investigated in VAN BENTHEM, VAN EIJCK & STEBLETSOVA [21], and the equivalence characterized by these modalities is called *simple branching bisimulation equivalence*. It happens to be the same as the *quasi-branching bisimulation* independently proposed by CHERIEF [47] and discussed in the Section 6.

In MILNER [134,135] the name observation equivalence is justified by the presentation of a scenario in which two processes are observation inequivalent iff there is an experiment on which one reacts differently than the other. A successful experiment on a process p corresponds with an HML formula satisfied by p , and, vice versa, for each HML formula satisfied by p there is an experiment showing so. In order to upgrade this testing scenario to one that distinguishes branching inequivalent processes by experiment, it suffices to find experiments corresponding to the 'until'-operators added to HML. For this purpose one may assume that a process leaves a continuous chain of core-dumps when it proceeds. Here each core-dump contains all information needed to reconstruct the process in the state were its core was dumped, and 'continuous' means that in every state at least one, but possibly many core-dumps are left. By examining all core-dumps of a process prior to the execution of an action a , it is possible to ascertain that a formula ϕ holds until a occurs. Similarly one can check that during a series of internal actions ϕ holds until ψ holds. More details can be found in VAN GLABBEEK [81].

Besides HML with until operators several other modal characterizations of branching bisimulation have been proposed. First of all each of the logics (A)CTL and (A)CTL* without nexttime operator, characterize (a variant of) branching bisimulation (DE NICOLA & VAANDRAGER [67,68]). See Argument 5 in the conclusion. These logics are more expressive than HML+'until'.

In DE NICOLA, MONTANARI & VAANDRAGER [66] it has been established that if in the definition of *-bisimulation, for $* \in \{\tau, b, \eta, d\}$, it is required that moves in the one process can be simulated by the other process, not only when going forward but also when going back in history, these modified notions all coincide with branching

bisimulation. This yields another modal characterization of branching bisimulation, namely HML with backward modalities.

A last possibility may be adding the eventually operator to HML. It remains to be determined for which classes of process graphs HML + 'eventually' is adequate.

CONCLUSION

In this chapter we introduced a new semantic equivalence for concurrent systems that we called *branching bisimulation equivalence*. We compared branching bisimulation with the coarser notion of weak bisimulation equivalence and two intermediate notions. Our main motivation for introducing branching bisimulation is that it preserves the branching structure of processes. Although we believe that this has been demonstrated in Sections 1 and 3, this chapter does not contain a formal definition of the 'branching structure' of a process. Such a definition will be offered in VAN GLABBEEK [82].

Here we would like to stress that this feature of branching bisimulation is of more than philosophical interest. Actually, we think that of all the equivalences in the linear time - branching time spectrum, branching bisimulation is most suited for verifying the correctness of concurrent systems in applications where the final word on what exactly is observable behaviour has not been pronounced. Below we list our arguments.

1. In applications where the notion of observable behaviour is clear, the most suitable equivalence is usually the one which is *fully abstract* with respect to this notion of observable behaviour, i.e. identifies two processes if and only if their observable behaviour is the same. However, if there is no clarity on what is observable, a verification in a fully abstract semantics w.r.t. any notion of observability needs to be redone every time one discovers that a little bit more can be observed than what was originally accounted for. Moreover, the soundness of the verification depends crucially on the right estimation of what can be observed. A verification (of the equivalence of two processes) in a semantics that preserves the internal structure of processes, on the other hand, does not depend on considerations of observability (as long as it is clear that no more can be observable than this internal structure), and is automatically valid in any semantics that is fully abstract w.r.t. some notion of observable behaviour. Now for the simple kind of processes that are the subject of this chapter, the best formalization of the internal structure of a process appears to be its branching structure. (However, for more complex systems the internal structure may involve more, e.g. the *causal structure* of processes.)

This argument is by no means new, and also appeared in the original version of this text. But the elaboration paraphrased above stems from VAN GLABBEEK [82].

2. The above argument says that branching bisimulation equivalence (or congruence) is suited for applications in which there is no certainty on what constitutes observable behaviour. However, it does not show that it is the only such notion. It could be that there exists a coarser equivalence that still preserves the upperbound on observable behaviour. In fact the name *observation equivalence* suggests that weak (or delay?) bisimulation equivalence is the coarsest equivalence with this property. This argument is enforced by Milner's testing scenario for

observation equivalence, presented in [134] and [135]. However, the testing scenario briefly sketched Section 8 and elaborated in VAN GLABBEEK [81] (together with argument 1 above) shows that in fact branching bisimulation represents the limit of observable behaviour, and hence no coarser equivalence shares the advantage mentioned in argument 1.

Of course one could argue that the testing scenario of Section 8 is not very realistic. However, the same can be said of the testing scenario for weak bisimulation, and the question of what is a realistic testing scenario is exactly the one that we want to avoid.

For those readers that believe that some variant of *ready simulation equivalence* represents the limit of observable behaviour and is therefore the right notion (cf. BLOOM, ISTRAIL & MEYER [35] and ULIDOWSKI [166]) we refer to the *liveness* argument in VAN GLABBEEK [81].

3. This leaves us with the question whether a finer equivalence than branching bisimulation might be more or equally suitable. In the absence of silent moves the only candidate appears to be *tree equivalence* (VAN GLABBEEK [82]), as even finer equivalences, such as graph isomorphism, are clearly useless from the point of view of practical applications. But tree semantics has the disadvantage that the standard operational and denotational interpretations of CCS-like system description languages do not coincide. Moreover the operational interpretation is not compositional, and the most plausible fix requires an upgrade of the underlying graph model into a multigraph model (allowing more than one equally labelled edges between two nodes).

Thus it is more tempting to search for generalizations of bisimulation equivalence to a setting with silent moves that are finer than branching bisimulation. Such generalizations undoubtedly exists. But in many application we are interested in three useful properties:

- i) The equivalence is abstract in the sense that it satisfies at least $\alpha\tau x = \alpha x$ (Milner's first τ -law). This criterion was mentioned in the historical note in Section 1.
- ii) It is a congruence for the operators of CCS and CSP.
- iii) The merge or parallel composition satisfies the expansion theorem of MILNER [134,135,137,138], i.e. interleaving semantics is employed, and also the other laws of strong bisimulation are satisfied.

In such circumstances branching bisimulation on finite closed terms is completely axiomatized by the first τ -law ($\alpha\tau x = \alpha x$) and the laws of strong bisimulation, and hence is the finest congruence possible (VAN GLABBEEK [80]).

Returning to tree equivalence, a similar argument as above applies: all recursion-free closed instances of the law $a(\tau(y+z)+y) = a(y+z)$ are derivable from $\alpha\tau x = \alpha x$ and the laws for (strong) tree equivalence (which are the laws of strong bisimulation without $x+x=x$). Thus, in applications where properties i-iii above are desired, a 'weak' version of tree congruence would be needed, that satisfies $a(\tau(y+z)+y) = a(y+z)$. Such a weak tree-congruence would not have the intuitive appeal of strong tree equivalence, and does not appear to have any advantages over branching bisimulation congruence.

4. The crucial difference between branching bisimulation and weak bisimulation is that the first one better takes into account the intermediate states of two equivalent processes as they progress through a computation. As argued already by HENNESSY & MILNER [111] on the occasion of the introduction of the first version of observation equivalence, it is useful to do so, "because different intermediate states can be exploited in different program context to provide different overall behaviour". Here we present an example of a program context that exploits the different intermediate states of two observation equivalent (i.e. weakly bisimilar) processes to provide different overall behaviour. It is a context (operator) that allows a process (its argument) to proceed normally, but in addition can report that the process is ready to perform a visible action, without actually doing it. Thus the states of this context are the same as the states of its argument, and the transitions are the transitions of its argument, together with a transition labelled "can do 'a'" form a state to itself, whenever the argument can do an 'a' from that state. Observe that the two processes $a+\tau b$ and $a+\tau b+b$ displayed in Figure 12 are τ -bisimilar.




Figure 12. An operator exploiting different intermediate states of weakly bisimilar processes.

They are even delay bisimilar, but not branching bisimilar. After placing them in the described context (the results are displayed below them) they are no longer τ -bisimilar. In fact they even have different traces, as only the second one has a trace "can do 'b'" followed by "a".

In BLOOM [34] a class of operators is given for which rooted branching bisimulation equivalence is a congruence. He also presents a class of operators for which rooted weak bisimulation equivalence is a congruence. The latter class is a subclass of the former, and the operator described above falls in the difference. Also all equivalences in the linear time - branching time spectrum that are situated between weak trace equivalence and delay bisimulation fail to be congruences for this operator. On the other hand, we know of no useful operator for which some abstract equivalence in the linear time - branching time spectrum is a congruence, but rooted branching bisimulation is not.

5. As mentioned in the introduction, τ -bisimulation semantics is not adequate for a modal logic with eventually operator. From the examples in the introduction one can see that the problem originates from the circumstance that τ -bisimulation equivalence does not preserve the branching structure of processes. They also apply to all other semantic equivalences that do not preserve branching, and indeed one can easily prove that such an operator would cause no problems in branching bisimulation semantics, at least not in the variant with explicit divergence. In fact, a much stronger result has been proved in DE NICOLA & VAANDRAGER [67].

The Computation Tree Logic CTL* (EMERSON & HALPERN [74]) is a very powerful logic, combining both branching time and linear time operators. It is a generalization of CTL (CLARKE & EMERSON [50]), that contains only branching time operators. CTL* is interpreted on Kripke structures (directed graphs of which the nodes are labelled with sets of atomic propositions). DE NICOLA & VAANDRAGER [67] established a translation from process graphs to Kripke structures, so that CTL* can also be regarded as a logic on process graphs. In fact, in DE NICOLA & VAANDRAGER [68] they introduced a counterpart ACTL* of CTL* on process graphs and supplied translations in both directions. One of the operators of (A)CTL/CTL*, the nexttime operator X, makes it possible to see when an (invisible) action takes place, and is therefore incompatible with abstraction. This operator was also criticized by LAMPORT [124]. BROWNE, CLARKE & GRÜMBERG [43] found that CTL without X and CTL* without X induce the same equivalence on Kripke structures, which they characterized as *stuttering equivalence*. In DE NICOLA & VAANDRAGER [67] branching bisimulation, after being translated to Kripke structures, is shown to coincide with stuttering equivalence. (To be precise, they consider two variants of CTL*, that correspond to two variants of stuttering equivalence and two variants of branching bisimulation, namely *divergence blind branching bisimulation* (our notion with fair abstraction) and *divergence sensitive branching bisimulation* (defined as branching bisimulation with explicit divergence as in Section 7, but also considering endnodes to be divergent). The stuttering equivalence of BROWNE, CLARKE & GRÜMBERG [43] is the divergence sensitive variant.) Hence (divergence sensitive) branching bisimulation is adequate for (A)CTL*-X. Since the eventually operator of GRAF & SIFAKIS [102] can be expressed in (A)CTL*, this implies that it causes no problems in branching bisimulation semantics.

6. The extra identifications made in τ -bisimulation semantics on top of branching bisimulation semantics can be cumbersome in certain applications of the theory. An example concerns the work of JONSSON & PARROW [119], mentioned in the introduction. On the other hand we are not aware of a single application where τ -bisimulation semantics can be successfully applied, but the extra distinctions made in branching bisimulation semantics pose a problem.

In BOUALI [36] an example is given where τ -bisimulation semantics works better than branching bisimulation semantics. The example concerns the minimization of PETERSON's mutual exclusion algorithm [147]. Here *minimization* means finding an equivalent process that is as small as possible. In branching bisimulation semantics this yields a process with 17 states, whereas in weak bisimulation semantics a process with only 14 states is obtained. (In fact, using the quasi-branching bisimulation mentioned in Sections 6 and 8 would bring the number of states to 14 already.) It should be noted however, that weak bisimulation is still far from optimal for this purpose. *Coupled simulation*, proposed by PARROW & SJÖDIN [146], is a generalization of bisimulation semantics to a setting with silent moves that is coarser than weak bisimulation. It is completely axiomatized by the laws of weak bisimulation together with $\tau(\tau x + y) = \tau x + y$. Minimization of Peterson's algorithm in coupled simulation semantics would yield a process with no more than 9 states, and using failure semantics (BROOKES, HOARE & ROSCOE [42], DE NICOLA & HENNESSY [64]) would bring it down to 4.

We conjecture that this is illustrative for a general tendency. Coupled simulation has distinct advantages over weak (and branching) bisimulation in applications where the latter notions are too fine. Examples of such applications can be found in PARROW & SJÖDIN [146] and Section 4.5 of Chapter II. However, whenever weak bisimulation performs better than branching bisimulation, it turns out to be the case that neither of the two notions are really suitable, and coupled simulation, or an even coarser equivalence, is called for.

7. No abstract semantic equivalence used in concurrency theory is as easy to decide as branching bisimulation congruence. For context-free process without silent actions, i.e. processes that can be specified with recursive equations over the language BPA of Section 2.1, bisimulation equivalence is the only equivalence in the linear time - branching time spectrum of Chapter I that is decidable at all (cf. GROOTE & HÜTTEL [105] and CHRISTENSEN, HIRSHFELD & STIRLING [49]). For finite-state processes, i.e. processes that can be specified with recursive equations over the language BCCS of Section 2.1, bisimulation equivalence can be decided in polynomial time, whereas most other equivalences in the linear time - branching time spectrum are PSPACE-complete (cf. KANELAKIS & SMOLKA [120]).

In GROOTE & VAANDRAGER [106] an algorithm is presented for deciding branching bisimulation equivalence between finite-state processes, with (time) complexity $O(l+n \cdot m)$. Here l is the size of Act, n is the number of nodes in the investigated process graphs and m the number of edges. The fastest algorithm for τ -bisimulation equivalence up till then had time complexity $O(l \cdot n^{2.376})$. In general $n \leq m \leq l \cdot n^2$, so it depends on the density of edges in a graph which algorithm is

faster. In a trial implementation of the scheduler of MILNER [134], reported in GROOTE & VAANDRAGER [106], branching bisimulation turned out to be much faster. Furthermore, it turned out that in such automatic verifications the space complexity was a much more serious handicap than the time complexity (the τ -bisimulation tools suffered from lack of memory already when applied to processes with 15.000 states). The space complexity of the algorithm of GROOTE & VAANDRAGER [106] is $O(n+m)$, which is less than the space complexity of the τ -bisimulation algorithm. Recently, BOUALI [36] proposed another algorithm for τ -bisimulation that has the same time complexity as the branching bisimulation algorithm. An trial with Milner's scheduler shows that it is somewhat slower, but with a constant factor only. The space complexity of BOUALI's algorithm is $O(n+m^+)$, where m^+ is the number of edges after taking the τ^+ transitive closure, obtained by adding an edge $p \rightarrow^\tau q$ whenever there is a nonempty path $p \Rightarrow q$. This improves the space complexity of the old algorithm for τ -bisimulation, but is not quite as good as the branching bisimulation algorithm.

8. For sequential processes, branching bisimulation is preserved under refinement of actions, whereas τ -bisimulation is not. This was established in Section 6, which appeared before as VAN GLABBEK & WEIJLAND [95]. A proof can also be found in DARONDEAU & DEGANI [56].
9. All $*$ -bisimulations ($*$ $\in \{\tau, b, \eta, d\}$) have relatively simple equational characterizations (see Section 2). However, the listed axioms are in no way self-evident, but arise from the semantic presentation of the respective notions. In the presence of the (plausible) axioms of strong equivalence, branching bisimulation congruence on finite closed terms is completely axiomatized by the (equally plausible) first τ -law ($atx=ax$) alone (VAN GLABBEK [80]). This means that the algebra of branching bisimulation can be understood without appeal to semantic notions at all.
10. The axiom system for branching bisimulation can easily be turned into a complete term rewriting system, which is not the case for the other abstract bisimulation semantics. Work in this direction has been done in AKKERMAN & BAETEN [8] (in the framework of ACP) and DE NICOLA, INVERARDI & NESI [65] (in the framework of CCS).
11. Branching bisimulation equivalence has a nice characterization as weak *back-and-forth bisimulation* (see the end of Section 8). No matter whether the weak, delay, η - or branching mode is selected, if it is required that moves in the one process can be simulated by the other process, not only when going forward but also when going back in history, these modified notions all coincide with branching bisimulation (DE NICOLA, MONTANARI & VAANDRAGER [66]).

This argument is balanced however by a nice characterization of weak bisimulation of which there is no analogy for branching bisimulation. Namely if (root-unwound) process graphs are saturated by adding an edge $s \rightarrow^a s'$ whenever there is a path $s \Rightarrow \rightarrow^a \Rightarrow s'$ (cf. Definition 4.8) and by adding a τ -loop in every state except the root, then two graphs are weakly equivalent iff their saturated graphs are strongly equivalent.

Chapter IV

Refinement of Actions in Causality Based Models

Rob van Glabbeek & Ursula Goltz

In this chapter we consider an operator for refinement of actions to be used in the design of concurrent systems. Actions on a given level of abstraction are replaced by more complicated processes on a lower level. This is done in such a way that the behaviour of the refined system may be inferred compositionally from the behaviour of the original system and from the behaviour of the processes substituted for actions. We define this refinement operation for causality based models like event structures and Petri nets. For Petri nets, we relate it to other approaches for refining transitions.

Contents


Introduction	173
1. Refinement of actions in prime event structures	181
2. Refinement of actions in flow event structures	186
3. Configuration structures and refinement of actions	193
4. Refinement of transitions in Petri nets	197
Related work	211

Introduction


In this chapter we consider the design of concurrent systems in the framework of approaches where the basic building blocks are the actions which may occur in a system. By an action we understand here any activity which is considered as a conceptual entity on a chosen level of abstraction. This allows to design systems in a top-down style, changing the level of abstraction by interpreting actions on a higher level by more complicated processes on a lower level. We refer to such a step in the design of a system as *refinement of actions*. An action could be refined by the sequential execution of several sub-actions, or by activities happening independently in parallel. One could also implement an action by a set of alternatives, of which only one should be taken.

0.1 Example

Consider the design of a sender, repeatedly reading data and sending them to a certain receiver. A first description of this system is given by the Petri net shown below. An introduction to Petri nets and the way they model concurrent systems can be found in REISIG [156]; the refinement mechanism used in this example will be treated formally in Section 4.




On a slightly less abstract description level the action “send data to receiver” might turn out to consist of two parts “prepare sending” and “carry out sending”, to be executed sequentially. This corresponds to the following refined Petri net.




Refinement by a sequential process

Then the action “prepare sending” may be decomposed in two independent activities “prepare data for transmission” and “get permission to send”, to be executed on different processors:




Refinement by a parallel process

Furthermore it may turn out that there are two alternative channels for sending messages. Each time the sender should choose one of them to send a message, perhaps depending on which one is available at the moment.



Refinement by alternative actions

On an even more concrete level of abstraction, channel 2 may happen to be rather unreliable, and getting a message at the other end requires the use of a communication protocol. On the other hand, channel 1 may be found to be reliable, and does not need such a precaution.



Refinement by an infinite process

Here we see that it may happen that the process we have substituted for the action “send on channel 2” does not terminate. It may happen that the attempt of sending data always fails and this prevents the system of reaching its initial state again.


Our aim is to define an operator for refinement of actions, taking as arguments a system description on a given level of abstraction and an interpretation of (some of) the actions on this level by more complicated processes on a lower level, and yielding a system description on the lower level. This should be done in such a way that the behaviour of the refined system may be inferred compositionally from the behaviour of the original system and from the behaviour of the processes substituted for actions.

As illustrated above, we want to allow to substitute rather general kinds of behaviours for actions. We even allow the refinement of an action by an infinite behaviour. This contradicts a common assumption that an action takes only a finite amount of time. It means that when regarding a sequential composition $a;b$ we can not be sure that b occurs under all circumstances; it can only occur if the action a really terminates successfully.

There is one type of refinement that we do not want to allow, namely to “forget” actions by replacing them with the empty process.

0.2 Example

Continuing Example 0.1 we could imagine that getting permission to send turns out to be unnecessary and can be skipped. Hence we replace the corresponding action by the empty behaviour, thus obtaining




Forgetful refinement


Even though this operation seems natural when applied as in the above example, it may cause drastic changes in the possible behaviours of a system. It may happen that executing a certain action a prevents another action from happening. This property should be preserved under refinement of a . However, if a is completely removed, it cannot prevent anything any more, which can remove a deadlock possibility from the system. Thus “forgetful” refinements can not be explained by a change in the level of abstraction at which systems are regarded. For this reason they will not be considered here.

0.3 Example

Consider the Petri net



and the net obtained when refining a by the empty behaviour:




In the first net it is possible to execute a and b , and by this reach a state where no further action is possible. If we try to deduce the behaviour after refinement from the behaviour of N , we would expect that the refined system may reach a state, by executing b , where no more action is possible. However, this is not the case for N' . After b , it is always possible to execute c in N' .

In order to define a suitable refinement operator, one first has to select a model for the description of concurrent systems. The models of concurrency found in the literature can roughly be distinguished in two kinds: those in which the independent execution of two processes is modelled by specifying the possible interleavings of their (atomic) actions, and those in which the causal relations between the actions of a system are represented explicitly. The interleaving based models were devised to describe systems built from actions that are assumed to be instantaneous or indivisible. Nevertheless, one might be tempted to use them also for the description of systems built from actions that may have a duration or structure. However, the following example shows that it is not possible to define the desired compositional refinement operator on such models of concurrency without imposing some restrictions (as already observed in PRATT [154] and CASTELLANO, DE MICHELIS & POMELLO [46]).

0.4 Example


The systems $P = a \parallel b$, executing the actions a and b independently, and $Q = a;b + b;a$, executing either the sequence ab or the sequence ba , cannot be distinguished in interleaving models; they are represented by the same tree in the model of synchronisation trees (MILNER [134]).



After refining a into the sequential composition of a_1 and a_2 , thereby obtaining the systems

$$P' = (a_1; a_2) \parallel b \quad \text{and} \quad Q' = (a_1; a_2); b + b; (a_1; a_2),$$

their tree representations are different:




The two systems are even non-equivalent, according to any reasonable semantic equivalence, since only P' can perform the sequence of actions a_1ba_2 . Hence, in the model of synchronisation trees the semantic representation of the refined systems is not derivable from the semantic representation of the original systems. The same holds for other interleaving models.

There are still ways left to define a compositional refinement operator on interleaving based models. First of all one could restrict the kind of refinements that are allowed in such a way that situations as in Example 0.4 cannot occur. Of course this would exclude the possibility of refining a in $a_1; a_2$ in either P or Q (or both). This idea is investigated in CZAJA, VAN GLABBEK & GOLTZ [51]. Although we consider this to be an interesting option, in this thesis we choose to allow rather general refinements, including at least the one of Example 0.4. Furthermore, some approaches have been proposed which are based on a concept of “atomic” actions; refining an atomic action would then result in an “atomic” process that cannot be “interrupted” by other activities (the refinement of P in Example 0.4 would not have the execution a_1ba_2). We will comment on these approaches in the concluding section. In this work we choose


not to assume action atomicity in any way, and to allow the parallel or independent execution of actions. Hence interleaving based models are unsuited for our approach. On the other hand we will show that the desired compositional refinement operator *can* be defined on causality based models of concurrency without imposing such restrictions. We will do this for semantic models like Petri nets and event structures. Since these models are being used as a semantics of languages like CCS, we hope that this will lead also to extending these languages by a mechanism for refinement.

0.5 Example

The systems $P = a \parallel b$ and $Q = a;b + b;a$ from Example 0.4 may be represented by the (labelled) Petri nets



The Petri net representations of the refined systems P' and Q' , where a is replaced by the sequence a_1a_2 , are then derivable by transition refinement from the nets for the original systems. We obtain



We will use two kinds of semantic models. Both of them are based on the idea of PETRI [149] to model causalities in concurrent systems explicitly and thereby also representing independence of activities. Additionally, the models we use represent the choice structure of systems; they show where decisions between alternative behaviours are taken.

We will not distinguish external and internal actions here; we do not consider abstraction by hiding of actions.

The more basic model, in particular when being concerned more with actions than with states, are *event structures*. We will consider three types of event structures here: *prime event structures* with a binary conflict relation [142], *flow event structures*, which are particularly suited as a semantic model of CCS [40], and, as a more abstract and general model, *configuration structures* (*families of configurations* [183]), where a system is represented by its subsets of events which determine possible executions.

The models considered so far are usually not applied to model systems directly, but rather as the underlying semantics of system description languages like CCS. One of the reasons for this is that infinite behaviours can only be represented by infinite structures (with an infinite set of events). So, finally, we will consider Petri nets as a framework which is directly applicable in the design process. Event structures may be derived from Petri nets as a particularly simple case, but Petri nets are more powerful. For example, infinite behaviours may be represented as finite net structures together with the “token game”. However, causality is then no longer a basic notion but has to be derived. Petri nets with their appealing graphical representation are being used extensively for the — more or less formal — representation of systems and — mostly less formal — during the design process. A disciplined way for developing net models systematically by refinement is therefore very important.

We start in Section 1 by recalling the basic notions for prime event structures and by showing how to refine actions by finite, conflict-free behaviours. We show that, for refining actions with more general behaviours, it is convenient to use more expressive models. In Section 2, we introduce flow event structures and show how to refine actions also by (possibly infinite) behaviours with conflicts. We show that, as for prime event structure refinement, the behaviour of a refined flow event structure may be deduced compositionally. In Section 3, we introduce configuration structures and a refinement operation for them. We show that the more “syntactic” constructions in the previous sections are consistent with this general notion. Finally, we give an overview on the work on refinement in Petri nets, and we suggest a rather general notion of refinement of transitions which is still modular with respect to behaviour. Related work is discussed in the concluding section.

1 Refinement of actions in prime event structures

In this section we show how to refine actions in the most simple form of event structures, prime event structures with a binary conflict relation (NIELSEN, PLOTKIN & WINSKEL [142]). Furthermore we motivate our move to more general structures in the next two sections, because of the limitations of this approach.

We consider systems that are capable of performing actions from a given set *Act* of action names.

1.1 Definition

A (labelled) prime event structure (over an alphabet *Act*) is a 4-tuple $\mathcal{E} = (E, \leq, \#, l)$ where

- E is a set of events,
- $\leq \subseteq E \times E$ is a partial order (the *causality relation*) satisfying the *principle of finite causes*:

$$\forall e \in E : \{d \in E \mid d \leq e\} \text{ is finite,}$$

- $\# \subseteq E \times E$ is an irreflexive, symmetric relation (the *conflict relation*) satisfying the *principle of conflict heredity*:

$$\forall d, e, f \in E : d \leq e \wedge d \# f \Rightarrow e \# f,$$

- $l : E \rightarrow \text{Act}$ is a labelling function.

The components of a prime event structure $\mathcal{E}$ will be denoted by $E_{\mathcal{E}}, \leq_{\mathcal{E}}, \#_{\mathcal{E}}$ and $l_{\mathcal{E}}$. If clear from the context, the index $\mathcal{E}$ will be omitted. As usual, we write $d < e$ for $d \leq e \wedge d \neq e$, etc.

A prime event structure represents a concurrent system in the following way: action names $a \in \text{Act}$ represent actions the system might perform, an event $e \in E$ labelled with a represents an occurrence of a during a possible run of the system, $d < e$ means that d is a prerequisite for e and $d \# e$ means that d and e cannot happen both in the same run.

Causal independence (*concurrency*) of events is expressed by the derived relation $co \subseteq E \times E : d \text{ co } e \text{ iff } \neg(d < e \vee e < d \vee d \# e)$. By definition, $<, >, \#$ and co form a partition of $E \times E$.

Throughout this work we assume a fixed set *Act* of action names as labelling set. Let $\mathbb{E}_{\text{prime}}$ denote the domain of prime event structures labelled over *Act*.

A prime event structure $\mathcal{E}$ is *finite* if $E_{\mathcal{E}}$ is finite; $\mathcal{E}$ is *conflict-free* if $\#_{\mathcal{E}} = \emptyset$. O denotes the empty event structure $(\emptyset, \emptyset, \emptyset, \emptyset)$.

For $X \subseteq E_{\mathcal{E}}$, the *restriction of $\mathcal{E}$ to X* is defined as

$$\mathcal{E} \upharpoonright X = (X, \leq \cap (X \times X), \# \cap (X \times X), l \upharpoonright X).$$

Two prime event structures $\mathcal{E}$ and $\mathcal{F}$ are *isomorphic* ($\mathcal{E} \cong \mathcal{F}$) iff there exists a bijection between their sets of events preserving $\leq$, $\#$ and labelling. Generally, we will not distinguish isomorphic event structures.

Isomorphism classes of conflict-free event structures are called *pomsets* (PRATT [154]). They have been considered under the name *partial words* in GRABOWSKI [101]. Pomsets generated by certain subsets of events may be considered as possible “executions” of the system represented by the event structure. The partial order between action occurrences then represents causal dependencies in the execution. Subsets of events representing executions (called *configurations*) have to be conflict-free; furthermore they must be left-closed with respect to $\leq$ (all prerequisites for any event occurring in the “execution” must also occur). It is assumed that in a finite period only finitely many actions are performed. We will consider only finite executions when describing the behaviour of systems. So, unlike WINSKEL [183], we require configurations to be finite. We will comment on this point in Section 3.


1.2 Definition

- i. A subset $X \subseteq E$ of events in a prime event structure $\mathcal{E}$ is *left-closed* in $\mathcal{E}$ iff, for all $d, e \in E$, $e \in X \wedge d \leq e \Rightarrow d \in X$.
 X is *conflict-free* in $\mathcal{E}$ iff $\mathcal{E} \upharpoonright X$ is conflict-free.
- ii. A subset $X \subseteq E$ will be called a (*finite*) *configuration* of a prime event structure $\mathcal{E}$ iff X is finite, left-closed and conflict-free in $\mathcal{E}$. $\text{Conf}(\mathcal{E})$ denotes the set of all configurations of $\mathcal{E}$. A configuration $X \in \text{Conf}(\mathcal{E})$ is called *complete* iff $\forall d \in E : d \notin X \Rightarrow \exists e \in X \text{ with } d \# e$.

Configurations may be considered as possible states of the system; they determine the remaining behaviour of the system as being the set of all events which have not yet occurred and are not excluded because of conflicts. Note that a configuration X is complete iff it is maximal, i.e. $X \subseteq Y \in \text{Conf}(\mathcal{E})$ implies $X = Y$.

1.3 Example

The system $a \parallel b + a; b$, executing either a and b independently or a and b sequentially, may be represented by the prime event structure



where the causality relation is represented by arcs.

The configurations of $\mathcal{E}$ are

$$\emptyset, \{e_1\}, \{e_2\}, \{e_3\}, \{e_1, e_3\}, \{e_2, e_4\},$$

corresponding to the pomsets

$$\emptyset, a, b, \begin{smallmatrix} a \\ b \end{smallmatrix} \text{ and } a \longrightarrow b.$$

$\begin{smallmatrix} a \\ b \end{smallmatrix}$ and $a \longrightarrow b$ correspond to complete configurations.


In forthcoming graphical representations of prime event structures, only immediate conflicts—not the inherited conflicts—are indicated. The $\leq$ -relation is represented by arcs, omitting those derivable by reflexivity and transitivity. Furthermore, instead of events only their labels are displayed; if a label occurs twice it represents two different events. Thus these pictures determine event structures only up to isomorphism.

We will now define a refinement operation substituting actions by non-empty, finite, conflict-free event structures. As discussed in the introduction, we will not allow forgetful refinements replacing actions by the empty event structure. We will later explain why we have to restrict to finite and conflict-free refinements of actions.

A refinement function will be a function ref specifying, for each action a , an event structure $ref(a)$ which is to be substituted for a . Interesting refinements (and also the refinements in our examples) will mostly refine only certain actions, hence replace most actions by themselves. However, for uniformity (and for simplicity in proofs) we consider all actions to be refined.

Given an event structure $\mathcal{E}$ and a refinement function ref , we construct the refined event structure $ref(\mathcal{E})$ as follows. Each event e labelled by a is replaced by a disjoint copy, $\mathcal{E}_e$, of $ref(a)$. The causality and conflict structure is inherited from $\mathcal{E}$: every event which was causally before e will be causally before all events of $\mathcal{E}_e$, all events which causally followed e will causally follow all the events of $\mathcal{E}_e$, and all events in conflict with e will be in conflict with all the events of $\mathcal{E}_e$.

Graphically, the idea may be sketched as follows



1.4 Definition

- (i) A function $ref: Act \rightarrow \mathbb{E}_{prime} - \{O\}$ is called a *refinement function (for prime event structures)* if $\forall a \in Act : ref(a)$ is finite and conflict-free.
- (ii) Let $\mathcal{E} \in \mathbb{E}_{prime}$ and let ref be a refinement function.
Then $ref(\mathcal{E})$ is the prime event structure defined by
 - $E_{ref(\mathcal{E})} = \{(e, e') | e \in E_{\mathcal{E}}, e' \in E_{ref(l_{\mathcal{E}}(e))}\}$,
 - $(d, d') \leq_{ref(\mathcal{E})} (e, e')$ iff $d <_{\mathcal{E}} e$ or $(d = e \wedge d' \leq_{ref(l_{\mathcal{E}}(d))} e')$,
 - $(d, d') \#_{ref(\mathcal{E})} (e, e')$ iff $d \#_{\mathcal{E}} e$,
 - $l_{ref(\mathcal{E})}(e, e') = l_{ref(l_{\mathcal{E}}(e))}(e')$.

We show that refinement is a well-defined operation on prime event structures, even when isomorphic prime event structures are identified.


1.5 Proposition

- (i) If $\mathcal{E} \in \mathbb{E}_{prime}$ and ref is a refinement function then $ref(\mathcal{E})$ is a prime event structure indeed.
- (ii) If $\mathcal{E} \in \mathbb{E}_{prime}$ and ref, ref' are refinement functions with $ref(a) \cong ref'(a)$ for all $a \in Act$ then $ref(\mathcal{E}) \cong ref'(\mathcal{E})$.
- (iii) If $\mathcal{E}, \mathcal{F} \in \mathbb{E}_{prime}$ and ref is a refinement function then $ref(\mathcal{E}) \cong ref(\mathcal{F})$.

Proof Straightforward. ■

1.6 Example

We consider a simplified version of the sender (Example 0.1) from the introduction. We assume that the sender reads and sends only once. We may carry out the first two steps of the design in terms of prime event structures as follows.



The next refinement step would require a refinement of an action by conflicting behaviours. This is not possible in our framework up to now.

The reason that we can only refine actions by conflict-free event structures is the axiom of conflict heredity and the notion of configuration in prime event structures. They imply that any event will always occur with a unique history (in terms of its causal predecessors) [183].

Now consider e.g. $\mathcal{E} = \downarrow_b^a$. Replacing a by $c\#d$ would require to duplicate the event labelled by b in some way, since b should then occur either caused by c or by d . Since this would lead to a complicated definition, we will consider more general forms of event structures that do not require duplication in Section 2 and 3.

The restriction to refinement of actions by finite event structures is necessary to ensure that the resulting event structure will obey the axiom of finite causes. In

the more general models we will consider later, we will not assume this axiom, and this will allow also refinements by infinite behaviours as discussed in the introduction.

Finally, we show how the behaviour of the refined event structure $\text{ref}(\mathcal{E})$ is determined by the behaviour of $\mathcal{E}$ and by the behaviour of the event structures which are substituted for actions.

1.7 Proposition

Let $\mathcal{E} \in \mathcal{E}_{\text{prime}}$, let ref be a refinement function.

We call $\tilde{X}$ a *refinement of configuration* $X \in \text{Conf}(\mathcal{E})$ by ref iff

- $\tilde{X} = \bigcup_{e \in X} \{e\} \times X_e$ where $\forall e \in X : X_e \in \text{Conf}(\text{ref}(l_{\mathcal{E}}(e))) - \{\emptyset\}$,
- $e \in \text{busy}(\tilde{X}) \implies e$ maximal in X with respect to $\leq_{\mathcal{E}}$
where $\text{busy}(\tilde{X}) := \{e \in X \mid X_e \text{ not complete}\}$.

Then $\text{Conf}(\text{ref}(\mathcal{E})) = \{\tilde{X} \mid \tilde{X} \text{ is a refinement of a configuration } X \in \text{Conf}(\mathcal{E})\}$.

Proof [84] or as a special case of Proposition 2.8. ■

Hence the configurations of $\text{ref}(\mathcal{E})$ are exactly those configurations which are refinements of configurations of $\mathcal{E}$. A refinement of a configuration X of $\mathcal{E}$ is obtained by replacing each event e in X by a non-empty configuration X_e of $\text{ref}(l_{\mathcal{E}}(e))$. Events which are causally necessary for other events in X may only be replaced by complete configurations.

2 Refinement of actions in flow event structures

In the previous section, we have indicated that for refining actions by event structures with conflicts more general models than prime event structures are appropriate. In BOUDOL & CASTELLANI [40] a form of event structures, called *flow event structures*, is suggested which is particularly suited for giving semantics to languages like CCS. Flow event structures are more general than prime events in the following sense: they do not assume conflict heredity and the axiom of finite causes, they allow inconsistent (self-conflicting) events and the causality relation is not required to be transitive and may even contain (syntactic) cycles. This makes it very easy to define operations like parallel composition and restriction, and we will show here that they are also well suited to deal with refinement of actions.

2.1 Definition

A (labelled) flow event structure (over an alphabet Act) is a 4-tuple $\mathcal{E} = (E, \prec, \#, l)$ where

- E is a set of events,
- $\prec \subseteq E \times E$ is an irreflexive relation, the *flow relation*,
- $\# \subseteq E \times E$ is a symmetric relation, the *conflict relation*,
- $l : E \rightarrow Act$ is the *labelling function*.

Let $\mathbb{E}$ denote the domain of flow event structures labelled over Act . The components of $\mathcal{E} \in \mathbb{E}$ will be denoted by $E_{\mathcal{E}}$, $\prec_{\mathcal{E}}$, $\#_{\mathcal{E}}$ and $l_{\mathcal{E}}$. The index $\mathcal{E}$ will be omitted if clear from the context. $\mathcal{E}$ is *conflict-free* if $\#_{\mathcal{E}} = \emptyset$. For $X \subseteq E_{\mathcal{E}}$, $\mathcal{E} \upharpoonright X = (X, \prec_{\mathcal{E}} \upharpoonright X, \#_{\mathcal{E}} \upharpoonright X, l_{\mathcal{E}} \upharpoonright X)$ is the *restriction of $\mathcal{E}$ to X* .

Two flow event structures $\mathcal{E}$ and $\mathcal{F}$ are *isomorphic* ($\mathcal{E} \cong \mathcal{F}$) iff there exists a bijection between their sets of events preserving $\prec$, $\#$ and labelling.

The interpretation of the conflict and the flow relation is formalised by defining configurations of flow event structures. Configurations must be conflict free; in particular, self-conflicting events will never occur in any configuration. $d \prec e$ will mean that d is a *possible immediate cause* for e . For an event to occur it is necessary that a *complete* non-conflicting set of its causes has occurred. Here a set of causes is complete if for any cause which is not contained there is a conflicting event which is contained. Finally, no cycles with respect to causal dependence may occur.


2.2 Definition Let $\mathcal{E} \in \mathbb{E}$.

- (i) $X \subseteq E$ is *left-closed in $\mathcal{E}$ up to conflicts* iff $\forall d, e \in E$: if $e \in X, d \prec e$ and $d \notin X$ then there exists an $f \in X$ with $f \prec e$ and $d \# f$.
 $X \subseteq E$ is *conflict-free* iff $\mathcal{E} \upharpoonright X$ is conflict-free.
- (ii) $X \subseteq E$ is a (*finite*) *configuration* of $\mathcal{E}$ iff X is finite, left-closed up to conflicts and conflict-free and does not contain a causality cycle: $\leq_X := (\prec \cap (X \times X))^*$ is an ordering. A configuration X is called *maximal* iff $X \subseteq Y \in \text{Conf}(\mathcal{E})$ implies $X = Y$. A configuration X is called *complete* iff $\forall d \in E : d \notin X \Rightarrow \exists e \in X$ with $d \# e$. $\text{Conf}(\mathcal{E})$ denotes the set of all configurations of $\mathcal{E}$.

The causal dependence between action occurrences in a configuration may again, as for prime event structures, be represented by a pomset; for $X \in \text{Conf}(\mathcal{E})$, we take the isomorphism class of $(X, \leq_X, l_{\mathcal{E}} \upharpoonright X)$.

2.3 Example

The system $((a+b) \parallel c); d$ may be represented by the flow event structure



(in graphical representations we omit names of events and represent $\prec$ by arcs of the form $\longrightarrow$).


The pomsets $\begin{smallmatrix} a \\ c \end{smallmatrix} \longrightarrow d$ and $\begin{smallmatrix} b \\ c \end{smallmatrix} \longrightarrow d$ correspond to complete configurations.

Note that prime event structures are special flow event structures defining $d \prec e$ iff $d < e$; the definition of configuration then coincides.

However, in contrast to prime event structures, not all maximal configurations are complete. Partly this is due to the fact that, in flow event structures, syntactic and semantic conflict not necessarily coincide, (two events are in *semantic conflict* if there is no configuration containing them both). Flow event structures where syntactic and semantic conflict coincide are called *faithful* in BOUDOL [38]. However, also in faithful flow event structures maximal configurations are not necessarily complete, either due to inconsistent events, but also in flow event structures without inconsistent events, as shown by the following example.

2.4 Example

Let $\mathcal{E} =$



The configuration $\{c_1, c_2, c_3\}$ is maximal but not complete.

Maximal but incomplete configurations may be interpreted as deadlocking behaviours. Assume that a semantic sequential composition is defined for flow event structures by putting all events in the first component in $\prec$ -relation with the events of the second component. Any incomplete maximal configuration of the first component would then disable the second component. Thus, in flow event structures, deadlock and termination may be distinguished.

2.5 Definition

A flow event structure $\mathcal{E}$ is *deadlock-free* iff every maximal configuration of $\mathcal{E}$ is complete.

Refinement of actions in flow event structures may now be defined as follows. We assume a refinement function $\text{ref} : \text{Act} \rightarrow \mathbb{E} - \{O\}$ (where O denotes the empty flow event structure) and replace each event labelled by a by a disjoint copy of $\text{ref}(a)$. The conflict and causality structure will just be inherited.

Hence, we may replace actions also by behaviours with conflicts and by infinite behaviours.

2.6 Definition

- (i) A function $\text{ref} : \text{Act} \rightarrow \mathbb{E} - \{O\}$ is called a *refinement function* (for flow event structures).
- (ii) Let $\mathcal{E} \in \mathbb{E}$ and let ref be a refinement function.
Then the *refinement of $\mathcal{E}$ by ref* , $\text{ref}(\mathcal{E})$, is the flow event structure defined by
 - $E_{\text{ref}(\mathcal{E})} = \{(e, e') \mid e \in E_{\mathcal{E}}, e' \in E_{\text{ref}(l_{\mathcal{E}}(e))}\},$
 - $(d, d') \prec_{\text{ref}(\mathcal{E})} (e, e')$ iff $d \prec e$ or $(d = e \wedge d' \prec_{\text{ref}(l_{\mathcal{E}}(d))} e'),$
 - $(d, d') \#_{\text{ref}(\mathcal{E})} (e, e')$ iff $d \#_{\mathcal{E}} e$ or $(d = e \wedge d' \#_{\text{ref}(l_{\mathcal{E}}(d))} e'),$
 - $l_{\text{ref}(\mathcal{E})}(e, e') = l_{\text{ref}(l_{\mathcal{E}}(e))}(e').$

As for prime event structures, we verify that $\text{ref}(\mathcal{E})$ is well-defined, even when isomorphic flow event structures are identified.

2.7 Proposition

- (i) If $\mathcal{E} \in \mathbb{E}$ and ref is a refinement function then $\text{ref}(\mathcal{E})$ is a flow event structure indeed.
- (ii) If $\mathcal{E} \in \mathbb{E}$ and ref, ref' are refinement functions with $\text{ref}(a) \cong \text{ref}'(a)$ for all $a \in \text{Act}$ then $\text{ref}(\mathcal{E}) \cong \text{ref}'(\mathcal{E})$.
- (iii) If $\mathcal{E}, \mathcal{F} \in \mathbb{E}$ and ref is a refinement function then $\text{ref}(\mathcal{E}) \cong \text{ref}(\mathcal{F})$.

Proof Straightforward. ■

Finally, we show that, analogously to prime event structures, the behaviour of a refined flow event structure $\text{ref}(\mathcal{E})$ may be deduced compositionally from the behaviour of $\mathcal{E}$ and the behaviour of the refinements of actions.

2.8 Proposition

Let $\mathcal{E} \in \mathbb{E}$, let ref be a refinement function for flow event structures.

We call $\tilde{X}$ a *refinement of configuration* $X \in \text{Conf}(\mathcal{E})$ by ref iff

- $\tilde{X} = \bigcup_{e \in X} \{e\} \times X_e$ where $\forall e \in X : X_e \in \text{Conf}(\text{ref}(l_{\mathcal{E}}(e))) - \{\emptyset\}$,
- $e \in \text{busy}(\tilde{X}) \implies e$ maximal in X with respect to $\leq_X$
where $\text{busy}(\tilde{X}) := \{e \in X \mid X_e \text{ not complete}\}$.

Then $\text{Conf}(\text{ref}(\mathcal{E})) = \{\tilde{X} \mid \tilde{X} \text{ is a refinement of a configuration } X \in \text{Conf}(\mathcal{E})\}$.

Proof

“ $\subseteq$ ” Let $\tilde{X} \in \text{Conf}(\text{ref}(\mathcal{E}))$.

First we show that $X := \text{pr}_1(\tilde{X}) \in \text{Conf}(\mathcal{E})$.

X is finite since $\tilde{X}$ is finite.

X is left-closed in $\mathcal{E}$ up to conflicts:

Let $e \in X, d \in E_{\mathcal{E}}$ with $d \prec_{\mathcal{E}} e$ and $d \notin X$.

We have to show that there exists an $f \in X$ with $f \prec_{\mathcal{E}} e$ and $f \#_{\mathcal{E}} d$.

Since $e \in X$ there must be some $(e, e') \in \tilde{X}$.

There exists $(d, d') \in E_{\text{ref}(\mathcal{E})}$, $(d, d') \notin \tilde{X}$ since $\text{ref}(d) \neq O$ and $d \notin X$.

Furthermore $(d, d') \prec_{\text{ref}(\mathcal{E})} (e, e')$ since $d \prec_{\mathcal{E}} e$.

So $\exists (f, f') \in \tilde{X}$ with $(f, f') \prec_{\text{ref}(\mathcal{E})} (e, e')$ and $(f, f') \#_{\text{ref}(\mathcal{E})} (d, d')$.

$f \neq d$ since $f \in X, d \notin X$; hence $f \#_{\mathcal{E}} d$.


If $f \neq e$ we have $f \prec_{\mathcal{E}} e$ and we are done.

Assume $f = e$ then $(d, d') \prec_{\text{ref}(\mathcal{E})} (f, f')$.

Then $\exists (g, g') \in \tilde{X}$ with $(g, g') \prec_{\text{ref}(\mathcal{E})} (f, f') = (e, e')$ and $(g, g') \#_{\text{ref}(\mathcal{E})} (d, d')$.

$g \#_{\mathcal{E}} d$ since $g \neq d$. Furthermore $g \in X$.

If $g \neq f = e$ then $g \prec_{\mathcal{E}} e$ and we are done. Since $\tilde{X}$ is finite, we will find (by repeating this), after finitely many steps, $(\tilde{f}, \tilde{f}') \in \tilde{X}$ with $\tilde{f} \#_{\mathcal{E}} d$ and $\tilde{f} \prec_{\mathcal{E}} e$. Hence X is left-closed up to conflicts.



X is conflict-free:

Assume $d, e \in X$ with $d \#_{\mathcal{E}} e$.

Then there exist $(d, d'), (e, e') \in \tilde{X}, (d, d') \#_{\text{ref}(\mathcal{E})} (e, e')$.

This is a contradiction since $\tilde{X}$ is conflict-free.

Finally we have to show that X does not contain a causality-cycle.

Assume $d, e \in X, d \neq e, d \leq_X e$ and $e \leq_X d$ (where $\leq_X$ is derived from $\prec_{\mathcal{E}}$).

It is straightforward to verify that this implies $\exists (d, d'), (e, e') \in \tilde{X}$ with $(d, d') \neq (e, e'), (d, d') \leq_{\tilde{X}} (e, e')$ and $(e, e') \leq_{\tilde{X}} (d, d')$. This is in contradiction with the cyclefreeness of $\tilde{X}$.

Hence $X = \text{pr}_1(\tilde{X}) \in \text{Conf}(\mathcal{E})$. We will show that $\tilde{X}$ is a refinement of X .

Let $e \in X$ and $X_e := \{e' \mid (e, e') \in \tilde{X}\}$. By construction $X_e \neq \emptyset$.

Let $\mathcal{E}_e := \text{ref}(l_{\mathcal{E}}(e))$. We want to show that $X_e \in \text{Conf}(\mathcal{E}_e)$.

Obviously $X_e \subseteq E_{\mathcal{E}_e}$.

X_e is finite, conflict-free and cycle-free since $\tilde{X}$ is finite, conflict-free and cycle-free. So it only remains to be shown that X_e is left-closed up to conflicts.

Let $d' \in \mathcal{E}_e, d' \prec_{\mathcal{E}_e} e' \in X_e, d' \notin X_e$.

Then $(e, d') \in E_{\text{ref}(\mathcal{E})}, (e, d') \prec_{\text{ref}(\mathcal{E})} (e, e') \in \tilde{X}$ and $(e, d') \notin \tilde{X}$.

So there exists $(f, f') \in \tilde{X}$ with $(f, f') \prec_{\text{ref}(\mathcal{E})} (e, e')$ and $(f, f') \#_{\text{ref}(\mathcal{E})} (e, d')$.

$f, e \in X \implies \neg(f \#_{\mathcal{E}} e) \implies f = e \wedge f' \#_{\mathcal{E}_e} d' \implies f' \in X_e$ and $f' \prec_{\mathcal{E}_e} e'$.

Hence $X_e \in \text{Conf}(\mathcal{E}_e)$.

From what we have shown by now it follows that $\tilde{X} = \bigcup_{e \in X} \{e\} \times X_e$

with $X \in \text{Conf}(\mathcal{E})$ and, for all $e \in X, X_e \in \text{Conf}(\text{ref}(l_{\mathcal{E}}(e))) - \{\emptyset\}$.

Now let $e \in \text{busy}(\tilde{X})$. We have to show that e is maximal in X w.r.t. $\leq_X$.

Suppose e is not maximal in X .

Then there exists $f \in X$ with $e \prec_{\mathcal{E}} f$, and there exists $(f, f') \in \tilde{X}$.

Since X_e is not complete there exists $d' \in E_{\mathcal{E}_e} - X_e$ with

$$(*) \quad \forall e' \in X_e : \neg(d' \#_{\mathcal{E}_e} e').$$

We have $(e, d') \prec_{\text{ref}(\mathcal{E})} (f, f'), (e, d') \notin \tilde{X}$.

Since $\tilde{X}$ is a configuration, there then exists $(g, g') \in \tilde{X}$ with $(g, g') \prec_{\text{ref}(\mathcal{E})} (f, f')$ and $(g, g') \#_{\text{ref}(\mathcal{E})} (e, d')$.

Since $g, e \in X$, we have $\neg(g \#_{\mathcal{E}} e)$.

Hence $g = e$ and $g' \in X_e, g' \#_{\mathcal{E}_e} d'$.

However this contradicts $(*)$.

“ $\supseteq$ ” Let $\tilde{X}$ be a refinement of $X \in \text{Conf}(\mathcal{E})$. We show that $\tilde{X} \in \text{Conf}(\text{ref}(\mathcal{E}))$.

It follows in a straightforward manner from the corresponding properties of X and the X_e 's that $\tilde{X}$ is finite and conflict-free and contains no causality cycles. Hence it suffices to show that $\tilde{X}$ is left-closed up to conflicts.

So let $(e, e') \in \tilde{X}$, let $(d, d') \in E_{\text{ref}(\mathcal{E})} - \tilde{X}$ with $(d, d') \prec_{\text{ref}(\mathcal{E})} (e, e')$.

We have to show that there exists $(f, f') \in \tilde{X}$ with $(f, f') \prec_{\text{ref}(\mathcal{E})} (e, e')$ and $(f, f') \#_{\text{ref}(\mathcal{E})} (d, d')$.

First assume $d = e$. Then this follows immediately from the corresponding property of X_e .

Now let $d \neq e$.

If $d \notin X$ then the requirement follows from the corresponding property of X .

So we now consider the remaining case that $d \neq e$ and $d \in X$. Then $d' \notin X_d$.

Since $d \neq e$ we have $d \prec_{\mathcal{E}} e$, hence d is not maximal in X .

Thus X_d must be complete.

So $d' \notin X_d$ implies $\exists f' \in X_d$ with $f' \#_{\text{ref}(l_{\mathcal{E}}(d))} d'$.

Hence $(d, f') \in \tilde{X}$, $(d, f') \prec_{\text{ref}(\mathcal{E})} (e, e')$ and $(d, f') \#_{\text{ref}(\mathcal{E})} (d, d')$. ■

We end this section with a lemma that will be useful later on.

2.9 Lemma Let $\mathcal{E} \in \mathbb{E}$, $X \in \text{Conf}(\mathcal{E})$ and $\text{busy} \subseteq X$.

Then $\forall e \in \text{busy} : e \text{ maximal in } X \text{ with respect to } \leq_X$
 $\iff \forall Y \subseteq \text{busy} : X - Y \in \text{Conf}(\mathcal{E})$.

Proof

" $\implies$ " Let $\mathcal{E} \in \mathbb{E}$, $X \in \text{Conf}(\mathcal{E})$, $Y \subseteq X$ and $\forall e \in Y : e \text{ maximal in } X \text{ w.r.t. } \leq_X$. It suffices to prove that $X - Y \in \text{Conf}(\mathcal{E})$. $X - Y$ is finite and conflict-free and does not contain causality cycles since X has these properties. It remains to be shown that $X - Y$ is left-closed up to conflicts.

Suppose $e \in X - Y$, $d \prec_{\mathcal{E}} e$ and $d \notin X - Y$. If $d \in Y$ then d would be maximal in X w.r.t. $\leq_X$, contradicting $d \prec_{\mathcal{E}} e$. Thus $d \notin X$. Hence there is an $f \in X$ with $f \prec_{\mathcal{E}} e$ and $d \#_{\mathcal{E}} f$. Since $f \prec_{\mathcal{E}} e$, f is not maximal in X w.r.t. $\leq_X$, so $f \in X - Y$, which had to be proven.

" $\impliedby$ " Let $\mathcal{E} \in \mathbb{E}$, $X \in \text{Conf}(\mathcal{E})$, $d \in X$ and $X - \{d\} \in \text{Conf}(\mathcal{E})$. It suffices to prove that d is maximal w.r.t. $\leq_X$.

Suppose it is not, then $\exists e \in X$ with $d \prec_{\mathcal{E}} e$. Since $X - \{d\} \in \text{Conf}(\mathcal{E})$, there exists an $f \in X - \{d\}$ with $f \prec_{\mathcal{E}} e$ and $d \#_{\mathcal{E}} f$, contradicting the conflict-freeness of X . ■

This means that, in Proposition 2.8, the condition " $e \in \text{busy}(\tilde{X}) \implies e \text{ maximal in } X \text{ w.r.t. } \leq_X$ " can be replaced by "for all $Y \subseteq \text{busy}(\tilde{X})$, $X - Y \in \text{Conf}(\mathcal{E})$ ".

3 Configuration structures and refinement of actions

In the previous section we have shown that flow event structures may be used for refinement of actions, even when substituting actions by behaviours with conflicts or by infinite behaviours. However, the refinement operation we have defined depends on the particular “syntax” of flow event structures. In this section, our aim is to define a refinement operation for a very general model of concurrent systems, such that refinement operations for particular representations, as flow event structures, are obtained as a special case.

We will consider a model where a system is represented by its set of configurations. As in the previous sections, occurrences of actions are represented by *events* labelled by the corresponding action names. A *configuration* is a set of events representing a state of the system where exactly its elements have happened. We only consider finite configurations here. Following ideas of WINSKEL [183] we represent a system by a family of configurations satisfying certain consistency requirements.

3.1 Definition

A (*labelled*) *configuration structure* (over an alphabet *Act*) is a pair $\mathcal{C} = (C, l)$ where C is a family of finite sets (*configurations*) such that

- $\emptyset \in C$,
- $X, Y, Z \in C, X \cup Y \subseteq Z \implies X \cup Y \in C$,
- $X \in C \wedge d, e \in X, d \neq e \implies \exists Y \in C$ with $Y \subseteq X$ and $(d \in Y \iff e \notin Y)$,

and $l : \bigcup_{X \in C} X \rightarrow \text{Act}$ is a *labelling function*.

The requirements for a family of sets of events to form a configuration structure may be explained as follows. The initial state of a system is the state where no action has been performed yet. Hence $\emptyset$ is always a configuration. Now, if two configurations X, Y are contained in a third configuration Z then $X \cup Y$ is consistent or *conflict-free*; e.g. all its elements can happen together in one run. Since both X and Y represent already possible runs, it should then also be possible to execute just the events in X and Y , hence $X \cup Y$ should be a configuration. If we consider two distinct events occurring in some run, then there must be an intermediate state where already one of them has occurred whereas the other has not yet occurred (coincidence can not be enforced). This is guaranteed by the third requirement.

Finally, a remark on our requirement that configurations should be finite. As usual, we assume that in a finite period only finitely many actions may be performed. Now the requirement says that we only consider states that are reachable in a finite period of time. WINSKEL [183] allows configurations to be infinite, thus representing also those states which can be reached in an infinite period of time. However, his infinite configurations are completely determined

by the finite ones. Hence configuration structures as defined here are equally expressive as Winskel's families of configurations.


Convention We will denote the components of a configuration structure $\mathcal{C}$ by $C_{\mathcal{C}}$ and $l_{\mathcal{C}}$ respectively. By abuse of language, $C_{\mathcal{C}}$ will also be denoted by $\mathcal{C}$. Furthermore the set $E_{\mathcal{C}}$ of events of $\mathcal{C}$ is defined by $E_{\mathcal{C}} = \bigcup_{X \in \mathcal{C}} X$.

Let $\mathcal{C}$ denote the domain of configuration structures labelled over Act .


3.2 Example

We consider the example referred to as a “parallel switch” in [183].

We have two actions 0 and 1 interpreted as closing switch 0 and closing switch 1, respectively, in an electric circuit. As soon as at least one of the switches is closed, a bulb lights up; this is represented as an action b .



This may be represented by the following configuration structure (with a unique correspondence between actions and events):



The b -event may occur here without a unique “causal history”; in the configuration $\{0, 1, b\}$ it is not clear whether b is caused by 0 or by 1.

Usually, the names of events are not important; hence we will not distinguish configuration structures which are isomorphic in the sense that they only differ with respect to names of events.

3.3 Definition

A configuration structure isomorphism between two configuration structures $\mathcal{C}, \mathcal{D} \in \mathcal{C}$ is a bijective mapping $f : E_{\mathcal{C}} \rightarrow E_{\mathcal{D}}$ such that

- $X \in \mathcal{C} \iff f(X) \in \mathcal{D}$ for $X \subseteq E_{\mathcal{C}}$,
- and $l_{\mathcal{D}}(f(e)) = l_{\mathcal{C}}(e)$ for $e \in E_{\mathcal{C}}$.

$\mathcal{C}$ and $\mathcal{D}$ are *isomorphic* — notation $\mathcal{C} \cong \mathcal{D}$ — if there exists a configuration structure isomorphism between them.

In configuration structures, completeness and maximality of configurations coincide. Deadlock and termination may not be distinguished.

3.4 Definition

A configuration X of a configuration structure $\mathcal{C}$ is called *complete* iff there is no $Y \neq X$ in $\mathcal{C}$ containing X .

We may now associate a configuration structure with each flow event structure (and via this also with each prime event structure).

3.5 Definition

Let $\mathcal{E} \in \mathbb{E}$.

The configuration structure of $\mathcal{E}$, $\mathcal{C}(\mathcal{E})$, is defined as

$$\mathcal{C}(\mathcal{E}) = (\text{Conf}(\mathcal{E}), l_{\mathcal{E}} \upharpoonright \bigcup_{X \in \text{Conf}(\mathcal{E})} X).$$

There is no unique correspondence in general: different flow event structures may have the same configuration structure (but not vice versa). In particular, the distinction between deadlock and termination is lost.

Next, we define refinement of actions for configuration structures. A refinement will be specified by a function ref specifying for each action a a configuration structure $\text{ref}(a)$ which is to be substituted for a . Again we only consider non-forgetful refinements here, hence $\text{ref}(a) \neq O$ for all $a \in \text{Act}$ where O denotes the empty configuration structure with $C_O = \{\emptyset\}$. Apart from this restriction, we may replace an action by any configuration structure.

3.6 Definition

(i) A function $\text{ref} : \text{Act} \rightarrow \mathcal{C} - \{O\}$ is called a *refinement function (for configuration structures)*.

(ii) Let $\mathcal{C}$ be a configuration structure and let ref be a refinement function.

We call $\tilde{X}$ a *refinement of a configuration* $X \in \mathcal{C}$ by ref iff

$$- \tilde{X} = \bigcup_{e \in X} \{e\} \times X_e \text{ where } \forall e \in X : X_e \in \text{ref}(l_{\mathcal{C}}(e)) - \{\emptyset\},$$

$$- \text{for all } Y \subseteq \text{busy}(\tilde{X}), X - Y \in \mathcal{C},$$

$$\text{where } \text{busy}(\tilde{X}) = \{e \in X \mid X_e \text{ not complete}\}.$$

The *refinement of $\mathcal{C}$ by ref* is defined as $\text{ref}(\mathcal{C}) = (C_{\text{ref}(\mathcal{C})}, l_{\text{ref}(\mathcal{C})})$ with

$$C_{\text{ref}(\mathcal{C})} := \{\tilde{X} \mid \tilde{X} \text{ is a refinement of some } X \in \mathcal{C} \text{ by } \text{ref}\}$$

and

$$l_{\text{ref}(\mathcal{C})}(e, e') = l_{\text{ref}(l_{\mathcal{C}}(e))}(e') \text{ for all } (e, e') \in \bigcup_{\tilde{X} \in C_{\text{ref}(\mathcal{C})}} \tilde{X}.$$

Intuitively, this definition may be explained as follows.

The configuration structure $ref(\mathcal{C})$ is obtained by taking all possible refinements of configurations of $\mathcal{C}$. A refinement of a configuration X of $\mathcal{C}$ is obtained by replacing each event e in X by a non-empty configuration X_e of $ref(l_{\mathcal{C}}(e))$. Events which are causally necessary for other events in X may only be replaced by complete configurations, hence it must be possible to take any subset of "uncompleted" or busy events out of X , again obtaining a configuration.

Next we show that refinement is a well-defined operation on configuration structures, even when isomorphic configuration structures are identified.

3.7 Proposition

- (i) If $\mathcal{C} \in \mathcal{C}$ and ref is a refinement function then also $ref(\mathcal{C})$ is a configuration structure.
- (ii) If $\mathcal{C} \in \mathcal{C}$ and ref, ref' are refinement functions with $ref(a) \cong ref'(a)$ for all $a \in Act$ then $ref(\mathcal{C}) \cong ref'(\mathcal{C})$.
- (iii) If $\mathcal{C}, \mathcal{D} \in \mathcal{C}$, ref is a refinement function and $\mathcal{C} \cong \mathcal{D}$ then $ref(\mathcal{C}) \cong ref(\mathcal{D})$.

Proof (i) cumbersome and omitted here, (ii) and (iii) straightforward. ■

Finally, we want to show that the easier syntactic refinement operation for flow event structures defined in section 2 is consistent with the refinement operation for configuration structures. However, since the distinction between deadlock and termination is lost in configuration structures, this is only true for deadlock-free refinements.

3.8 Theorem

Let $\mathcal{E} \in \mathcal{E}$, let ref be a refinement function for flow event structures with $\forall a \in Act : ref(a)$ deadlock-free.

Then $\mathcal{C}(ref(\mathcal{E})) = ref'(\mathcal{C}(\mathcal{E}))$

where $ref'(a) = \mathcal{C}(ref(a))$ for all $a \in Act$.

Proof

It has to be shown that $\mathcal{C}_{\mathcal{C}(ref(\mathcal{E}))} = \mathcal{C}_{ref'(\mathcal{C}(\mathcal{E}))}$ and $l_{\mathcal{C}(ref(\mathcal{E}))} = l_{ref'(\mathcal{C}(\mathcal{E}))}$. The first requirement translates to

$$Conf(ref(\mathcal{E})) = \{\tilde{X} | \tilde{X} \text{ is refinement of some } X \in Conf(\mathcal{E}) \text{ by } ref'\}.$$

From Proposition 2.8 we know

$$Conf(ref(\mathcal{E})) = \{\tilde{X} | \tilde{X} \text{ is a refinement of some } X \in Conf(\mathcal{E}) \text{ by } ref\}.$$

So it suffices to establish that a refinement by ref' is the same as a refinement by ref . This follows immediately from Proposition 2.8 and Lemma 2.9 in combination with Definition 3.6, provided that for $a \in Act : X$ is a complete configuration of $ref(a)$ iff X is complete in $ref'(a) = \mathcal{C}(ref(a))$. This is the case if ref is deadlock-free.

The second requirement is straightforward. ■

4 Refinement of transitions in Petri nets

We start by giving some basic definitions and notations for Petri nets; for explanations and concepts we refer to introductory texts on nets, such as REISIG [156] or PETERSON [148].

For simplicity we assume that there is a one to one correspondence between the transitions in the net and the actions that the system modelled by the net can perform; we do not consider nets with labelled transitions. However, we will show later that our approach can easily be extended to this case.

4.1 Definition

- $N = (S, T, F)$ is called a *net structure* iff
- S is a set (of *places*),
 - T is a set (of *transitions*), $S \cap T = \emptyset$,
 - $F \subseteq (S \times T) \cup (T \times S)$ such that
 $\forall t \in T : \exists s, s' \text{ with } sFt \text{ and } tFs'$ (transitions have non-empty pre- and postsets)
 $\text{and } \forall s \in S : sFt \implies \neg tFs$ (no self-loops).

The restrictions we have made here—non-empty pre- and postsets of transitions and no self-loops—will be needed for our refinement construction.

Two nets $N = (S, T, F)$ and $N' = (S', T', F')$ are *isomorphic* — notation $N \cong N'$ — if $T = T'$ and there exists a bijective mapping $f : S \rightarrow S'$ satisfying $sFt \iff f(s)F't$ and $tFs \iff tF'f(s)$.

Generally, we will not distinguish isomorphic net structures.

As usual, we use the following notations.

For $x \in S \cup T$, let $\bullet x := \{y \in S \cup T \mid yFx\}$ (*preset of x*),

$x^\bullet := \{y \in S \cup T \mid xFy\}$ (*postset of x*).

Let ${}^\circ N := \{x \in S \cup T \mid \bullet x = \emptyset\}$ (*initial places of N*),

$N^\circ := \{x \in S \cup T \mid x^\bullet = \emptyset\}$ (*final places of N*).

Note that ${}^\circ N, N^\circ \subseteq S$.

The components of a net N will be denoted by S_N, T_N, F_N (the index is omitted when clear from the context). We will sometimes use the characteristic mapping of F as a function $F : (S \times T) \cup (T \times S) \rightarrow \{0, 1\}$.

A concurrent system may be modelled by a net structure where the places carry tokens, indicating the state of the system. The dynamic behaviour of the system is derived by the so called *firing rule*. We assume that all places have unbounded capacities; any mapping $M : S_N \rightarrow \mathbb{N}$ will be called a *marking* of the net N . However, we will restrict our considerations to one-safe nets here. We will illustrate later why refinement in non-one-safe nets may lead to problems.

4.2 Definition

(N, M_o) is called a P/T-system or a *marked net* iff N is a net structure and $M_o : S \rightarrow \mathbb{N}$ (*initial marking*).

By abuse of notation, we will use N both for (N, M_o) (when M_o is clear from the context) and for the underlying net structure.

4.3 Definition Let (N, M_o) be a marked net, let $M, M' : S \rightarrow \mathbb{N}$ and $t \in T$.

- (i) t is *enabled by* M iff $\forall s \in \bullet t : M(s) > 0$.
- (ii) M' is reached from M by firing t ($M[t > M']$) iff
 - t is enabled by M and
 - $\forall s \in S : M'(s) = M(s) - F(s, t) + F(t, s)$.

The *marking class* $[N, M_o >$ of a marked net (N, M_o) is then defined as the set of all markings reachable from M_o by finitely many transition firings. A marked net is *one-safe* if $\forall M \in [N, M_o >, \forall s \in S : M(s) \leq 1$. In one-safe nets, we may use set notations for markings: $M \subseteq S$ is the marking where exactly the places in M carry a token.

Whenever referring to a marked net in the following, we assume it to be one-safe.

A conceptual framework for refinement in Petri nets are *net morphisms* [75]. A net morphism is a mapping between the elements of two net structures such that the distinction between places and transitions is observed to some extent. It is possible to map, for example, a place to a transition, but only if this place is surrounded by transitions with the same image.

4.4 Definition Let $N = (S, T, F)$, $N' = (S', T', F')$ be net structures.

- (i) A mapping $f : S \cup T \rightarrow S' \cup T'$ is called a *net morphism* iff
 - $\forall x, y \in S \cup T$ with $f(x) \neq f(y)$ and $(x, y) \in F : [(f(x), f(y)) \in F' \text{ and } x \in S \Leftrightarrow f(x) \in S']$.


- (ii) A net morphism $f : S \cup T \rightarrow S' \cup T'$ is called a *quotient* iff f is surjective and $(x', y') \in F' \implies \exists (x, y) \in F$ with $f(x) = x', f(y) = y'$ (surjectivity also with respect to arcs).

A quotient can be thought of as a factorisation. The net is partitioned such that sorts are preserved: each subset of elements forming a class in this partition must have a boarder consisting just of places or just of transitions and is then considered as one place or one transition, respectively. A quotient N_1 of a net N_2 is considered as an *abstraction* of N_2 (REISIG [157]). Conversely, N_2 is then called a *refinement* of N_1 . In this framework, transitions as well as places may be refined.


However, behavioural aspects are not taken precisely into account and this may lead to problems.

4.5 Example


Consider




The net N_1 is an abstraction of



by the quotient mapping all elements inside the broken line to r (and otherwise the identity). Conversely, N_2 is considered as a refinement of N_1 . However, consider the slightly enlarged systems



and



Again, N'_1 is a quotient of N'_2 , hence N'_2 may be considered as a refinement of N'_1 .

Assuming that places 1 and 2 are initially marked, we find that the net N'_1 is deadlock-free in the sense that it is possible to fire transitions until the two final places are both marked. However, even though the part of N'_2 corresponding to r is also deadlock-free (namely N_2 is deadlock-free), N'_2 may reach a deadlock situation by firing t and t' .

This shows that the notion of a net morphism or quotient is in general not strong enough to reason about the behaviour of refinements in a compositional way. An attempt to restrict it in such a way that behavioural aspects are taken more strongly into account has been made in DESEL & MERCERON [70]. They identify a subclass of morphisms they call *vicinity respecting*. The essential idea is that those net morphisms respect the impact of elements on their environment.

4.6 Definition

A net morphism $f : N \rightarrow N'$ is said to be *vicinity respecting* iff $\forall x \in SUT$:

- $f(\odot x) = \{f(x)\} \vee f(\odot x) = \odot f(x)$, and
- $f(x^\odot) = \{f(x)\} \vee f(x^\odot) = f(x)^\odot$,

where $\odot x := \{x\} \cup \bullet x$, $x^\odot := \{x\} \cup x^\bullet$, respectively.

The morphisms considered in Example 4.5 are not vicinity respecting. We will discuss later to what extent this notion does indeed characterise the refinements we are interested in.

In order to avoid confusion, we have to mention here another notion of morphism suggested for Petri nets by WINSKEL [183]. This notion is particularly tailored to take behavioural aspects into account, however it does not allow to contract for example a line of two transitions into one transition. So it is not suited for treating refinement. More recent approaches in the categorical framework [131, 122] have not yet been evaluated under this aspect.

For the case of refining transitions, which we are interested in here, also more constructive approaches are being considered explaining how to replace a transition in a net by a “refinement net”. The problem is to specify how to connect the “refinement net” to the environment of the refined transition, and to investigate what restrictions on refinement nets are then necessary for a sensible refinement operation.

One possibility is to require a one to one correspondence between “input/output places” of the refinement net and the surrounding places of the refined transition. In VOGLER [174], a construction for this case is proposed, and it is shown that it is then necessary to impose certain restrictions on refinement nets, in particular disallowing initial concurrency (otherwise a situation as in Example 4.5 might occur).

Most constructions for refining transitions are based on distinguishing initial and final transitions in a refinement net and connecting them to the preset and postset, respectively, of the refined transition (VALETTE [173] and subsequently SUZUKI & MURATA [164], VOGLER [175] and BEST, DEVILLERS, KIEHN & POMELLO [33]).


In these approaches, the main idea is that a transition may only be replaced by a net behaving like a transition with respect to its effect on the environment:

- it cannot move without being activated by the environment,
- it has the same possible behaviours whenever it is activated,
- it may not deadlock,
- it consumes and produces tokens in a coincident manner.

The final condition ensures that the problematic situation explained in Example 4.5 may not occur. VALETTE [173] and others ensure this property by allowing only refinements for transitions with at most one initial and at most one final transition. VOGLER [175] generalises this by allowing several initial transitions which must be in conflict (and, symmetrically, the same for final transitions). This means that we may not have initial or final concurrency in refinement nets.

The other requirements are usually ensured by extending the net which is supposed to be substituted for a transition by a new place supplying a token to the initial transition(s) and receiving a token from the final transition(s) and then analysing the behaviour of this net. This is illustrated on the next page.


The interesting problem discussed in Example 4.5 was to refine a transition by some behaviour exhibiting initial concurrency. Symmetrically, we also want to allow refinements with final concurrency. This may not be handled in these approaches (VOGLER [174] excludes only initial concurrency). A possibility to get rid of this restriction which has not yet been pursued further is to restrict the environment of transitions which are refined.




Here we propose a construction which generalises the approach of VALETTE [173] and VOGLER [175] for the class of one-safe nets without self-loops, and which offers the possibility of refining transitions also with initial and final concurrency. This will be achieved by extending these approaches by specifying explicitly which initial transitions should be concurrent or in conflict (additionally to constraints already imposed by the internal structure of the refinement net). For this, we extend the refinement net with initial places in the preset of initial transitions. Similarly, we add end places specifying the relationship between final transitions. Clearly, in the refinement net, initial places have no ingoing arcs and final places have no outgoing arcs. When analysing the behaviour of a refinement net, we assume that all initial places (and no final places) carry tokens. As in [173], we allow that also other places in a refinement net carry initial tokens. The approaches of [173] and [175] may be seen as a special case of our approach by splitting the start-place considered above into two places: one initial and one final place. Since we will require as [173] that a refinement net has the same possible behaviour whenever it is activated, it is reasonable to assume that the initial places are just those places without ingoing arcs and the final places just those without outgoing arcs. The initial and final places will then be used in the embedding construction to ensure that causal dependencies are preserved by the refinement operation.

4.7 Example

Consider again the net N'_1 of Example 4.5.




We tried to refine r by two concurrent transitions followed by another transition which causally depends on both of them. This refinement of r may be represented as




Places 7 and 8 are initial places, place 9 is the final place.

Now R is inserted into N'_1 for the transition r by taking the cartesian product of the preplaces of r with the initial places of R and of the postplaces of r with the final places of R . We obtain





N''_2 is again a quotient of N'_1 , however the mapping between places is no longer the identity. We see that, even though tokens are not removed coincidentally by the refinement of r , we have ensured that either both transitions in the refinement of r will fire or none of them, hence N''_2 will not deadlock. This has been achieved by preserving precisely the conflict and causality structure.

In contrast to the approaches similar to VALETTE [173], we do allow to refine transitions by deadlocking behaviours (where we use the word deadlock in the usual intuitive meaning rather than in the net theoretic sense). The reason is that we do not expect that the properties of the original net, like deadlock-freeness, are preserved by refinement. We only require that the properties of the resulting net are derivable in a compositional way. Whether or not a net to be inserted deadlocks is specified by its behaviour with respect to its final places. A refinement net *deadlocks* if it may reach a situation where no transition may fire but not all its final places are marked. This may be explained by putting the refinement net in a context by connecting its final and its initial places by a transition.




The refinement net deadlocks iff t may not occur.

4.8 Example




R will deadlock since not all its final places can get a token.

When replacing R for r , we get




where t will never occur.

However, replacing r by



gives




where t will occur.

The next example shows that it is not possible to consider places which have ingoing arcs as initial places of a refinement net.

4.9 Example


Let

$N =$




and

consider the net




If we would replace R for r , we would obtain




which has not the expected behaviour, since once the refined r has been chosen, no a should be possible any more.

This problem can be solved by using labelled nets and unfolding R into



Inserting R' into N yields



which has indeed the expected behaviour.

Next, we will define our construction formally and, in particular, describe formally the requirements on nets which may be inserted for transitions. We will then relate our construction to the notion of vicinity respecting net morphisms and to our approach for refinement in event structures.

4.10 Definition

(N, M_o) with $N = (S, T, F)$ is a *refinement net* iff

- ${}^oN \neq \emptyset$ and $N^o \neq \emptyset$,
- ${}^oN \subseteq M_o$ and $N^o \cap M_o = \emptyset$,
- no $t \in T$ is enabled by $M_o - {}^oN$,
- for any $M \in [N, M_o >$ with $N^o \subseteq M$ we have $M - N^o = M_o - {}^oN$,
(N will exhibit identical behaviour when reactivated).

4.11 Definition Let (N, M_o) be a marked net and $r \in T_N$.

Let (R, M_o^R) be a refinement net, w.l.o.g. $T_N \cap T_R = \emptyset$, $S_N \cap S_R = \emptyset$.

Then $N[R/r] := (S, T, F)$ is defined by

$$\begin{aligned}
 S &:= (S_N - (\bullet r \cup r \bullet)) \cup (S_R - ({}^o R \cup R^o)) \cup Int \\
 &\quad \text{where } Int := (\bullet r \times {}^o R) \cup (r \bullet \times R^o), \\
 T &:= (T_N - \{r\}) \cup T_R, \\
 F &:= (F_N \cup F_R) \upharpoonright (S \times T \cup T \times S) \\
 &\quad \cup \{((s_N, s_R), t) \mid (s_N, s_R) \in Int, \\
 &\quad (t \in T_N \setminus \{r\} \wedge (s_N, t) \in F_N) \vee (t \in T_R \wedge (s_R, t) \in F_R)\} \\
 &\quad \cup \{(t, (s_N, s_R)) \mid (s_N, s_R) \in Int, \\
 &\quad (t \in T_N \setminus \{r\} \wedge (t, s_N) \in F_N) \vee (t \in T_R \wedge (t, s_R) \in F_R)\}
 \end{aligned}$$


and $(N, M_o)[R/r] = (N[R/r], M_o^{[R/r]})$ with

$$\begin{aligned}
 M_o^{[R/r]}(s) &= M_o(s) \text{ iff } s \in S_N, \quad M_o^{[R/r]}(s) = M_o^R(s) \text{ iff } s \in S_R, \\
 M_o^{[R/r]}(s) &= M_o(s_N) \text{ iff } s = (s_N, s_R) \in Int.
 \end{aligned}$$


It is straightforward to verify that $N[R/r]$ is again a one-safe net.

The following example illustrates why we restrict ourselves to one-safe nets (a similar example is given in BEST, DEVILLERS, KIEHN & POMELLO [33]).

4.12 Example



When replacing r by R , we would obtain



However, this net has not the expected behaviour, since the two independent occurrences of the refined r -transition may now cooperate and execute d . As remarked in VALETTE [173], this problem can only occur if in N the refined transition can be “two-enabled”.

Next we show that the order in which transitions are replaced does not matter. In particular, this means (at least for finite nets) that we can extend our approach to non-injective labellings of transitions by action names by refining all transitions labelled by the same action one by one by disjoint copies of the corresponding refinement net.

4.13 Proposition

Let (N, M_o) be a marked net, $r_1, r_2 \in T_N$, $r_1 \neq r_2$, and let R_1, R_2 be refinement nets. Then $N_1 = ((N, M_o)[R_1/r_1])[R_2/r_2]$ is isomorphic to $N_2 = ((N, M_o)[R_2/r_2])[R_1/r_1]$.

Proof Straightforward. ■

We now show that, for any refinement $N[R/r]$, there exists a canonical vicinity respecting net morphism from $N[R/r]$ to N .

4.14 Proposition

Let (N, M_o) be a marked net, let $r \in T_N$, let R be a refinement net.

Then $f : N[R/r] \rightarrow N$ with

$$f(x) = \begin{cases} x & \text{iff } x \in (S_N - (\bullet r \cup r^\bullet)) \cup (T_N - \{r\}), \\ r & \text{iff } x \in (S_R - ({}^\circ R \cup R^\circ)) \cup T_R, \\ s_N & \text{iff } x = (s_N, s_R) \in \text{Int} \end{cases}$$


is a vicinity respecting morphism, in particular a quotient.

Proof Straightforward ■


We have shown that our construction may be understood in terms of vicinity respecting quotients. However, one could now pose the converse question. May any vicinity respecting quotient which refines only transitions, that is never maps a transition to a place, be generated by our construction? The answer is no, as shown in the following example. However, we would not consider the morphism in this example as a sensible transition refinement.

4.15 Example

Consider



and



The broken lines in N_2 indicate a quotient from N_2 to N_1 which is vicinity respecting and maps no transition to a place. However, we would not like to consider this as a transition refinement. To execute both transitions corresponding to r , an intermediate occurrence of u is necessary. N_2 may not be generated as a refinement of N_1 with our construction.

An interesting problem is to find a further restriction to obtain a class of net morphisms characterising refinement.

Finally, we would like to show that the construction for refinement of transitions we have presented is consistent with the refinement operation on event

structures. This would mean in particular, that this construction for nets indeed preserves precisely the conflict- and causality structure. We will show this for the special case of *occurrence nets*, nets with acyclic flow relation and only forward branched places. These nets correspond directly to prime event structures as defined in Section 1. As refinement nets, we will consider special (finite) occurrence nets, called *causal nets*, with only unbranched places. Causal nets correspond to conflict-free prime event structures. This yields precisely the class of refinements which we have considered in Section 1.

4.16 Definition

- (i) A net structure N is an *occurrence net* iff
 - the transitive closure of F is irreflexive,
 - $\forall s \in S_N : |\bullet s| \leq 1$,
 - $\#_N$ is irreflexive, where for $x, y \in S \cup T$,
 $x \#_N y \iff \exists t, t' \in T_N$ with $t \neq t', t^\bullet \cap t'^\bullet \neq \emptyset, t F^* x$ and $t' F^* y$,
 - $\forall t \in T_N : \{t' \in T_N \mid t' F^* t\}$ finite (axiom of finite causes).
- (ii) A net structure N is a *causal net* iff N is an occurrence net and $\forall s \in S_N : |s^\bullet| \leq 1$.

4.17 Definition Let N be an occurrence net.

The (*prime*) *event structure* of N , $Ev(N)$, is defined as
 $Ev(N) := (T_N, F^* \upharpoonright T_N, \#_N, id_{T_N})$.

It is straightforward to verify that $Ev(N)$ is indeed a prime event structure [142].

Using these notions, we may now show the consistency of transition refinement in this class of nets with prime event structure refinement as defined in Section 1.

4.18 Theorem

Let N be an occurrence net, let $r \in T_N$; let R be a finite causal net.
 Then $Ev(N[R/r]) \cong ref(Ev(N))$ where

$$\begin{aligned} ref(r) &:= Ev(R), \\ ref(t) &:= (\{t\}, \{(t, t)\}, \emptyset, \{(t, t)\}) \text{ for } t \neq r \\ &\quad \text{(identical refinement).} \end{aligned}$$

Proof Omitted. ■

More general consistency results, by unfolding marked nets or associating configuration structures with marked nets and relating with our refinement notion in Section 3, have to be left for further research.

Related work

In this chapter we defined a compositional refinement operator on three kinds of event structures and on Petri nets. Our operator on nets can be regarded as a generalisation of the refinement operators of VALETTE [173], SUZUKI & MURATA [164], BEST, DEVILLERS, KIEHN & POMELLO [33] and VOGLER [175] (although we use a less general kind of nets), and we have compared it with the notions of net morphism (REISIG [157]) and vicinity respecting quotients (DESEL & MERCERON [70]). The operator on *prime* event structures was introduced in VAN GLABBEEK & GOLTZ [84]. It has been defined on sets of pomsets—a linear time variant of the model of prime event structures—in GISCHER [76] and on process graphs modelling only sequential processes in VAN GLABBEEK & WEIJLAND [95] (Section 6 of the previous chapter). The operator is also defined on *causal trees* [53, 54] and on yet another kind of event structures in DARONDEAU & DEGANO [55, 57]. Operational definitions of action refinement can be found in DEGANO & GORRIERI [62], RENSINK [158] and BUSI, VAN GLABBEEK & GORRIERI [45].

In principle there are two ways to treat “syntactic” action refinement in system description languages like CCS. One of them is to use the CCS-actions for modelling the refinable actions of this thesis. In the absence of communication (or synchronisation) refinement can simply be defined as syntactic substitution of an action by a process expression. This approach has been taken in ACETO & HENNESSY [5] and NIELSEN, ENGBERG & LARSEN [141], and has also been mentioned in CASTELLANO, DE MICHELIS & POMELLO [46]. In the presence of communication defining such a refinement operator is much more difficult. A first proposal, for the simple case of an operator only splitting actions in two parts to be executed sequentially, can be found in VAN GLABBEEK & VAANDRAGER [91]. A more general approach is undertaken in ACETO & HENNESSY [6] and ACETO [3].

An alternative is to use the actions of CCS for modelling “atomic” or instantaneous actions that cannot be refined, and representing our refinable actions by means of variables or *parameters*. This approach requires a general sequential composition operator and has been carried out in BERGSTRA & TUCKER [32] in the setting of ACP. In particular [32] shows that there is no problem in defining a refinement operator while working in interleaving semantics: atomic actions a, b cannot be refined, so the equation $a \parallel b = a; b + b; a$ is harmless; parameters x, y can be refined, but there is no equation $x \parallel y = x; y + y; x$. Of course the refinement operator, ordinary substitution, is defined in the language (that still contains all information about causal dependence) and not in the associated interleaving model (which would be impossible according to Example 0.4).

A completely different approach is taken in GORRIERI, MARCHETTI & MONTANARI [100] and BOUDOL [37], see also GORRIERI [98] and DEGANO & GORRIERI [61]. There all actions are assumed to be “atomic”, and this property

should be preserved if they are refined. In [37] even two kinds of atomicity are proposed, corresponding with two kinds of refinement. In [100, 98, 61] this kind of refinement is carried out in an interleaving based model, as mentioned in the introduction.

Refinement in more concrete programming languages is treated in GRIBOMONT [103].

It is often argued that a concurrent system should not be represented just by a Petri net or an event structure, but rather by an equivalence class of such objects. Action refinement is only well-defined on a quotient domain induced by a semantic equivalence if this equivalence is a congruence for refinement, i.e. $P \equiv Q \implies \text{ref}(P) \equiv \text{ref}(Q)$. (Actually, we like to regard action refinement as a binary operator, rather than a family of unary operators. Doing so means that being a congruence for action refinement is an even stronger property, namely $P \equiv Q \wedge \forall a \in \text{Act}(\text{ref}(a) \equiv \text{ref}'(a)) \implies \text{ref}(P) \equiv \text{ref}'(Q)$.) The search for suitable equivalences will be the topic of the remaining chapters of this thesis.

Chapter V

Partial Order Semantics for Refinement of Actions

neither necessary nor always sufficient
but appropriate when used with care

Rob van Glabbeek & Ursula Goltz

Note: This chapter appeared originally in Bulletin of the EATCS 38, pp. 154–163. Here it serves as an informal summary of the remaining two chapters of this thesis. It uses Petri nets rather than event structures and contains no technicalities like definitions and proofs. Instead more attention has been paid to the examples.

Originally this chapter was written in continuation of a series of papers in the Bulletin of the EATCS about the relative merits of partial order semantics and interleaving semantics, starting with CASTELLANO, DE MICHELIS & POMELLO [46]. That paper pointed out a significant advantage of partial order semantics, by formulating a desirable property of semantic equivalences that is not met by interleaving equivalences. This property is *preservation under refinement of actions*. A semantic equivalence is *preserved under action refinement* if two equivalent processes remain equivalent after replacing all occurrences of an action a by a more complicated process $r(a)$. For example, $r(a)$ may be a sequence of two actions a_1 and a_2 . This property may be desirable in applications where concurrent systems are modelled at different levels of abstraction, and where the actions on an abstract level turn out to represent complex processes on a more concrete level. Therefore for example PRATT [154] and LAMPORT [125] already advocate the use of semantic equivalences that are not based on action atomicity.

CASTELLANO, DE MICHELIS & POMELLO [46] showed by means of a simple example that none of the interleaving equivalences - not even bisimulation - is preserved under action refinement. Furthermore they claim that ‘on the other hand, the approaches based on partial order are not constrained to the assumption of atomicity’. Indeed, they give a proof that “linear time” partial order semantics, where a system is identified with the set of its possible (partially ordered) runs, is preserved by refinement. They conclude that ‘interleaving

semantics is adequate only if the abstraction level at which the atomic actions are defined is fixed. Otherwise, partial order semantics should be considered'.


In this chapter we would like to point out that this conclusion is not so obvious. In particular we will argue

- that there are several equivalences based on partial orders which are *not* preserved by refinement (namely when taking the choice structure of systems into account);
- that nevertheless a “branching time” partial order equivalence can be found that is preserved under refinement;
- but that, in order to achieve preservation under refinement it is not necessary to employ partial order semantics: there exist equivalences that abstract from the causal structure of concurrent systems and are still preserved under refinement.

In interleaving semantics, the possible runs of a system are represented as sequences of action occurrences, modelling parallelism by arbitrary interleaving of actions. The example of [46] consisted of the two systems M and N which may not be distinguished in this kind of semantics:

$$\begin{aligned} M &= a \parallel b && \text{(two actions } a \text{ and } b, \text{ executed independently);} \\ N &= a; b + b; a && \text{(either the sequence } ab \text{ or the sequence } ba \text{ is executed).} \end{aligned}$$

They have the following Petri net representations (labelling transitions by action names):



It was shown that after refining a into the sequential composition of a_1 and a_2 , thereby obtaining the systems

$$M' = (a_1; a_2) \parallel b \text{ and } N' = (a_1; a_2); b + b; (a_1; a_2),$$

M' can perform the sequence of actions a_1ba_2 , while N' cannot do this. Hence M' and N' are not equivalent in interleaving semantics.


A first attempt to capture parallelism more precisely is made by so called *step semantics*. Here it is specified that in a run of a parallel system several independent actions may occur together in one *step*. We can think of a system having a global clock where at each clock tick several actions occur simultaneously. This view is taken in calculi like SCCS [136], CIRCAL [137] and MEIJE [10]. Step semantics also have been given to CCS in [59] and to TCSP in [165].

It is easy to see that the two systems M and N considered above are already distinguished in step semantics: In M it is possible to execute the step $\{a, b\}$ whereas in N it is not. So the example in [46] is not well chosen to advocate partial order semantics; already step semantics would be sufficient in this case. Therefore, we will now give a slightly more elaborate example. Consider the following two systems:

$$P = (a; b) \parallel c,$$

$$Q = a; (b \parallel c) + (a \parallel c); b.$$

In both of these systems the actions a, b and c are executed, and b occurs after completion of a . However, in P the c action occurs independently of both a and b whereas in Q c may only occur either “causally behind” a or “causally before” b . P and Q may be represented by the following Petri nets (using a construction explained for instance in [91] for implementing $+$).



P and Q are identified when considering their possible sequences of steps. Both of them take into account the five possibilities for c : occurring before a , simultaneous with a , between a and b , simultaneous with b , or after b . However, after substituting $(c_1; c_2)$ for c only the first system can perform the sequence of actions c_1abc_2 . Thus also this semantics is not preserved under refinement.

On the other hand, P and Q can be distinguished by considering the *partial orders of action occurrences* they allow (the order indicating causal dependence).

$$\begin{array}{ccc} a & \rightarrow & b \\ c & & \end{array} \quad \begin{array}{l} (a \text{ followed by } b \text{ and} \\ \text{independently } c) \end{array}$$

is a computation of P but not of Q . In [46] it was shown that partial order semantics—when identifying a system with its set of possible (partially ordered) runs—is preserved under action refinement.

However, when taking the choice structure of systems into account, the situation becomes less obvious.

Before discussing the problem in detail, we would like to give an overview, by classifying the equivalences being currently investigated (without claiming completeness). They may be positioned in a two dimensional diagram as shown below, distinguishing them firstly with respect to the preserved level of detail in runs of systems (as discussed above) and secondly with respect to the preserved level of detail of the choice structure of systems (we do not consider abstraction from internal actions here).

runs	sequences of actions	sequences of steps	partial orders
branching structure			
paths	interleaving trace equivalence	step trace equivalence	pomset trace equivalence
: e.g. testing			
bisimulation	interleaving bisimulation equivalence	step bisimulation equivalence	e.g. pomset bisimulation equivalence

In trace semantics (“linear time” semantics), a system is fully determined by its set of possible runs, thereby completely neglecting the branching structure.

On the other end, bisimulation semantics preserve the information where two different courses of action diverge (although branching of identical courses of action is still neglected). In between there are several “decorated trace semantics”, where part of the branching structure is taken into account. Mostly these are motivated by the observable behaviour of processes, according to some testing scenario (see Chapter I).


Up to now we have only considered the trace equivalences in the upper row of the diagram. We recalled from [46] that pomset trace equivalence is preserved under action refinement, while interleaving trace equivalence is not. Moreover we have shown that also step trace equivalence is not preserved under refinement. Next we will try to establish similar results for the corresponding branching time equivalences and for the testing equivalences in between.

In interleaving semantics this generalisation is quite simple. As observed in [46], the systems M and N are identified even in interleaving bisimulation semantics while the refined systems M' and N' are not even identified in interleaving trace semantics. So there is one single example showing that neither interleaving bisimulation equivalence nor interleaving trace equivalence is preserved under refinement. As a consequence, also none of the decorated trace equivalences based on interleaving, which are more discriminating than interleaving trace equivalence, but less discriminating than interleaving bisimulation equivalence, is preserved under refinement; in each of the decorated trace semantics based on interleaving, M and N are identified, while M' and N' are distinguished.

Our example against step trace equivalence however cannot be used to show that also step bisimulation equivalence is not preserved under refinement; the systems P and Q happen to be different in step bisimulation semantics already: after performing an a -action the system P is always able to continue with a b -action, whereas Q can perform an a -action and reach a state where it is not possible to continue with b . Nevertheless, the following example shows that also step bisimulation semantics is not preserved under refinement. Consider the two systems M and L which may not be distinguished in step bisimulation semantics:

$$\begin{aligned} M &= a \parallel b && \text{(two actions } a \text{ and } b, \text{ executed independently);} \\ L &= a \parallel b + a; b && \text{(either } a \text{ and } b \text{ are executed independently or} \\ &&& \text{the sequence } ab \text{ is executed).} \end{aligned}$$

They have the following Petri net representations:



The systems $M' = (a_1; a_2) \parallel b$ and $L' = (a_1; a_2) \parallel b + (a_1; a_2); b$ which are obtained by substituting $a_1; a_2$ for a are no longer step bisimulation equivalent; only L' can perform a_1 , and reach a state where it is not possible to continue with b .

Hence, neither step trace nor step bisimulation equivalence is preserved under refinement. However, M' and L' happen to be step trace equivalent, so none of the previous two examples is adequate for both equivalences. In order to tackle the whole range of equivalences included between step trace and step bisimulation equivalence we need yet another example, which simultaneously shows that both step trace and step bisimulation equivalence are not preserved under refinement. Consider the systems


$$Q = a; (b \parallel c) + (a \parallel c); b \quad \text{and} \\ R = Q + P = a; (b \parallel c) + (a \parallel c); b + (a; b) \parallel c.$$

The Petri net associated to Q has been shown before, and the net for R is drawn in [91], where it was also pointed out that Q and R are step bisimulation equivalent. However, after refining c into $c_1; c_2$ the two systems are not even interleaving trace equivalent; only the second system can perform the sequence of actions c_1abc_2 . As a consequence, none of the decorated trace equivalences based on steps, such as the step failure semantics of [165], is preserved under refinement.

A rather straightforward combination of the ideas of bisimulation and of capturing causal dependencies by partial orders has been proposed in BOUDOL & CASTELLANI [39]. They suggest to consider transition systems as for the usual interleaving bisimulation, but to label the arcs in these transition systems by *pomsets* (partially ordered multisets of action occurrences) instead of single actions. However, it turns out that the obtained equivalence, usually called

pomset bisimulation, is not preserved by refinement of actions.

Consider the two systems K and L below.



In both systems either a and b are executed independently or the sequence ab is executed. However, in L the choice between these two options is made at the beginning, while in K this choice can be postponed until the execution of a has been completed.

The system K can behave as follows:

- it performs the single action a and the remaining behaviour is $b+b$, which is identified with b ;
- it performs the single action b and the remaining behaviour is a ;
- it performs the pomset $\begin{smallmatrix} a \\ b \end{smallmatrix}$ (a and b executed independently) and there is no remaining behaviour;
- or it performs the pomset $a \rightarrow b$ (a followed by b) and again there is no remaining behaviour.

The behaviour of L can be described in exactly the same way and for this reason the two systems are pomset bisimulation equivalent.

Now let us imagine that a is refined into $a_1; a_2$. The systems K' and L' which are obtained in this way can be distinguished in pomset bisimulation semantics, and even in interleaving bisimulation: only L' can refuse to do a b -action after execution of a_1 .

Hence pomset bisimulation semantics is not preserved under refinement of ac-

tions. Another example for this are the two terms

$$a; (b + c) + (a \parallel b) \text{ and } a; (b + c) + (a \parallel b) + (a; b)$$

(again refining a into $a_1; a_2$). However the example given before can also be used to show that even the notion of generalised pomset bisimulation, as discussed in VAN GLABBEK & VAANDRAGER [91], is not preserved under refinement. Of course we cannot find an example tackling the whole range of equivalences included between pomset trace and pomset bisimulation semantics, since we already observed that pomset trace equivalence is preserved under refinement. However, the systems K' and L' can already be distinguished in interleaving failure semantics, as employed in [42, 64]. Thus no equivalence that is at least as discriminating as interleaving failure equivalence but less discriminating than pomset bisimulation equivalence can be preserved under refinement.

The interplay of equivalence notions and refinement of actions as discussed up to now is investigated in detail in Chapter VI. There all the equivalence notions and examples presented so far are given formally in the framework of event structures; refinement of actions is performed by replacing actions by non-empty pomsets. That paper concludes by showing that another “partial order bisimulation” is indeed preserved by refinement. In order to avoid technical details, we just outline these results here.

After we realised that pomset bisimulation is not preserved by refinement, another equivalence was considered, hoping that it would solve the problem (see e.g. DEVILLERS [71]). This equivalence had been considered before under the name NMS partial ordering equivalence in DEGANO, DE NICOLA & MONTANARI [58]. The main idea is to bisimulate transition systems where the states are labelled by their (partially ordered) histories. In the next chapter it is shown that this equivalence is indeed preserved by refinement when we restrict ourselves to systems *without autoconcurrency*, that is to systems which do not allow concurrent occurrences of the same action like in $a \parallel a$. However, for systems with autoconcurrency it turns out that NMS po equivalence is not preserved by refinement. Even more, it does not even respect pomset bisimulation equivalence in this case. The example showing both these facts was suggested to us by Alex Rabinovich who used it to show that this equivalence is not a congruence with respect to a TCSP-like parallel composition. To obtain a congruence, a stronger version of NMS po equivalence was suggested in RABINOVICH & TRAKHTENBROT [155] In the next chapter it is shown that *this* “partial order bisimulation equivalence” is indeed always preserved by refinement.

So we have shown that it is *not automatically sufficient* to move to partial order semantics for refinement of actions. When considering the choice structure, this has to be done with care. In the remaining part of this note, we argue that on the other hand it is *not even necessary* to move to partial orders (as one may conclude from [46]).

A branching time semantics lying strictly between step semantics and partial order semantics has been proposed in VAN GLABBEK & VAANDRAGER [91]. This *ST-bisimulation semantics* is based on the idea that actions have a duration, and may overlap in time. Contrary to step semantics, it recognises the possibility that, in $P = (a; b) \parallel c$, action c may have an overlap with both a and b , while b can only occur after completion of a . However when in a run of a system an action b happens after completion of a , it is not taken into account whether or not there is a causal link between the two actions.

Compare for instance the systems M and K that have been presented before. Both systems perform an a -action and a b -action. In M these actions are always independent, whereas in K it is possible to perform a b -action which causally depends on a : so M and K are distinguished in partial order semantics. However, in ST-bisimulation semantics the only execution of K which is not possible in M (first a and then the b which is causally dependent on this a) can not be distinguished from another execution of K (and of M), namely: first a and then the b which is independent of this a . In K , the choice between both runs is only made after completion of a , and in that state the remaining part of both executions is the same: just b . Hence M and K are identified in ST-bisimulation semantics.

So ST-bisimulation equivalence abstracts from the causal structure of concurrent systems. Nevertheless it is preserved under refinement, as will be shown in Chapter VII. A similar result can be proved for linear time semantics as well. A variant of this can be found in NIELSEN, ENGBERG & LARSEN [141]. Furthermore a variant of failure semantics, based on the same ideas that underly ST-bisimulation semantics has been proposed in VOGLER [175]. There it is proven that also this equivalence respects refinement.

This shows that indeed partial order semantics (in the strong sense¹) are not necessary for the type of refinement we have considered. Nevertheless, we need partial order bisimulation semantics when it is required to model the interplay of causality and branching in full detail.

We hope that this note, and the formal versions of it in the next chapters, help to clarify the relationship between various equivalences being currently considered. However, we do not intend to advocate any particular type of equivalence here. We just want to illustrate that the appropriate equivalence notion has to be chosen carefully with regard to the considered questions.

¹We refer here to semantics that use partial orders for representing causal dependence. Interleaving semantics can also be described by means of partial orders (that happen to be total), and ST-bisimulation semantics, although its original definition did not manifestate partial orders, can be described by means of so-called *interval orders* [176, 178]. As interval orders are truly partial, WALTER VOGLER [176, 178] argues that ST-bisimulation semantics counts as a partial order semantics. His conclusion that partial order semantics is necessary for refinement of actions is, apart from the chosen terminology, consistent with ours.

Chapter VI

Equivalence Notions for Concurrent Systems and Refinement of Actions

Rob van Glabbeek & Ursula Goltz

In this chapter we investigate equivalence notions for concurrent systems. We consider “linear time” approaches where the system behaviour is characterised as the set of possible runs as well as “branching time” approaches where the conflict structure of systems is taken into account. We show that the usual interleaving equivalences, and also the equivalences based on *steps* (multisets of concurrently executed actions) are not preserved by refinement of actions. We prove that linear time partial order semantics, where causality in runs is explicit, is invariant under refinement. Finally, we consider various bisimulation equivalences based on partial orders and show that the strongest one of them is preserved by refinement whereas the others are not.

Notes: This chapter is based on our paper [84]. However, the introduction of prime event structures and the refinement operator as well as the proof of its compositionality has been omitted since this is covered by Section 1 of Chapter IV already. Furthermore our results are reformulated in such a way that they are not only valid for prime event structures but also for flow event structures, as presented in Section 2 of Chapter IV. On flow event structures, infinite refinements and refinements with conflicts are no longer excluded, which strengthens our refinement theorems.

Contents

Introduction	224
1. Interleaving semantics	225
2. Step semantics	227
3. “Linear time” partial order semantics	230
4. “Branching time” partial order semantics	231
4.1. Pomset bisimulation equivalence	232
4.2. History preserving bisimulation	233
Conclusion	238

Introduction

A large body of research is devoted to equivalence notions for concurrent systems. Most of the equivalence notions currently being considered are based on a semantics where concurrency is modelled by arbitrary interleaving of atomic actions. In PRATT [154] and in CASTELLANO, DE MICHELIS & POMELLO [46] it is pointed out that this approach has a severe drawback. It leads to complications when changing the level of atomicity of events; "... we would like a theory of processes to be just as usable for events having a duration or structure, where a single event can be atomic from one point of view and compound from another" ([154]). In [46], an example is given, showing that the usual interleaving equivalence is not invariant under refinement of actions when this is simply modelled by textual replacement. Both [154] and [46] claim that modelling concurrency by expressing causal dependencies explicitly using partial orders could help to solve this problem. However, the two systems considered in [46] can already be distinguished by considering interleavings of "steps" (multisets of concurrently executable actions). So their example does not show that it is indeed necessary to consider partially ordered executions. Furthermore, their proof of the claim that partial order equivalence is preserved by refinement is only valid for "linear time" partial order semantics, where the set of all possible executions of a system is considered, without taking into account where conflicts are resolved. This is also the model considered by Pratt.

In this chapter, we will consider various equivalence notions based on steps and on partial orders. We will discuss linear time semantics, but we will also take the conflict structure of systems into account by considering various forms of bisimulation ("branching time" semantics). We will show that the known equivalences based on steps are not invariant under action refinement. We will rephrase in our framework the proof of [46], showing that linear time partial order semantics is indeed robust against changing the level of atomicity. Then we consider several equivalence notions based on branching time partial order semantics. We give examples, showing that pomset bisimulation equivalence of BOUDOL & CASTELLANI [39] and also the NMS partial ordering equivalence suggested in DEGANO, DE NICOLA & MONTANARI [58], are not preserved by refinement of actions. An equivalence notion for Petri nets which coincides with the notion of NMS partial ordering equivalence was suggested in DEVILLERS [71] where the refinement problem has also been discussed. We also show that NMS partial ordering equivalence does not imply pomset bisimulation and vice versa; these notions are incomparable. Finally we show that a finer equivalence notion, proposed in RABINOVICH & TRAKHTENBROT [155] under the name BS-bisimulation, is indeed preserved by refinement. This equivalence does respect pomset bisimulation.

We do not intend to advocate any particular equivalence notion here, the purpose of this investigation is to find out about the consequences of the different

approaches. There will certainly be a tradeoff between simplicity and distinguishing power. We just want to illustrate that the appropriate notion has to be chosen carefully with regard to the questions considered.

1 Interleaving semantics

In this paper, concurrent systems are represented by event structures. It is written in such a way that the text applies both to prime event structures as to flow event structures (but in case of prime event structures $\leq_X$ should be read as $\leq$). The reader is referred to Section 1—for prime event structures—or Section 2—for flow event structures—of Chapter IV, for an introduction to event structures and action refinement. There it is also explained how configurations model the states of a concurrent system.

We may now ask which actions may occur in a configuration and which configuration is then obtained.

Definition Let $\mathcal{E}$ be an event structure.

1. $X \longrightarrow_{\mathcal{E}} X'$ iff $X, X' \in \text{Conf}(\mathcal{E})$ and $X \subseteq X'$.
2. $X \xrightarrow{a} X'$ iff $a \in \text{Act}$, $X \longrightarrow_{\mathcal{E}} X'$ and $X' \setminus X = \{e\}$ with $l(e) = a$.

Note that $X \longrightarrow_{\mathcal{E}} X'$ implies that $\mathcal{E} \upharpoonright (X' \setminus X)$ is finite and conflict-free.

Here $X \xrightarrow{a} X'$ says that if $\mathcal{E}$ is in the state represented by X , then it may perform an action a and reach a state represented by X' . Likewise, $X \longrightarrow_{\mathcal{E}} X'$ says that $\mathcal{E}$ may evolve from X to X' .

Considering transitions $X \xrightarrow{a} X'$ only, one can define the usual *interleaving semantics*. The simplest form is that of comparing just the possible sequences of action occurrences.

Definition

$w = a_1 \cdots a_n \in \text{Act}^*$ is a (*sequential*) *trace* of an event structure $\mathcal{E}$ iff there exist configurations $X_0, \dots, X_n$ of $\mathcal{E}$ such that $X_0 = \emptyset$ and $X_{i-1} \xrightarrow{a_i} X_i$ ($i = 1, \dots, n$).

$\text{SeqTraces}(\mathcal{E})$ denotes the set of all sequential traces of $\mathcal{E}$.

Two event structures $\mathcal{E}, \mathcal{F}$ are called *interleaving trace equivalent* ($\mathcal{E} \approx_{it} \mathcal{F}$) iff $\text{SeqTraces}(\mathcal{E}) = \text{SeqTraces}(\mathcal{F})$.

With the concept of labelled transition systems, we obtain a stronger equivalence notion based on the idea of bisimulation [145, 134]. For example, the systems $a(b + c)$ and $ab + ac$ have the same traces but are distinguished by bisimulation equivalence.

Definition Let $\mathcal{E}, \mathcal{F}$ be event structures.

A relation $R \subseteq \text{Conf}(\mathcal{E}) \times \text{Conf}(\mathcal{F})$ is called an *interleaving bisimulation between $\mathcal{E}$ and $\mathcal{F}$* iff $(\emptyset, \emptyset) \in R$ and if $(X, Y) \in R$ then

- $X \xrightarrow{a} X' \Rightarrow \exists Y' \text{ with } Y \xrightarrow{a} Y' \text{ and } (X', Y') \in R,$
- $X \xrightarrow{a} Y' \Rightarrow \exists X' \text{ with } X \xrightarrow{a} X' \text{ and } (X', Y') \in R.$

$\mathcal{E}$ and $\mathcal{F}$ are *interleaving bisimulation equivalent* ($\mathcal{E} \approx_{ib} \mathcal{F}$) iff there exists an interleaving bisimulation between $\mathcal{E}$ and $\mathcal{F}$.

Clearly, $\mathcal{E} \approx_{ib} \mathcal{F}$ implies $\mathcal{E} \approx_{it} \mathcal{F}$.

Example 1.1

We now recall the example of [46], showing that both $\approx_{it}$ and $\approx_{ib}$ are not preserved by refinement.

They considered the two systems $P = a||b$ and $Q = a;b + b;a$, representable by the following event structures.

$$\begin{array}{ccc} \mathcal{E}_P & = & a \quad b \quad , \quad \mathcal{E}_Q = \quad a - \# - b \\ & & \downarrow \qquad \qquad \downarrow \\ & & b \qquad \qquad \quad a \end{array}$$

In all known interleaving semantics, P and Q are considered equivalent, we have $\mathcal{E}_P \approx_{ib} \mathcal{E}_Q$. However, if we allow to refine the action a into the pomset $a_1 \rightarrow a_2$, this gives rise to the two systems

$$\begin{array}{ccc} \mathcal{E}_{P'} & = & a_1 \quad b \quad , \quad \mathcal{E}_{Q'} = \quad a_1 - \# - b \\ & \downarrow & \downarrow \qquad \qquad \downarrow \\ & a_2 & a_2 \qquad \qquad a_1 \\ & & \downarrow \qquad \qquad \downarrow \\ & & b \qquad \qquad \quad a_2 \end{array}$$

and they are not interleaving equivalent; indeed they are not even interleaving trace equivalent: $\mathcal{E}_{P'}$ allows for the sequence $a_1 b a_2$ whereas $\mathcal{E}_{Q'}$ doesn't.

This shows that both interleaving trace equivalence and interleaving bisimulation equivalence are not preserved by action refinement. Even more, the same can be said for all equivalences identifying P and Q and respecting interleaving trace equivalence, e.g. failure equivalence [42], testing equivalence [64].

As an event structure equivalence which is indeed preserved by refinement one could consider event structure isomorphism (cf. Proposition 1.5(iii)/2.7(iii) of Chapter IV). However, the main purpose of introducing an equivalence notion is to abstract from certain details in a system representation. For example, we

would like to express that the processes a and $a + a$ exhibit the same behaviour. Furthermore, we would like to identify processes like

$$(a \parallel (b + c)) + (a \parallel b) + ((a + c) \parallel b) \quad \text{and} \quad (a \parallel (b + c)) + ((a + c) \parallel b)$$

(absorption law, see [39]). This is not possible when using event structure isomorphism.

Hence, in the sequel we will consider various equivalence notions in between these two extremes (interleaving trace equivalence and event structure isomorphism), taking into account the concurrency and the conflict structure (“branching-time” semantics) in more and more detail.

2 Step semantics

A more discriminating view of concurrent systems than that offered by interleaving semantics is obtained by modelling concurrency as either arbitrary interleaving or simultaneous execution. This view is taken in calculi like SCCS [136], CIRCAL [133] and MEIJE [10]. In TAUBNER & VOGLER [165], this idea is applied to give a non-interleaving semantics to theoretical CSP, called *step failure semantics*. The word *step* originates from Petri net theory where it denotes a set (or multiset) of concurrently executable transitions. Recently, a step semantics for CCS has been defined [59], inspired by [10]. Step semantics give a more precise account of concurrency than interleaving semantics, e.g. the systems $a \parallel b$ and $a; b + b; a$ are distinguished. This means that the example given in [46] constitutes an argument against interleaving semantics but not against step semantics. We will formalise some step equivalence notions and then discuss examples which show that even these equivalences are not preserved by refinement.

Step semantics are defined by generalising the single action transitions $X \xrightarrow{a} X'$ from Section 1 to transitions of the form $X \xrightarrow{A} X'$ where A is a multiset over Act , representing actions occurring concurrently. In particular, we allow actions to occur concurrently with themselves (“autoconcurrency”). Using this new kind of transitions, *step trace equivalence* and *step bisimulation equivalence* are straightforward generalisations of the corresponding interleaving equivalences, see e.g. POMELLO [153].

Definition Let $\mathcal{E}$ be an event structure.

$X \xrightarrow{A} X'$ iff $A \in \mathcal{N}^{Act}$ (A is a multiset over Act), $X \longrightarrow_{\mathcal{E}} X'$ and $X' \setminus X = G$ such that $\forall e, e' \in G \ e \text{ co } e'$ and $l(G) = A$ where $l(G)(a) = |\{e \in G \mid l(e) = a\}|$.

Definition

$W = A_1 \cdots A_n$ where $A_i \in N^{Act}$ ($i = 1, \dots, n$) is a *step trace* of an event structure $\mathcal{E}$ iff there exist configurations $X_0, \dots, X_n$ of $\mathcal{E}$ such that

$X_0 = \emptyset$ and $X_{i-1} \xrightarrow{A_i} X_i$ ($i = 1, \dots, n$).

$StepTraces(\mathcal{E})$ denotes the set of all step traces of an event structure $\mathcal{E}$.

Two event structures $\mathcal{E}, \mathcal{F}$ are called *step trace equivalent* ($\mathcal{E} \approx_{st} \mathcal{F}$) iff $StepTraces(\mathcal{E}) = StepTraces(\mathcal{F})$.

Definition Let $\mathcal{E}, \mathcal{F}$ be event structures.

A relation $R \subseteq Conf(\mathcal{E}) \times Conf(\mathcal{F})$ is called a *step bisimulation between $\mathcal{E}$ and $\mathcal{F}$* iff $(\emptyset, \emptyset) \in R$ and if $(X, Y) \in R$ then

– $X \xrightarrow{A} X' \implies \exists Y' \text{ with } Y \xrightarrow{A} Y' \text{ and } (X, Y') \in R,$

– $Y \xrightarrow{A} Y' \implies \exists X' \text{ with } X \xrightarrow{A} X' \text{ and } (X', Y') \in R.$

$\mathcal{E}$ and $\mathcal{F}$ are *step bisimulation equivalent* ($\mathcal{E} \approx_{sb} \mathcal{F}$) iff there exists a step bisimulation between $\mathcal{E}$ and $\mathcal{F}$.

As for interleaving semantics, $\mathcal{E} \approx_{sb} \mathcal{F}$ implies $\mathcal{E} \approx_{st} \mathcal{F}$. Moreover (as far as we know) all other interesting step equivalence notions are positioned somewhere in between (recall that we do not consider abstraction from internal actions).

Considering the two systems $P = a||b$ and $Q = a;b + b;a$ from [46], represented as event structures $\mathcal{E}_P$ and $\mathcal{E}_Q$ in Example 1.1, we find that $\mathcal{E}_P$ and $\mathcal{E}_Q$ are not equivalent in step semantics. The step $\{a, b\}$ is possible in $\mathcal{E}_P$ but not in $\mathcal{E}_Q$. So the example in [46] is not adequate for step semantics. Here we give examples showing that both $\approx_{st}$ and $\approx_{sb}$ are not invariant under refinement of actions, as well as all equivalences included between them, e.g. step failure equivalence.

The following example shows that step trace semantics is in general not invariant under refinement.

Example 2.1

We consider the two systems $P = (a;b)||c$ and $Q = (a||c);b + a;(b||c)$:

$$\mathcal{E}_P = \begin{array}{c} a \\ \downarrow \\ b \end{array} \quad c \quad \text{and} \quad \mathcal{E}_Q = \begin{array}{c} a \quad c \\ \searrow \quad \swarrow \\ b \end{array} + \begin{array}{c} a \\ \swarrow \quad \searrow \\ b \quad c \end{array}.$$

The $+$ -sign in the second system is supposed to indicate that this

system either behaves like $\begin{array}{c} a \quad c \\ \searrow \quad \swarrow \\ b \end{array}$ or like $\begin{array}{c} a \\ \swarrow \quad \searrow \\ b \quad c \end{array}$, that is

either performs a and c in parallel and then b , or first a and then b and c in parallel. The $+$ -sign may easily be “implemented” by

indicating that all events in the first component are in conflict with all events in the second component and vice versa.

These two systems are step trace equivalent. However, when refining c into $c_1 \rightarrow c_2$, the resulting systems

$$\mathcal{E}_{P'} = \begin{array}{cc} a & c_1 \\ \downarrow & \downarrow \\ b & c_2 \end{array} \quad \text{and} \quad \mathcal{E}_{Q'} = \begin{array}{c} c_1 \\ \downarrow \\ c_2 \\ \swarrow \searrow \\ a \quad b \end{array} + \begin{array}{cc} & a \\ & \swarrow \searrow \\ b & c_1 \\ & \downarrow \\ & c_2 \end{array}$$

are not step trace equivalent (not even interleaving trace equivalent).

This example shows that $\approx_{st}$ is not preserved by refinement. However, the example is not adequate for step bisimulation equivalence since P and Q are not step bisimulation equivalent (after performing a , the b is always possible in P but not always in Q). The next example shows that also $\approx_{sb}$ is not preserved by refinement.

Example 2.2

Consider $P = a \parallel b$ and $Q = (a \parallel b) + a;b$,

$$\mathcal{E}_P = a \quad b, \quad \mathcal{E}_Q = \begin{array}{c} a \text{---} \# \text{---} a \text{---} \# \text{---} b \\ \downarrow \\ b \end{array}.$$

It is easy to verify that $\mathcal{E}_P \approx_{sb} \mathcal{E}_Q$. However, refining a into $a_1 \rightarrow a_2$ yields

$$\mathcal{E}_{P'} = \begin{array}{cc} a_1 & \\ \downarrow & b \\ a_2 & \end{array}, \quad \mathcal{E}_{Q'} = \begin{array}{cc} a_1 \text{---} \# \text{---} a_1 \text{---} \# \text{---} b \\ \downarrow & \downarrow \\ a_2 & a_2 \\ & \downarrow \\ & b \end{array}$$

After the step $\{a_1\}$, the step $\{b\}$ is always possible in $\mathcal{E}_{P'}$. However, in $\mathcal{E}_{Q'}$ it may be the case that the step $\{b\}$ is impossible after executing a_1 (choosing the branch $a_1 \rightarrow a_2 \rightarrow b$). Hence $\mathcal{E}_{P'}$ and $\mathcal{E}_{Q'}$ are not step bisimulation equivalent (not even interleaving bisimulation equivalent).

However, this example is still not suitable for disqualifying the whole range of equivalence notions included between $\approx_{st}$ and $\approx_{sb}$, as the example of [46] does in the interleaving case (see Example 1.1), since the refined systems $\mathcal{E}_{P'}$ and $\mathcal{E}_{Q'}$ turn out to be step trace equivalent. A slightly more complicated example may be given, disqualifying all equivalence notions between $\approx_{sb}$ and $\approx_{st}$.

Example 2.3

First consider the following three systems:

$$\mathcal{E}_1 = \begin{array}{c} a \\ \swarrow \searrow \\ b \quad c \end{array}, \quad \mathcal{E}_2 = \begin{array}{c} a \quad c \\ \searrow \swarrow \\ b \end{array}, \quad \mathcal{E}_3 = \begin{array}{c} a \\ \downarrow \\ b \end{array} \quad c.$$

Now we consider the two composed systems

$$\mathcal{E} = \mathcal{E}_1 + \mathcal{E}_2, \quad \mathcal{F} = \mathcal{E}_1 + \mathcal{E}_2 + \mathcal{E}_3.$$

We have $\mathcal{E} \approx_{sb} \mathcal{F}$ [91]. However, when refining c into $c_1 \rightarrow c_2$ only the refinement of $\mathcal{F}$ may perform the sequence of actions $c_1 a b c_2$. The resulting systems $\mathcal{E}'$ and $\mathcal{F}'$ are not even interleaving trace equivalent.

So let $\approx$ be an equivalence included between $\approx_{st}$ and $\approx_{sb}$, then also $\mathcal{E} \approx \mathcal{F}$, but $\mathcal{E}' \not\approx \mathcal{F}'$.

Thus we have shown that all the currently known versions of step equivalence are not preserved by refinement.

3 “Linear time” partial order semantics

In [46] it was claimed that equivalence based on considering partially ordered executions is preserved by refinement. In this section we will make this claim more precise. We will show that this is indeed true when considering the set of all possible executions of systems (*traces*), formalising the proof sketch from [46] in terms of event structures. However, in the next section, we will consider equivalence notions taking account of the timing of choices, based on the idea of bisimulation, and we will show that in this case this claim is not so obvious.

In Chapter IV, we discussed that the possible executions of a system may be represented as isomorphism classes of labelled partial orders (*pomsets*), thus taking full account of the causality relation for event occurrences.

Definition

- (i) Let $X = (X, \leq_X, l_X)$ and $Y = (Y, \leq_Y, l_Y)$ be partial orders which are labelled over Act . X and Y are *isomorphic* ($X \cong Y$) iff there exists a bijection between X and Y respecting the ordering and the labelling. The isomorphism class of a partial order labelled over Act is called a *pomset over Act*.
- (ii) Let $\mathcal{E}$ be an event structure.
 $Pomsets(\mathcal{E}) := \{ [(X, \leq_X, l_{\mathcal{E}} \upharpoonright X)]_{\cong} \mid X \in Conf(\mathcal{E}) \}.$

- (iii) Two event structures $\mathcal{E}$ and $\mathcal{F}$ are *pomset trace equivalent* ($\mathcal{E} \approx_{pt} \mathcal{F}$) if $Pomsets(\mathcal{E}) = Pomsets(\mathcal{F})$.

Clearly, pomset trace equivalence implies step trace equivalence. Example 2.1 shows that pomset trace equivalence is strictly finer than step trace equivalence. On the other hand, pomset trace equivalence and step (or interleaving) bisimulation equivalence are incomparable: $a;(b+c) \approx_{pt} a;b+a;c$ and for $\mathcal{E}_P$ and $\mathcal{E}_Q$ of Example 2.2, $\mathcal{E}_P \approx_{sb} \mathcal{E}_Q$ but $\mathcal{E}_P \not\approx_{pt} \mathcal{E}_Q$.

Theorem Let $\mathcal{E}, \mathcal{F}$ be flow event structures.

Then $\mathcal{E} \approx_{pt} \mathcal{F}$ implies $ref(\mathcal{E}) \approx_{pt} ref(\mathcal{F})$ for any refinement function ref .

Proof

Let $\mathcal{E} \approx_{pt} \mathcal{F}$ and let ref be a refinement function. We have to show

$$Pomsets(ref(\mathcal{E})) = Pomsets(ref(\mathcal{F})).$$

” $\subseteq$ ” Let $u \in Pomsets(ref(\mathcal{E}))$.

Then $u = [(\tilde{X}, \leq_{\tilde{X}}, l_{ref(\mathcal{E})} \upharpoonright \tilde{X})]_{\cong}$ where $\tilde{X} \in Conf(ref(\mathcal{E}))$.

With Proposition 1.7 or 2.8 from Chapter IV, we have that $\tilde{X}$ is a refinement of some configuration X of $\mathcal{E}$. Since $Pomsets(\mathcal{E}) = Pomsets(\mathcal{F})$, there exists $Y \in Conf(\mathcal{F})$ such that $(X, \leq_X, l_{\mathcal{E}} \upharpoonright X)$ and $(Y, \leq_Y, l_{\mathcal{F}} \upharpoonright Y)$ are isomorphic. Since isomorphism preserves labelling, we can refine Y to a configuration $\tilde{Y}$ (by choosing identical refinements for corresponding events) such that

$$(\tilde{X}, \leq_{\tilde{X}}, l_{ref(\mathcal{E})} \upharpoonright \tilde{X}) \cong (\tilde{Y}, \leq_{\tilde{Y}}, l_{ref(\mathcal{F})} \upharpoonright \tilde{Y}),$$

hence $u \in Pomsets(ref(\mathcal{F}))$.

” $\supseteq$ ” by symmetry. ■

4 “Branching time” partial order semantics

In this section, we discuss several suggestions to define equivalence notions based on partial orders and recording where choices are made. We show that most of these fail in general to be preserved by refinement. Finally we show that the last and strongest notion is indeed invariant with respect to refinement.

4.1 Pomset bisimulation equivalence

In BOUDOL & CASTELLANI [39] it was suggested to generalise the idea of bisimulation by considering transitions labelled by pomsets. So we now consider transitions $X \xrightarrow{u} X'$ where u is a pomset over Act .

Definition Let $\mathcal{E}$ be an event structure.

$$X \xrightarrow{u} X' \text{ iff } X \longrightarrow_{\mathcal{E}} X' \text{ and } u \text{ is the isomorphism class of } \mathcal{E} \upharpoonright (X' \setminus X).$$

Definition Let $\mathcal{E}, \mathcal{F}$ be event structures.

A relation $R \subseteq Conf(\mathcal{E}) \times Conf(\mathcal{F})$ is called a *pomset bisimulation between $\mathcal{E}$ and $\mathcal{F}$* iff $(\emptyset, \emptyset) \in R$ and if $(X, Y) \in R$ then

- $X \xrightarrow{u} X' \Rightarrow \exists Y' \text{ with } Y \xrightarrow{u} Y' \text{ and } (X', Y') \in R,$
- $Y \xrightarrow{u} Y' \Rightarrow \exists X' \text{ with } X \xrightarrow{u} X' \text{ and } (X', Y') \in R.$

$\mathcal{E}$ and $\mathcal{F}$ are *pomset bisimulation equivalent* ($\mathcal{E} \approx_{pb} \mathcal{F}$) iff there exists a pomset bisimulation between $\mathcal{E}$ and $\mathcal{F}$.

This equivalence notion is clearly finer than both step bisimulation equivalence and pomset trace equivalence: $\mathcal{E} \approx_{pb} \mathcal{F}$ implies $\mathcal{E} \approx_{sb} \mathcal{F}$ and $\mathcal{E} \approx_{pt} \mathcal{F}$; moreover, the processes $a||b$ and $(a||b) + a;b$ considered in Example 2.2 are *sb*-equivalent but not *pb*-equivalent; $a;(b+c)$ and $a;b+a;c$ are pomset trace equivalent but not *pb*-equivalent.

However, *pb*-equivalence is not preserved by refinement.

Example 4.1

Consider $P = a;(b+c) + (a||b)$ and $Q = a;(b+c) + (a||b) + a;b$. We have $P \approx_{pb} Q$. However, when refining a into $a_1 \rightarrow a_2$ and executing a_1 , we may arrive in a situation in the second system where a_2 and b may only be executed sequentially and where c is excluded. This is not possible in the first system.

In VAN GLABBEEK & VAANDRAGER [91], the pomset bisimulation was criticized for violating “the real combination of causality and branching time”. The criticism is that only the first system of Example 4.1 has the property that any action a that is causally preceding b is also preceding the choice between b and c . Therefore they suggested a generalised pomset bisimulation equivalence, that is finer than pomset bisimulation equivalence, does not identify the two systems of Example 4.1, and still satisfies $a = a + a$ and the absorption law of Section 1.

However, generalised pomset bisimulation equivalence is also not preserved by refinement.

Example 4.2

$$\mathcal{E} = \begin{array}{c} a \\ \downarrow \\ b \# b \end{array} \qquad \mathcal{F} = \begin{array}{c} a \\ \downarrow \\ b \# b \end{array} + \begin{array}{c} a \\ \downarrow \\ b \end{array}$$

These two systems are generalised pomset bisimulation equivalent [91]. However, when refining a into $a_1 \rightarrow a_2$, the resulting systems

$$\mathcal{E}' = \begin{array}{c} a_1 \\ \downarrow \\ a_2 \\ \downarrow \\ b \# b \end{array} \qquad \text{and} \qquad \mathcal{F}' = \begin{array}{c} a_1 \\ \downarrow \\ a_2 \\ \downarrow \\ b \# b \end{array} + \begin{array}{c} a_1 \\ \downarrow \\ a_2 \\ \downarrow \\ b \end{array}$$

are not even interleaving bisimulation equivalent. After the action a_1 the action b is always possible in $\mathcal{E}'$. However in $\mathcal{F}'$ it may be the case that b is impossible after executing a_1 (choosing the branch $a_1 \rightarrow a_2 \rightarrow b$).

4.2 History preserving bisimulation

Another equivalence notion based on the idea of bisimulation with partial orders that might be preserved by refinement was suggested independently by Devillers and Van Glabbeek at the workshop on “Combining Compositionality and Concurrency” [143, 71]. It turned out that this notion coincides with the NMS partial ordering equivalence suggested earlier in DEGANO, DE NICOLA & MONTANARI [58]. We rephrase the definition here in terms of event structures as follows.

Definition Let $\mathcal{E}, \mathcal{F}$ be event structures.

- A relation $R \subseteq \text{Conf}(\mathcal{E}) \times \text{Conf}(\mathcal{F})$ is called a *weak history preserving bisimulation between $\mathcal{E}$ and $\mathcal{F}$* iff $(\emptyset, \emptyset) \in R$ and if $(X, Y) \in R$ then
 - there is an isomorphism between $(X, \leq_X, l_{\mathcal{E}} \upharpoonright X)$ and $(Y, \leq_Y, l_{\mathcal{F}} \upharpoonright Y)$,
 - $X \rightarrow_{\mathcal{E}} X' \Rightarrow \exists Y' \text{ with } Y \rightarrow_{\mathcal{F}} Y' \text{ and } (X', Y') \in R$,
 - $Y \rightarrow_{\mathcal{F}} Y' \Rightarrow \exists X' \text{ with } X \rightarrow_{\mathcal{E}} X' \text{ and } (X', Y') \in R$.
- $\mathcal{E}$ and $\mathcal{F}$ are *weakly history preserving equivalent* ($\mathcal{E} \approx_{wh} \mathcal{F}$) iff there exists a weak history preserving bisimulation between $\mathcal{E}$ and $\mathcal{F}$.

Note that the isomorphism requirement guarantees that the labels of the events in $X' \setminus X$ and $Y' \setminus Y$ correspond as well.

As observed in [71], it is sufficient to consider only those transitions $X \rightarrow_{\mathcal{E}} X'$ (resp. $Y \rightarrow_{\mathcal{F}} Y'$) where $X'(Y')$ is obtained from $X(Y)$ by executing exactly one event.

The two systems considered in Example 4.1 are pomset bisimulation equivalent but not weakly history preserving equivalent. However, *wh*-equivalence is not stronger than pomset bisimulation, as shown by the following example; the two notions are in general incomparable. We will show later that *wh*-equivalence does respect pomset bisimulation for systems without autoconcurrency.

The following example will also show that *wh*-equivalence is in general not preserved by refinement. This example was suggested to us by Rabinovich. He used it for showing that $\approx_{wh}$ is not a congruence with respect to a TCSP-like parallel composition.

Example 4.3


$$\text{Let } \mathcal{E} = \begin{array}{c} a - \# - a \quad a \\ \downarrow \quad \downarrow \\ b - \# - b \end{array} \quad \text{and } \mathcal{F} = \begin{array}{c} a - \# - a \quad a \\ \downarrow \quad \# \quad \downarrow \\ b \quad \quad b \end{array}.$$

It is straightforward to check that $\mathcal{E} \approx_{wh} \mathcal{F}$. However, $\mathcal{E}$ and $\mathcal{F}$ are not pomset bisimulation equivalent. After executing a , it is always possible to execute $a \rightarrow b$ in $\mathcal{E}$, in $\mathcal{F}$ it may be impossible to execute $a \rightarrow b$ after a . When refining a into $a_1 \rightarrow a_2$, the resulting systems are no longer *wh*-equivalent, not even interleaving bisimulation equivalent. This can be proven by providing a formula in Hennessy-Milner logic (Section 1.13 of Chapter I) that is satisfied by the refinement of $\mathcal{F}$, but not by the refinement of $\mathcal{E}$. Such a formula is:

$$a_1 a_2 (b T \wedge a_1 \neg b T).$$

An equivalence respecting both pomset bisimulation and *wh*-equivalence may be considered by extending the definition of pomset bisimulation with the first requirement in the definition of *wh*-bisimulation (the other two requirements in the definition of *wh*-bisimulation are then redundant). However, the following example shows that also this equivalence would not preserve refinement.

Example 4.4



As is quite difficult to check, $\mathcal{E}$ and $\mathcal{F}$ are equivalent according to the equivalence notion proposed above, but after refining a into $a_1 \rightarrow a_2$ they are not even bisimulation equivalent. The formula $a_1 a_1 \neg a_2 \neg b T$ is satisfied by $\mathcal{E}$, but not by $\mathcal{F}$.

We finally define a finer version of history preserving equivalence which will respect pomset bisimulation. This notion was first suggested by Trakhtenbrot, Rabinovich & Hirshfeld in terms of behaviour structures (see [155]). We will show that this equivalence is preserved by refinement. For systems without autoconcurrency, this equivalence coincides with $\approx_{wh}$. This will imply the result that $\approx_{wh}$ is invariant under refinement for systems without autoconcurrency.

Definition Let $\mathcal{E}, \mathcal{F}$ be event structures.

A relation $R \subseteq \text{Conf}(\mathcal{E}) \times \text{Conf}(\mathcal{F}) \times \mathcal{P}(E_{\mathcal{E}} \times E_{\mathcal{F}})$ is called a *history preserving bisimulation* between $\mathcal{E}$ and $\mathcal{F}$ if $(\emptyset, \emptyset, \emptyset) \in R$ and whenever $(X, Y, f) \in R$ then

- f is an isomorphism between $(X, \leq_X, l_{\mathcal{E}} \upharpoonright X)$ and $(Y, \leq_Y, l_{\mathcal{F}} \upharpoonright Y)$,
- $X \rightarrow_{\mathcal{E}} X' \Rightarrow \exists Y', f'$ with $Y \rightarrow_{\mathcal{F}} Y'$, $(X', Y', f') \in R$ and $f' \upharpoonright X = f$,
- $Y \rightarrow_{\mathcal{F}} Y' \Rightarrow \exists X', f'$ with $X \rightarrow_{\mathcal{E}} X'$, $(X', Y', f') \in R$ and $f' \upharpoonright X = f$.

$\mathcal{E}$ and $\mathcal{F}$ are *history preserving equivalent* ($\mathcal{E} \approx_h \mathcal{F}$) iff there exists a history preserving bisimulation between $\mathcal{E}$ and $\mathcal{F}$.

Clearly, we have $\mathcal{E} \approx_h \mathcal{F} \Rightarrow \mathcal{E} \approx_{wh} \mathcal{F}$. However the two systems of Example 4.3 are not h -equivalent.

Proposition $\mathcal{E} \approx_h \mathcal{F} \Rightarrow \mathcal{E} \approx_{pb} \mathcal{F}$.

Proof

We show that any history preserving bisimulation between $\mathcal{E}$ and $\mathcal{F}$ is also a pomset bisimulation between $\mathcal{E}$ and $\mathcal{F}$ (after leaving out the isomorphism component). Let R be an h -bisimulation, and suppose $(X, Y, f) \in R$ and $X \xrightarrow{u} X'$. Then $X \rightarrow_{\mathcal{E}} X'$, thus $\exists Y', f'$ with $Y \rightarrow_{\mathcal{F}} Y'$, $(X', Y', f') \in R$ and $f' \upharpoonright X = f$. Since f' is an isomorphism and $f' \upharpoonright X = f$, $\text{range}(f' \upharpoonright (X' \setminus X)) = \text{range}(f') \setminus \text{range}(f) = Y' \setminus Y$, so $f' \upharpoonright (X' \setminus X)$ is an isomorphism between $X' \setminus X$ and $Y' \setminus Y$. Hence $Y \xrightarrow{u} Y'$, so R satisfies the first clause of a pomset bisimulation. The second clause follows by symmetry. ■

From this proof we learn that h -bisimulation not only respects pomset bisimulation but even the previous proposal combining weak history preserving equivalence and pomset bisimulation. Thus $\approx_h$ is the strongest equivalence considered so far (except for event structure isomorphism of course). Nevertheless it is possible to abstract from certain details in a system representation: we have $a \approx_h a + a$ and $(a \parallel (b + c)) + (a \parallel b) + ((a + c) \parallel b) \approx_h (a \parallel (b + c)) + ((a + c) \parallel b)$ (absorption law).

We now show that considering only those transitions $X \rightarrow_{\mathcal{E}} X'$, $Y \rightarrow_{\mathcal{F}} Y'$, respectively, where $X'(Y')$ is obtained from $X(Y)$ by executing exactly one event yields the same equivalence. We write $X \triangleright_{\mathcal{E}} X'$ for $X \rightarrow_{\mathcal{E}} X'$ and $|X' \setminus X| = 1$. Let $\approx_{oh}$ be the equivalence notion obtained by replacing $\rightarrow$ by $\triangleright$ in the definition of $\approx_h$.

Proposition For event structures $\mathcal{E}, \mathcal{F} : \mathcal{E} \approx_h \mathcal{F}$ iff $\mathcal{E} \approx_{oh} \mathcal{F}$.

Proof

The implication $\mathcal{E} \approx_h \mathcal{F} \Rightarrow \mathcal{E} \approx_{oh} \mathcal{F}$ is trivial. The implication $\mathcal{E} \approx_{oh} \mathcal{F} \Rightarrow \mathcal{E} \approx_h \mathcal{F}$ immediately follows from the observation that whenever $X \rightarrow_{\mathcal{E}} X'$, there exist configurations $X_1, \dots, X_n (n \in \mathbb{N})$ such that $X = X_1 \triangleright_{\mathcal{E}} \dots \triangleright_{\mathcal{E}} X_n = X'$. ■

Next we show that $\approx_h$ is preserved by refinement.

Theorem Let $\mathcal{E}, \mathcal{F} \in \mathcal{E}$ and let ref be a refinement function.

Then $\mathcal{E} \approx_h \mathcal{F} \implies ref(\mathcal{E}) \approx_h ref(\mathcal{F})$.

Proof

Let $R \subseteq Conf(\mathcal{E}) \times Conf(\mathcal{F}) \times \mathcal{P}(E_{\mathcal{E}} \times E_{\mathcal{F}})$ be a history preserving bisimulation between $\mathcal{E}$ and $\mathcal{F}$. Define the relation $\tilde{R}$ by:

$$\begin{aligned} \tilde{R} = \{ & (\tilde{X}, \tilde{Y}, \tilde{f}) \in Conf(ref(\mathcal{E})) \times Conf(ref(\mathcal{F})) \times \mathcal{P}(E_{ref(\mathcal{E})} \times E_{ref(\mathcal{F})}) \mid \\ & \exists (X, Y, f) \in R \text{ such that} \\ & - \tilde{X} \text{ is a refinement of } X, \\ & - \tilde{Y} \text{ is a refinement of } Y \\ & - \text{ and } \tilde{f} : \tilde{X} \longrightarrow \tilde{Y} \text{ is a bijection, satisfying } \tilde{f}(e, e') = (f(e), e') \}. \end{aligned}$$

We show that $\tilde{R}$ is a history preserving bisimulation between $ref(\mathcal{E})$ and $ref(\mathcal{F})$.

- i. $(\emptyset, \emptyset, \emptyset) \in \tilde{R}$ since $(\emptyset, \emptyset, \emptyset) \in R$.
- ii. Suppose $(\tilde{X}, \tilde{Y}, \tilde{f}) \in \tilde{R}$. Take $(X, Y, f) \in R$ such that
 - $\tilde{X}$ is a refinement of X ,
 - $\tilde{Y}$ is a refinement of Y
 - and $\tilde{f} : \tilde{X} \longrightarrow \tilde{Y}$ is bijection, satisfying $\tilde{f}(e, e') = (f(e), e')$.

Now three things have to be established:

1. $\tilde{f}$ satisfies $(d, d') \leq_{\tilde{X}} (e, e') \iff \tilde{f}(d, d') \leq_{\tilde{Y}} \tilde{f}(e, e')$ and $l_{ref(\mathcal{F})}(\tilde{f}(e, e')) = l_{ref(\mathcal{E})}(e, e')$.
2. $\tilde{X} \longrightarrow_{ref(\mathcal{E})} \tilde{X}' \implies \exists \tilde{Y}', \tilde{f}'$ such that $\tilde{Y} \longrightarrow_{ref(\mathcal{F})} \tilde{Y}', \tilde{f}' \upharpoonright \tilde{X} = \tilde{f}$ and $(\tilde{X}', \tilde{Y}', \tilde{f}') \in \tilde{R}$.
3. $\tilde{Y} \longrightarrow_{ref(\mathcal{F})} \tilde{Y}' \implies \exists \tilde{X}', \tilde{f}'$ such that $\tilde{X} \longrightarrow_{ref(\mathcal{E})} \tilde{X}', \tilde{f}' \upharpoonright \tilde{X} = \tilde{f}$ and $(\tilde{X}', \tilde{Y}', \tilde{f}') \in \tilde{R}$.

ad 1. Straightforward.

ad 2. Suppose $\tilde{X} \longrightarrow_{ref(\mathcal{E})} \tilde{X}'$, i.e. $\tilde{X}' \in Conf(ref(\mathcal{E}))$ and $\tilde{X} \subseteq \tilde{X}'$.

We have $\tilde{X}' = \bigcup_{e \in \tilde{X}'} \{e\} \times X'_e$ where $X'_e \in Conf(\mathcal{E})$ and

$$\forall e \in X' : X'_e \in \text{Conf}(\text{ref}(l_{\mathcal{E}}(e))) - \{\emptyset\}.$$

Then $X = \text{pr}_1(\widetilde{X})$ and $X' = \text{pr}_1(\widetilde{X}')$, so $X \longrightarrow_{\mathcal{E}} X'$.

Since R is a history preserving bisimulation,

$$\exists Y', f' \text{ with } Y \longrightarrow_{\mathcal{F}} Y', f' \vdash X = f \text{ and } (X', Y', f') \in R.$$

$$\text{Let } \widetilde{Y}' = \{(f'(e), e') \mid (e, e') \in \widetilde{X}'\}$$

$$\text{and } \widetilde{f}' = \{((e, e'), (f'(e), e')) \mid (e, e') \in \widetilde{X}'\}.$$

It now suffices to show that $\widetilde{Y}'$ is a refinement of Y' , since then it follows immediately with Proposition 1.7 or 2.8 from Chapter IV that $\widetilde{Y}' \in \text{Conf}(\text{ref}(\mathcal{F}))$, $\widetilde{Y} \longrightarrow_{\text{ref}(\mathcal{F})} \widetilde{Y}'$ (using that $f' \vdash X = f$), $\widetilde{f}' \vdash X = \widetilde{f}$ (likewise) and $(\widetilde{X}', \widetilde{Y}', \widetilde{f}') \in \widetilde{R}$.

$$\begin{aligned} - \text{ By construction } \widetilde{Y}' &= \bigcup_{e \in X'} \{f'(e)\} \times X'_e = \bigcup_{e \in Y'} \{e\} \times Y'_e \text{ where} \\ \forall e \in Y' : Y'_e &= X'_{f'^{-1}(e)} \in \text{Conf}(\text{ref}(l_{\mathcal{E}}(f'^{-1}(e)))) - \{\emptyset\} = \\ &\text{Conf}(\text{ref}(l_{\mathcal{F}}(e))) - \{\emptyset\}. \end{aligned}$$

$$\begin{aligned} - e \in \text{busy}(\widetilde{Y}') &= \{e \in Y' \mid Y'_e \text{ not complete}\} \iff \\ f'^{-1}(e) \in \text{busy}(\widetilde{X}') &= \{e \in X' \mid X'_e \text{ not complete}\} \text{ by construc-} \\ &\text{tion.} \end{aligned}$$

Furthermore, e maximal in $Y' \iff f'^{-1}(e)$ maximal in X' , since f' is an isomorphism.

Hence $e \in \text{busy}(\widetilde{Y}')$ implies e maximal in Y' , since $\widetilde{X}'$ is a refinement of X' .

From this it follows that $\widetilde{Y}'$ is a refinement of Y' .

ad 3. By symmetry. ■

Finally we show that $\approx_{wh}$ and $\approx_h$ coincide for event structures where concurrent events may not carry the same label. As a corollary we then have that also $\approx_{wh}$ is preserved by refinement in this case and respects pomset bisimulation.

Definition $\mathcal{E}$ is an event structure *without autoconcurrency* iff

$$\forall d, e \in E_{\mathcal{E}} : d \text{ co } e \text{ and } l(d) = l(e) \Rightarrow d = e.$$

Theorem For event structures $\mathcal{E}, \mathcal{F}$ without autoconcurrency,

$$\mathcal{E} \approx_{wh} \mathcal{F} \Leftrightarrow \mathcal{E} \approx_h \mathcal{F}.$$

Proof

First note that a wh -bisimulation can be regarded as a h -bisimulation without the requirements that $f' \vdash X = f$.

Now “ $\Leftarrow$ ” is trivial.

In order to establish “ $\Rightarrow$ ” we will show that any wh -bisimulation between event structures $\mathcal{E}, \mathcal{F}$ without autoconcurrency is also a h -bisimulation. So let R be a h -bisimulation between such $\mathcal{E}$ and $\mathcal{F}$, without the requirements $f' \vdash X = f$. We

proof that these requirements are met nevertheless. Assume that $(X, Y, f) \in R$ and $X \rightarrow_{\mathcal{E}} X'$. Then there exists $(X', Y', f') \in R$ with $Y \rightarrow_F Y'$. Suppose $f' \upharpoonright X \neq f$. Then there exists an $e \in X$ with $f'(e) \neq f(e)$.

Now observe that if g is an isomorphism between two arbitrary partial orders $(X, \leq_X, l_X)$ and $(Y, \leq_Y, l_Y)$, and $g(e_1) = e_2$ then

$$|\{e' \in X \mid e' \leq_X e_1\}| = |\{e' \in Y \mid e' \leq_Y e_2\}|.$$

Hence we cannot have $f'(e) < f(e)$ or $f(e) < f'(e)$.

Since Y' is conflict-free we conclude $f'(e) \text{ co } f(e)$.

Moreover, f' and f preserve labelling, so $l_{\mathcal{F}}(f'(e)) = l_{\mathcal{E}}(e) = l_{\mathcal{F}}(f(e))$.

This is a contradiction since $\mathcal{F}$ was assumed to have no autoconcurrency. ■

Corollary

Let $\mathcal{E}, \mathcal{F}$ be event structures without autoconcurrency and let ref be a refinement.

Then $\mathcal{E} \approx_{wh} \mathcal{F} \Rightarrow \text{ref}(\mathcal{E}) \approx_{wh} \text{ref}(\mathcal{F})$.

Conclusion

In this chapter we have shown that equivalences based on interleaving of actions or of steps (multisets of concurrently executable actions) are not preserved when changing the level of atomicity of actions. However, we could show that certain equivalences based on modelling causal relations explicitly by partial orders are indeed preserved by refinement of actions. We considered “linear time” approaches, where the behaviour of a system is equated to the set of its possible runs, and “branching time” approaches, where the conflict structure of systems is taken into account. We could show the negative results about the interleaving approaches regardless of the level of detail in modelling conflict behaviour. However, for the positive results about the partial order approaches, the conflict structure turned out to be crucial. It turned out that linear time partial order semantics is preserved under action refinement, but the existing approaches to branching time partial order semantics are not. However, we found that a new and finer branching time partial order semantics due to RABINOVICH & TRAKHTENBROT [155]—which we call *history preserving bisimulation*—is robust under refinement. Subsequently, DEGANO, DE NICOLA & MONTANARI [60] show that this equivalence fits smoothly in their NMS framework [58] by combining totally and partially ordered observations into mixed orders, and VAANDRAGER [170] shows that the new equivalence coincides with bisimulation equivalence on causal trees [53, 54]. BEST, DEVILLERS, KIEHN & POMELLO [33] define this equivalence on Petri nets and also prove a preservation result for action refinement. In DEVILLERS [72, 73] this work is generalised to a setting with internal moves, whereas such a generalisation of our work appears

in VOGLER [179, 178]. Furthermore, VOGLER settles the decidability of history preserving bisimulation on finite safe nets in [177] and generalises this result to a setting with internal actions in [180].

A seemingly different equivalence, defined on a syntactical level on a subset of CCS, is proposed by HENNESSY [109]. In ACETO & HENNESSY [5] it is established that this *timed equivalence* is preserved by action refinement on this subset. In the next chapter of this thesis it is argued that timed equivalence coincides with history preserving equivalence on that subset. However, on the class of systems considered here, or on full CCS, timed equivalence is not preserved under refinement, as shown in VAN GLABBEK & VAANDRAGER [94] and ACETO & HENNESSY [6]. Therefore a finer equivalence, that is still coarser than history preserving bisimulation, is proposed in [6], see also ACETO [3]. This *refine equivalence* turns out to coincide (for τ -free systems) with the *ST-bisimulation* introduced in the next chapter. The relations between timed and refine equivalence—or in the terminology of [94] *split*- and ST-bisimulation equivalence—are further explored in GORRIERI & LANEVE [99]. A linear time variant of timed/refine equivalence is proven robust under refinement in NIELSEN, ENGBERG & LARSEN [141], also on a simple language where the difference between the two versions does not show up. An overview on refinement theorems in the setting of Petri nets appears in BROUWER, GOLD & VOGLER [41].

For an overview of our results consider the following diagram:

runs conflict structure	sequences of actions	sequences of steps	pomsets
paths	$\approx_{it}$	$\approx_{st}$	$\approx_{pt}$
e.g. testing			?
bisimulation	$\approx_{ib}$	$\approx_{sb}$	$\approx_{pb} \approx_{gpb}$ $\approx_{wh} \approx_{whpb}$ $\approx_h$

 means: not preserved by refinement

This diagram is not at all complete. A naturally arising question is to what extent it is actually necessary to move to partial orders to achieve invariance of equivalence under refinement (here we have only shown that steps are not sufficient). This question will be addressed in the last chapter of this thesis.

An interesting topic for further research (that has in the meantime been pursued in VOGLER [175, 178], HENNESSY [110] and JATEGAONKAR & MEYER [117]) is to investigate testing equivalences based on partial orders, taking the conflict structure in a weaker form into account. In each case, since the systems $\mathcal{E}'$ and $\mathcal{F}'$ of Example 4.2 are not even interleaving failure equivalent, no equivalence that is included between interleaving failure and pomset bisimulation equivalence, such as the partial order equivalence of ACETO, DE NICOLA & FANTECHI [4], can be preserved under action refinement.


When the text of this chapter was written (as part of [84]) it referred to conflict-free and finitary refinements using prime event structures as our system model. Here we have applied the same text to refinements with conflicts and with possibly infinite behaviours. For this it was convenient to use a more expressive model. We have chosen here *flow event structures* which are also very suited for giving semantics to CCS-like languages. We could have used simpler forms of event structures where just the axiom of finite causes is dropped, and the axiom of conflict heredity is dropped [39] or weakened [55, 57]. These forms of event structures may be seen—like prime event structures—as special cases of flow event structures. On two such models DARONDEAU & DEGANI [55, 57] also state a refinement theorem for history preserving bisimulation and the same class of refinements as considered here.

The refinement operation we have considered replaced actions by arbitrary non-empty event structures. As remarked in Chapter IV, it is debatable whether one should consider refinements where replacing actions by the empty event structure is allowed (*forgetful refinements*). Such refinements can drastically change the structure of processes; they can not be explained by a change in the level of abstraction at which processes are regarded. Nevertheless, in the setting of prime event structures our results hold also for forgetful refinements (with slightly more complicated proofs).

A problem which we have not addressed in detail is the treatment of deadlocks by our equivalences. Since the notion of configuration does not distinguish deadlock and termination and the equivalences considered here are based on this notion, these equivalences will not be congruences for sequential composition and for refinement (when viewed as a binary operation). For example, we may have $\forall a \in Act : ref(a) \approx_h ref'(a)$, however $ref(\mathcal{E}) \not\approx_h ref'(\mathcal{E})$. This problem is solved in VAN GLABBEEK & GOLTZ [88], where a deadlock sensitive variant of history preserving bisimulation is proposed.

Finally we would like to address the question whether history preserving bisimulation as defined here is the coarsest equivalence respecting pomset bisimu-

lation and being preserved by refinement. We conjecture that this is not the case, in particular, that for



$\mathcal{E} \not\approx_h \mathcal{F}$, but for any refinement ref , $ref(\mathcal{E}) \approx_{pb} ref(\mathcal{F})$. This conjecture has in the meantime be proven in VOGLER [179], using an equivalence that combines the concept of ST-bisimulation, as proposed in the next chapter, with pomset bisimulation.

Nevertheless, if it is required to model the interplay of causality and branching in full detail, history preserving bisimulation seems to be the coarsest suitable equivalence.

Chapter VII

The Refinement Theorem for ST-bisimulation Semantics

Rob van Glabbeek

In this chapter I prove that ST-bisimulation equivalence, as introduced in [91], is preserved under refinement of actions. This implies that it is possible to abstract from the causal structure of concurrent systems without assuming action atomicity.

TABLE OF CONTENTS

Introduction	243
1. Concurrent systems and refinement of actions	247
2. The behaviour of concurrent systems I	249
3. Equivalence notions for concurrent systems I	250
4. The behaviour of concurrent systems II	255
5. Equivalence notions for concurrent systems II	257
6. The refinement theorems	265
Concluding remarks	268

INTRODUCTION

Virtually all semantic equivalences employed in theories of concurrency are defined in terms of *actions* that concurrent systems may perform. Mostly, these actions are taken to be *atomic*, meaning that they are considered not to be divisible into smaller parts. In this case, the defined equivalences are said to be based on *action atomicity*.


However, in the top-down design of distributed systems it might be fruitful to model processes at different levels of abstraction. The actions on an abstract level then turn out to represent complex processes on a more concrete level. This methodology does not seem to be compatible with non-divisibility of actions and for this reason PRATT [154], LAMPORT [125] and others plead for the use of semantic equivalences that are not based on action atomicity.

As indicated in CASTELLANO, DE MICHELIS & POMELLO [46], the concept of action atomicity can be formalized by means of the notion of *refinement of actions*. A semantic equivalence is *preserved under action refinement* if two equivalent processes remain equivalent after replacing all occurrences of an action a by a more complicated process $r(a)$. In particular, $r(a)$ may be a sequence of two actions a_1 and a_2 . An equivalence is strictly based on action atomicity if it is not preserved under refinement.

Most semantic equivalences can be positioned in a two dimensional classification diagram, such as the one of Figure 1. On the x -axis equivalences are ordered with respect to the preserved level of detail of runs of processes. Three well-known points on this axis are *interleaving semantics*, where runs are represented by sequences of action occurrences, *step semantics*, where runs are represented by sequences of multisets of action occurrences - the multisets (or *steps*) representing simultaneous occurrences - and *partial order semantics*, in which all causal dependencies between action occurrences in runs of processes are preserved. On the y -axis the equivalences are ordered with respect to the preserved level of detail of the branching structure of these runs. Two well-known points on this axis are *trace semantics*, where a process is fully determined by the set of its possible (partial) runs, thereby completely neglecting the branching structure of processes, and *bisimulation semantics*, where also the information is preserved where two different courses of action diverge (although branching of identical courses of action is still neglected). In between there are several *decorated trace semantics*, where part of the branching structure is taken into account. Mostly these are motivated by the observable behaviour of processes, according to some testing scenario. In Figure 1 the equivalences become finer, or more discriminating, when moving upwards or to the right.

In [46], CASTELLANO, DE MICHELIS & POMELLO show by means of a simple example that none of the interleaving equivalences - not even bisimulation - is preserved under action refinement. Furthermore they claim that ‘on the other hand, the approaches based on partial order are not constrained to the assumption of atomicity’. Therefore they conclude that ‘interleaving semantics is adequate only if the abstraction level at which the atomic actions are defined is fixed. Otherwise, partial order semantics should be considered’.

In [84] (the previous chapter of this book), URSULA GOLTZ & I elaborated on this argument by providing examples, showing that also none of the step equivalences is preserved under refinement, and by formalizing the proof sketch of [46] that trace equivalence based on partial orders is invariant under refinement. We also wanted to prove this for bisimulation equivalence based on partial orders, but surprisingly we found that none of the partial order bisimulation equivalences proposed before publication of [46] is preserved under action refinement. However, we *did* prove a refinement theorem for a new notion of bisimulation equivalence based on partial orders, proposed recently by Hirshfeld, RABINOVICH & TRAKHTENBROT [155]. We chose to call

FIGURE 1. *Semantic equivalences*

this equivalence *history preserving bisimulation equivalence*, notation $\approx_h$. Hence, even in bisimulation semantics, the requirements of preservation under action refinement and capturing causal dependencies in processes by means of partial orders can be conciliated. But of course, this still does not show that in case preservation under refinement is required, it is *necessary* to employ partial order semantics. In this chapter I will show that it is not.¹

Event structures and Petri nets have been established as suitable domains for modelling (both branching and causal aspects of) concurrent systems. Usually a state of a concurrent system is represented by a *configuration* of the associated event structure, or by a *marking* of the associated net. In this chapter I argue that when events or transitions are considered to have a duration or structure, configurations or markings do not properly represent all the states of concurrent systems. Instead I propose to use so-called *ST-configurations* or *ST-markings*. The idea to model a state in a safe labelled marked net as the set of places (*Stellen*) containing a token, together with the set of transitions (*Transitionen*) which are currently firing (an *ST-marking*) originates from VAN GLABBEEK & VAANDRAGER [91]. In this chapter I translate this idea to the realm of event structures by introducing *ST-configurations*.

All interleaving, step and partial order equivalences on event structures or Petri nets considered so far, have been defined in terms of configurations or markings. If the constructions from interleaving semantics are applied on ST-configurations instead of ordinary configurations two new points on the *x*-axis of Figure 1 emerge. *Split-semantics* is just interleaving semantics, but based on

1. WALTER VOGLER [176,179] reads in the text of this chapter a claim that a purely interleaving semantics gives a congruence for action refinement. No such claim is made. A reaction to Vogler's position that partial order semantics *is* necessary for action refinement is added to the concluding remarks.

interleaving of beginnings and ends of events, instead of entire events; *ST-semantics* is a refinement of split semantics where in addition a link is required between the beginning and the end of any event. Split semantics is more discriminating than step semantics, whereas ST-semantics is as least as discriminating as split semantics. Furthermore ST-trace semantics is less discriminating than trace semantics based on partial orders and ST-bisimulation semantics is less discriminating than history preserving bisimulation semantics (but incomparable with the other bisimulation semantics based on partial orders proposed so far). Hence the situation is as indicated in Figure 2.




FIGURE 2. More semantic equivalences

ST-bisimulation equivalence was introduced by FRITS VAANDRAGER & ME in [91]. In the same paper we observed that for systems without autoconcurrency ST-bisimulation equivalence coincides with split bisimulation equivalence and provided a complete axiomatization on closed ACP-terms for the latter notion. Split bisimulation equivalence was proposed in HENNESSY [109] for a subset of CCS. ACETO & HENNESSY [5] proved that on this subset split bisimulation equivalence is preserved under action refinement. HENNESSY [109] also provided a complete axiomatization for split bisimulation equivalence on this subset. Since - if one forgets about τ -moves - this proof system is sound for ST-bisimulation equivalence, and even for history preserving bisimulation equivalence, it follows that on the domain considered in [109] the three equivalences coincide. In combination with the refinement theorem for history preserving bisimulation equivalence of the previous chapter this yields an alternative proof of Aceto & Hennessy's refinement theorem. Split trace equivalence has been considered in VAANDRAGER [172]. In a joint paper with FRITS VAANDRAGER [94] we will show that on the domain of labelled event structures (prime event structures with binary conflict), or on full CCS, split semantics is not proof against refinement. In fact the equivalences obtained by splitting an event into two parts (its beginning and its end) turned out to be different from the equivalences obtainable by splitting an event into three parts. This was established by means of a rather complicated example (the *owl*

example), that also shows that split semantics is strictly less discriminating than ST-semantics. By means of even more complicated examples we established that for each $n \in \mathbb{N}$ split- n semantics is also different from split- $n+1$ semantics.

The result contributed by the present chapter is that ST-bisimulation semantics as well as ST-trace semantics are preserved under action refinement. In [91] it was shown that these semantics do not respect causality. It follows that it is possible to abstract from the causal structure of concurrent systems without assuming action atomicity.

1. CONCURRENT SYSTEMS AND REFINEMENT OF ACTIONS

Also in this chapter I consider systems that are capable of performing actions from a given set Act of action names. Following [84], as my model for this kind of systems I have chosen labelled prime event structures with a binary conflict relation as introduced in NIELSEN, PLOTKIN & WINSKEL [142]; I could have chosen other models like Petri nets or behaviour structures [155] as well. In this chapter I will not distinguish external and internal actions; I do not consider abstraction by hiding of actions.

DEFINITION. A (labelled) event structure (over an alphabet Act) is a 4-tuple $\mathfrak{E} = (E, <, \#, l)$, where

- E is a set of events;
- $< \subseteq E \times E$ is a partial order (the causality relation) satisfying the principle of finite causes:

$$LACe' \in EVSTe' < eRAC \text{ is finite for } e \in E;$$

- $\# \subseteq E \times E$ is an irreflexive, symmetric relation (the conflict relation) satisfying the principle of conflict heredity:

$$e_1 \# e_2 < e_3 \Rightarrow e_1 \# e_3;$$

- $l: E \rightarrow Act$ is a labelling function.

An event structure represents a concurrent system in the following way: action names $a \in Act$ represent actions the system may perform, an event $e \in E$ labelled with a represents an occurrence of a during a possible run of the system, $e' < e$ means that e' is a prerequisite for e and $e' \# e$ means that e' and e cannot happen both in the same run.

One usually writes $e' \leq e$ for $e' < e \vee e' = e$, $>$ for $<^{-1}$ and $\geq$ for $\leq^{-1}$. Causal independence (concurrency) of events is expressed by the derived relation $\sim \subseteq E \times E$ defined by: $e' \sim e$ iff $\neg(e' \# e \vee e' < e \vee e' > e \vee e' = e)$. By definition $<$, $=$, $>$, $\#$ and $\sim$ form a partition of $E \times E$. The concurrency relation $co \subseteq E \times E$, originating from Petri net theory, is defined slightly different from $\sim$: $e' co e$ iff $e' \sim e \vee e' = e$.

The components of an event structure $\mathfrak{E}$ will be denoted by respectively

E_δ , $<_\delta$, $\#_\delta$ and l_δ . The derived relations will be denoted $\sim_\delta$, co_δ , $\leq_\delta$, $>_\delta$ and $\geq_\delta$.

Throughout the chapter, I assume a fixed set Act of action names as labelling set. Let E denote the domain of event structures labelled over Act .

DEFINITION. An *event structure isomorphism* between two event structures $\delta, \mathcal{F} \in E$ is a bijective mapping $f: E_\delta \rightarrow E_{\mathcal{F}}$ such that

- $f(e) <_{\mathcal{F}} f(e') \Leftrightarrow e <_\delta e'$,
- $f(e) \#_{\mathcal{F}} f(e') \Leftrightarrow e \#_\delta e'$ and
- $l_{\mathcal{F}}(f(e)) = l_\delta(e)$.

δ and $\mathcal{F}$ are *isomorphic* - notation $\delta \cong \mathcal{F}$ - if there exists an event structure isomorphism between them. Generally, one does not distinguish isomorphic event structures.

DEFINITION. The *restriction* of an event structure δ to a set $X \subseteq E_\delta$ of events is the event structure $\delta \upharpoonright X = (X, <_\delta \cap (X \times X), \#_\delta \cap (X \times X), l_\delta \upharpoonright X)$.

An event structure δ is *finite* if E_δ is finite; δ is *conflict-free* if $\#_\delta = \emptyset$. 0 denotes the *empty* event structure $(\emptyset, \emptyset, \emptyset, \emptyset)$.

In [46] it is shown that equivalence notions based on interleaving are not preserved when replacing an action in a system by a sequence of two actions. In section 1 of Chapter IV we considered a more general version of this operation, which I will also use in the present chapter: replacing actions by finite, conflict-free, non-empty event structures. Replacing actions by infinite event structures could in general invalidate the principle of finite causes. As explained in Chapter IV, replacing actions by event structures containing conflicts would require a more sophisticated notion of refinement or, alternatively, a more general form of event structures where the axiom of conflict heredity is dropped, e.g. flow event structures as in Section 2 of Chapter IV. The generalization of the results of this chapter to flow event structures seems to be completely straightforward, but has still to be carried out. Finally, replacing actions by the empty event structure can drastically change the structure of processes; it can not be explained by a change in the level of abstraction at which processes are regarded (Chapter IV). In the concluding section I will discuss possible extensions of my result to these cases.

A refinement will be a function r specifying for each action a an event structure $r(a)$ which is to be substituted for a . Interesting refinements will mostly refine only certain actions, hence replace most actions by themselves. However, for uniformity (and for simplicity in proofs) I consider all actions to be refined.

Given an event structure δ and a refinement r , the refined event structure $r(\delta)$ is constructed as follows. Each event e labelled by a is replaced by a disjoint copy, $r(e)$, of $r(a)$. The causality and conflict structure is inherited from δ : every event which was causally before e will be causally before all events of

$r(e)$, all events which causally followed e will causally follow all the events of $r(e)$, and all events in conflict with e will be in conflict with all the events of $r(e)$.

DEFINITION. A *refinement* $r: Act \rightarrow \mathbb{E} - \{0\}$ is a function that takes any action $a \in Act$ into a finite, conflict-free, non-empty event structure $r(a) \in \mathbb{E}$. If $\mathfrak{E} \in \mathbb{E}$ and r is a refinement, then $r(\mathfrak{E})$ is the event structure defined by:

- $E_{r(\mathfrak{E})} = \{(e, e') \mid e \in E_{\mathfrak{E}}, e' \in E_{r(l_{\mathfrak{E}}(e))}\};$
- $(e_1, e_1') <_{r(\mathfrak{E})} (e_2, e_2')$ iff $e_1 <_{\mathfrak{E}} e_2$ or $(e_1 = e_2 \wedge e_1' <_{r(l_{\mathfrak{E}}(e_1))} e_2')$;
- $(e_1, e_1') \#_{r(\mathfrak{E})} (e_2, e_2')$ iff $e_1 \#_{\mathfrak{E}} e_2$;
- $l_{r(\mathfrak{E})}(e, e') = l_{r(l_{\mathfrak{E}}(e))}(e')$.

PROPOSITION 1:

- i. If $\mathfrak{E} \in \mathbb{E}$ and r is a refinement then $r(\mathfrak{E})$ is an event structure indeed.
- ii. If $\mathfrak{E} \in \mathbb{E}$ and r, r' are refinements with $r(a) \cong r'(a)$ for $a \in Act$, then $r(\mathfrak{E}) \cong r'(\mathfrak{E})$.
- iii. If $\mathfrak{E}, \mathfrak{F} \in \mathbb{E}$, r is a refinement and $\mathfrak{E} \cong \mathfrak{F}$, then $r(\mathfrak{E}) \cong r(\mathfrak{F})$.

PROOF: Straightforward. □

This proposition says that refinement is a well-defined operation on event structures, even when isomorphic event structures are identified.

2. THE BEHAVIOUR OF CONCURRENT SYSTEMS I

Let $\mathfrak{E}$ be an event structure, modelling the behaviour of a concurrent system P . Classically, a state S of P is given by a set of events from $\mathfrak{E}$. Such a set is called a *configuration*. Its elements represent the occurrences of actions that happened before P reached the state S . If two events e' and e cannot happen both in the same run ($e' \#_{\mathfrak{E}} e$) then they also cannot occur in the same configuration. So configurations have to be *conflict-free*. Furthermore, if e occurs in a configuration C and e' is a prerequisite for e ($e' <_{\mathfrak{E}} e$) then also e' must occur in C . Hence configurations must be *left-closed* with respect to $<_{\mathfrak{E}}$. Finally, as is usual, in this book it is assumed that in a finite period only finitely many actions are performed. Therefore, unlike in many other papers, configurations are required to be finite here.

DEFINITION. A set $X \subseteq E_{\mathfrak{E}}$ of events in an event structure $\mathfrak{E}$ is *left-closed* in $\mathfrak{E}$ if for all $e, e' \in E_{\mathfrak{E}}$

$$e' <_{\mathfrak{E}} e \in X \Rightarrow e' \in X.$$

X is *conflict-free* in $\mathfrak{E}$ if $\mathfrak{E} \upharpoonright X$ is conflict-free. A *configuration* of $\mathfrak{E}$ is a finite, left-closed, conflict-free subset of $E_{\mathfrak{E}}$. Let $\mathcal{C}(\mathfrak{E})$ be the set of configurations of $\mathfrak{E}$. Write $X \rightarrow_{\mathfrak{E}} X'$ if $X, X' \in \mathcal{C}(\mathfrak{E})$ and $X \subseteq X'$.

$X \rightarrow_{\mathfrak{E}} X'$ says that both X and X' represent states of the concurrent system

represented by $\mathfrak{E}$, and that this system may evolve from the state represented by X to the one represented by X' .

As the lemma below will show, the behaviour of a refined event structure $r(\mathfrak{E})$ may be deduced from the behaviour of $\mathfrak{E}$ and from the behaviour of the event structures which are substituted for actions. On the other hand, one may derive information about the behaviour of $\mathfrak{E}$ from the behaviour of $r(\mathfrak{E})$.

Let $r(e)$ abbreviate $r(l_{\mathfrak{E}}(e))$ and let pr_1 denote projection to the first component.

LEMMA 2: Let $\mathfrak{E}$ be an event structure and r a refinement.

i. $\tilde{C} \subseteq E_{r(\mathfrak{E})}$ is a configuration of $r(\mathfrak{E})$ iff

$$\tilde{C} = \{(e, e') \mid e \in C, e' \in C_e\} \text{ where}$$

C is a configuration of $\mathfrak{E}$,

C_e is a configuration of $r(e)$ for $e \in C$,

$C_e = E_{r(e)}$ if e is not maximal in C with respect to $<_{\mathfrak{E}}$.

ii. If $\tilde{C} \rightarrow_{r(\mathfrak{E})} \tilde{C}'$ then $pr_1(\tilde{C}) \rightarrow_{\mathfrak{E}} pr_1(\tilde{C}')$.

PROOF: See [84]. A similar lemma has been proved in Chapter IV. $\square$

3. EQUIVALENCE NOTIONS FOR CONCURRENT SYSTEMS I

In this section the semantic equivalences of Figure 1 are defined in terms of configurations.

The interleaving equivalences can be defined by means of the *single action transition relations* $\xrightarrow{a}_{\mathfrak{E}} \subseteq \mathcal{Q}(\mathfrak{E}) \times \mathcal{Q}(\mathfrak{E})$ for $a \in Act$ and $\mathfrak{E} \in \mathbf{E}$.

DEFINITION. $C \xrightarrow{a}_{\mathfrak{E}} C'$ iff $C \rightarrow_{\mathfrak{E}} C'$ and $C' - C = \{e\}$ with $l_{\mathfrak{E}}(e) = a$.

Here $C \xrightarrow{a}_{\mathfrak{E}} C'$ says that if the system represented by $\mathfrak{E}$ is in the state represented by C , then it may perform an action a and reach the state represented by C' .

DEFINITION. A sequence $a_1 \cdots a_n \in Act^*$ is a (*sequential*) *trace* of an event structure $\mathfrak{E}$ if there exist configurations $C_0, \dots, C_n$ of $\mathfrak{E}$ such that $C_0 = \emptyset$ and

$$C_{i-1} \xrightarrow{a_i}_{\mathfrak{E}} C_i \quad (i = 1, \dots, n).$$

$SeqTraces(\mathfrak{E})$ denotes the set of all sequential traces of $\mathfrak{E}$.

Two event structures $\mathfrak{E}$ and $\mathfrak{F}$ are *interleaving trace equivalent* - notation $\mathfrak{E} \approx_{it} \mathfrak{F}$ - if $SeqTraces(\mathfrak{E}) = SeqTraces(\mathfrak{F})$.

DEFINITION. Let $\mathfrak{E}, \mathfrak{F} \in \mathbf{E}$. A relation $R \subseteq \mathcal{Q}(\mathfrak{E}) \times \mathcal{Q}(\mathfrak{F})$ is called a (*sequential*) *bisimulation* between $\mathfrak{E}$ and $\mathfrak{F}$ if $(\emptyset, \emptyset) \in R$ and whenever $(C, D) \in R$ then for $a \in Act$:

$$- \quad C \xrightarrow{a}_{\mathfrak{E}} C' \Rightarrow \exists D' \text{ with } D \xrightarrow{a}_{\mathfrak{F}} D' \text{ and } (C', D') \in R;$$

- $D \xrightarrow{a}_{\mathfrak{F}} D' \Rightarrow \exists C' \text{ with } C \xrightarrow{a}_{\mathfrak{E}} C' \text{ and } (C', D') \in R.$

$\mathfrak{E}$ and $\mathfrak{F}$ are *interleaving bisimulation equivalent* - $\mathfrak{E} \approx_{ib} \mathfrak{F}$ - if there exists a sequential bisimulation between them.

Step equivalences can be defined by generalizing the single action transition relations $\xrightarrow{a}_{\mathfrak{E}} \subseteq \mathcal{Q}(\mathfrak{E}) \times \mathcal{Q}(\mathfrak{E})$ to *step transition relations* $\xrightarrow{A}_{\mathfrak{E}} \subseteq \mathcal{Q}(\mathfrak{E}) \times \mathcal{Q}(\mathfrak{E})$, where A is a multiset over Act .

DEFINITION. Let $\mathfrak{E}$ be an event structure and $A : Act \rightarrow \mathbb{N}$ a multiset over Act . For $X \subseteq E_{\mathfrak{E}}$ let $l_{\mathfrak{E}}(X) \in \mathbb{N}^{Act}$ be the multiset of labels of the events from X , defined by $l_{\mathfrak{E}}(X)(a) = |\{e \in X \mid l_{\mathfrak{E}}(e) = a\}|$.

Then $C \xrightarrow{A}_{\mathfrak{E}} C'$ iff $C \rightarrow_{\mathfrak{E}} C'$ and $C' - C = G \subseteq E_{\mathfrak{E}}$ such that $\forall e, e' \in G: e \text{ co}_{\mathfrak{E}} e'$ and $l_{\mathfrak{E}}(G) = A$.

Here $C \xrightarrow{A}_{\mathfrak{E}} C'$ says that if the system represented by $\mathfrak{E}$ is in the state represented by C , then it may concurrently perform the multiset of actions A and reach the state represented by C' . Since A is a multiset rather than a set, actions may occur concurrently with themselves ('autoconcurrency').

DEFINITION. A sequence $A_1 \cdots A_n$ of multisets $A_i \in \mathbb{N}^{Act}$ ($i = 1, \dots, n$) is a *step trace* of an event structure $\mathfrak{E}$ if there exist configurations $C_0, \dots, C_n$ of $\mathfrak{E}$ such that $C_0 = \emptyset$ and $C_{i-1} \xrightarrow{A_i}_{\mathfrak{E}} C_i$ ($i = 1, \dots, n$).

$StepTraces(\mathfrak{E})$ denotes the set of all step traces of $\mathfrak{E}$.

Two event structures $\mathfrak{E}$ and $\mathfrak{F}$ are *step trace equivalent* - $\mathfrak{E} \approx_{st} \mathfrak{F}$ - if $StepTraces(\mathfrak{E}) = StepTraces(\mathfrak{F})$.

DEFINITION. Let $\mathfrak{E}, \mathfrak{F} \in \mathbb{E}$. A relation $R \subseteq \mathcal{Q}(\mathfrak{E}) \times \mathcal{Q}(\mathfrak{F})$ is called a *step bisimulation* between $\mathfrak{E}$ and $\mathfrak{F}$ if $(\emptyset, \emptyset) \in R$ and whenever $(C, D) \in R$ then for $A \in \mathbb{N}^{Act}$:

- $C \xrightarrow{A}_{\mathfrak{E}} C' \Rightarrow \exists D' \text{ with } D \xrightarrow{A}_{\mathfrak{F}} D' \text{ and } (C', D') \in R;$

- $D \xrightarrow{A}_{\mathfrak{F}} D' \Rightarrow \exists C' \text{ with } C \xrightarrow{A}_{\mathfrak{E}} C' \text{ and } (C', D') \in R.$

$\mathfrak{E}$ and $\mathfrak{F}$ are *step bisimulation equivalent* - $\mathfrak{E} \approx_{sb} \mathfrak{F}$ - if there exists a step bisimulation between them.

A trace equivalence preserving causal dependencies between action occurrences in runs of processes is the *pomset trace equivalence* as implicitly employed, for instance, in PRATT [154].

DEFINITION. A *partially ordered multiset (pomset)* is an isomorphism class of conflict-free event structures. A pomset u is a *pomset trace* of an event structure $\mathfrak{E}$ if u is the isomorphism class of $\mathfrak{E} \upharpoonright C$ for some configuration $C \in \mathcal{Q}(\mathfrak{E})$. $Pomsets(\mathfrak{E})$ denotes the set of all pomset traces of $\mathfrak{E}$.

Two event structures $\mathfrak{E}$ and $\mathfrak{F}$ are *pomset trace equivalent* - $\mathfrak{E} \approx_{pt} \mathfrak{F}$ - if $Pomsets(\mathfrak{E}) = Pomsets(\mathfrak{F})$.

Sequential traces, step traces as well as pomset traces of an event structure $\mathcal{E}$ represent possible (partial) runs of the system represented by $\mathcal{E}$. A trace of each of these three types specifies a multiset of actions, executed during such a run. However, whereas sequential and step traces in addition only specify a possible order in which these actions may occur (with and without the possibility of simultaneous occurrences), a pomset trace specifies all causal dependencies between the occurrences of these actions, through the partial order inherited from $\mathcal{E}$. From this information all the possible orders in which the actions may occur can be derived.

Like pomset trace equivalence, most of the equivalences that preserve causal dependencies between occurrences of actions are defined by means of partial orders. Therefore, such equivalences are called *partial order equivalences*. It happens that on $\mathbb{E}$ there is only one reasonable trace equivalence based on partial orders - namely $\approx_{pt}$ - and the same can be said about trace equivalences based on steps and on interleaving and about bisimulation equivalences based on steps and on interleaving. However, of late years several bisimulation equivalences based on partial orders have been defined on $\mathbb{E}$:

- 1986: the *NMS partial ordering equivalence* of DEGANO, DE NICOLA & MONTANARI [58],
- 1986: the *pomset bisimulation equivalence* or *equipollence* of BOUDOL & CASTELLANI [39],
- 1987: the *generalized pomset bisimulation equivalence* of VAN GLABBEK & VAANDRAGER [91] and
- 1988: the *behaviour structure bisimulation equivalence* of RABINOVICH & TRAKHTENBROT [155].

In my opinion only the last - and finest - one fully captures the interplay of causality and branching and is most worthy of filling up the right upper corner of Figure 1. Originally it was defined on behaviour structures [155], but in [84] (Chapter VI of this book) the notion was defined on event structures as well, under the name *history preserving bisimulation equivalence*.

DEFINITION. Let $\mathcal{E}, \mathcal{F} \in \mathbb{E}$. A relation $R \subseteq \mathcal{C}(\mathcal{E}) \times \mathcal{C}(\mathcal{F}) \times \mathcal{P}(E_{\mathcal{E}} \times E_{\mathcal{F}})$ is called a *history preserving bisimulation* between $\mathcal{E}$ and $\mathcal{F}$ if $(\emptyset, \emptyset, \emptyset) \in R$ and whenever $(C, D, f) \in R$ then:

- $f : C \rightarrow D$ is an isomorphism between $\mathcal{E} \upharpoonright C$ and $\mathcal{F} \upharpoonright D$;
- $C \rightarrow_{\mathcal{E}} C' \Rightarrow \exists D', f' \text{ with } D \rightarrow_{\mathcal{F}} D', (C', D', f') \in R \text{ and } f' \upharpoonright C = f$;
- $D \rightarrow_{\mathcal{F}} D' \Rightarrow \exists C', f' \text{ with } C \rightarrow_{\mathcal{E}} C', (C', D', f') \in R \text{ and } f' \upharpoonright C = f$.

$\mathcal{E}$ and $\mathcal{F}$ are *history preserving bisimulation equivalent* - $\mathcal{E} \approx_h \mathcal{F}$ - if there exists a history preserving bisimulation between them.

PROPOSITION 3: For all equivalences $\approx_1$ and $\approx_2$ defined in this section, the formula

$$\forall \mathcal{E}, \mathcal{F} \in \mathbb{E}: \mathcal{E} \approx_1 \mathcal{F} \Rightarrow \mathcal{E} \approx_2 \mathcal{F}$$

holds iff there is a path $\approx_1 \rightarrow \dots \rightarrow \approx_2$ in Figure 1.

PROOF: The implications follow directly from the definitions; in order to prove the absence of other implications, it suffices to provide counterexamples against $\approx_{pt} \rightarrow \approx_{ib}$, $\approx_{ib} \rightarrow \approx_{st}$ and $\approx_{sb} \rightarrow \approx_{pt}$.

COUNTEREXAMPLES. In the graphical representations of event structures below, the conventions of [172] are followed: the conflict relation is denoted by means of dotted lines, only immediate conflicts - not the inherited ones - are indicated; the causality relation is represented by arrows, omitting those derivable by transitivity; and instead of events only their labels are displayed, if a label occurs twice it represents two different events. Thus these pictures determine event structures only up to isomorphism.



FIGURE 3. *Pomset trace equivalent but not interleaving bisimulation equivalent (standard example)*

The two event structures of Figure 3 are pomset trace equivalent: their pomset traces are $a \rightarrow b$, $a \rightarrow c$, a and the empty pomset. However, they are not interleaving bisimulation equivalent: both systems represented perform first the action a and then either b or c , but the first system makes the choice between b and c after the execution of a whereas the second one starts with making this choice.



FIGURE 4. *Interleaving bisimulation equivalent but not step trace equivalent (standard example)*

The first system represented in Figure 4 performs two actions a and b concurrently. The second one either performs b after completion of a or vice versa. In interleaving semantics these systems are identified. However, they are not step trace equivalent: only the first system can perform a and b simultaneously.




FIGURE 5. *Step bisimulation equivalent but not pomset trace equivalent (new)*

The two systems represented in Figure 5 are step bisimulation equivalent: both systems perform the actions a , b and c exactly once; in both cases a is a prerequisite for b , and c can happen before a , simultaneous with a , between a and b , simultaneous with b , or after b ; and in both cases all choices between alternative courses of action are made only when one of the alternatives actually occurs. However, they are not pomset trace equivalent: the pomset resembling the first event structure of Figure 5 is a pomset trace of this first event structure, but not of the second one. $\square$

THEOREM: *Of all equivalences mentioned in this section, only $\approx_{pt}$ and $\approx_h$ are preserved under action refinement.*

PROOF: The two event structures of Figure 5 are step bisimulation equivalent. However, after refining c in $c_1 \rightarrow c_2$ the resulting event structures (below) are not even interleaving trace equivalent.




FIGURE 6. *Refined event structures*

Only the first one has a trace $c_1 a b c_2$. This shows that no equivalence that is at least as fine as interleaving trace equivalence and at least as coarse as step bisimulation equivalence is preserved under refinement of actions. More counterexamples and the refinement theorems for $\approx_{pt}$ and $\approx_h$ can be found in Chapter VI. $\square$

4. THE BEHAVIOUR OF CONCURRENT SYSTEMS II

A configuration of an event structure $\mathfrak{E}$ represents a state S of the system represented by $\mathfrak{E}$ by considering two kinds of events with respect to S : those that happened before the system reached this state and those that did not happen (yet). I argue that when events or transitions are considered to have a duration or structure, such configurations do not properly represent all the states of the represented system. Instead I propose to consider a third kind of events with respect to S : those that are currently happening when the system is in state S . This gives rise to the introduction of *ST-configurations* (a name explained in the introduction).

DEFINITION. An *ST-configuration* of $\mathfrak{E}$ is a pair (C, P) of subsets of $E_{\mathfrak{E}}$, such that $P \subseteq C$, C is finite and conflict-free and

$$e' <_{\mathfrak{E}} e \in C \Rightarrow e' \in P.$$

Thus both P and C are configurations and $C - P$ contains only maximal elements in C . An ST-configuration (C, P) represents the state of a concurrent system where C is the set of events whose execution has been started and P (the *past*) is the set of events whose execution has been completed. An ordinary configuration can be regarded as an ST-configuration with $P = C$. Let $\mathfrak{S}(\mathfrak{E})$ be the set of ST-configurations of $\mathfrak{E}$. Write $(C, P) \rightarrow_{\mathfrak{E}} (C', P')$ if $(C, P), (C', P') \in \mathfrak{S}(\mathfrak{E})$, $C \subseteq C'$ and $P \subseteq P'$.

As in Section 2, the behaviour of a refined event structure $r(\mathfrak{E})$ may be deduced from the behaviour of $\mathfrak{E}$ and from the behaviour of the event structures which are substituted for actions.

NOTATION. For each pair $(\tilde{C}, \tilde{P}) \in \mathcal{P}(E_{r(\mathfrak{E})}) \times \mathcal{P}(E_{r(\mathfrak{E})})$ with $\tilde{P} \subseteq \tilde{C}$, there are unique sets $C_e, P_e \subseteq E_{r(e)}$ for every $e \in pr_1(\tilde{C})$ such that $\tilde{C} = \{(e, e') \mid e \in pr_1(\tilde{C}), e' \in C_e\}$ and $\tilde{P} = \{(e, e') \mid e \in pr_1(\tilde{C}), e' \in P_e\}$. In fact $C_e = \{e' \mid (e, e') \in \tilde{C}\}$ and $P_e = \{e' \mid (e, e') \in \tilde{P}\}$. Now $r^{-1}(\tilde{C}, \tilde{P})$ denotes the unique pair $(C, P) \in \mathcal{P}(E_{\mathfrak{E}}) \times \mathcal{P}(E_{\mathfrak{E}})$ such that $C = pr_1(\tilde{C})$ and $P = \{e \in C \mid P_e = E_{r(e)}\}$.

LEMMA 4: Let $\mathfrak{E}$ be an event structure and r a refinement.

- i. $(\tilde{C}, \tilde{P}) \in \mathcal{P}(E_{r(\mathfrak{E})}) \times \mathcal{P}(E_{r(\mathfrak{E})})$ is an ST-configuration of $r(\mathfrak{E})$ iff
 - $\tilde{C} = \{(e, e') \mid e \in C, e' \in C_e\}$ and $\tilde{P} = \{(e, e') \mid e \in C, e' \in P_e\}$ where
 - (C, P) is an ST-configuration of $\mathfrak{E}$,
 - (C_e, P_e) is an ST-configuration of $r(e)$ for $e \in C$,
 - $P_e = E_{r(e)}$ iff $e \in P$.
- ii. If $(\tilde{C}, \tilde{P}) \rightarrow_{r(\mathfrak{E})} (\tilde{C}', \tilde{P}')$ then $r^{-1}(\tilde{C}, \tilde{P}) \rightarrow_{\mathfrak{E}} r^{-1}(\tilde{C}', \tilde{P}')$.

PROOF: i. " $\Rightarrow$ ". Let $(\tilde{C}, \tilde{P}) \in \mathcal{S}(r(\mathbb{G}))$. First I show that $(C, P) := r^{-1}(\tilde{C}, \tilde{P}) \in \mathcal{S}(\mathbb{G})$.

- $P \subseteq C$ by definition.
- C is finite and conflict-free since $\tilde{C}$ is finite and conflict-free.
- Suppose $d <_{\mathbb{G}} e \in C$. I have to show that $d \in P$.
 Since $e \in C = pr_1(\tilde{C})$ there exists $(e, e') \in \tilde{C}$;
 since $r(d)$ is non-empty there exists $(d, d') \in E_{r(\mathbb{G})}$;
 since $d <_{\mathbb{G}} e$ one has $(d, d') <_{r(\mathbb{G})} (e, e') \in \tilde{C}$;
 and since $(\tilde{C}, \tilde{P})$ is an ST-configuration it follows that $(d, d') \in \tilde{P} \subseteq \tilde{C}$.
 Thus $d \in C$. So it remains to be proven that $P_d = E_{r(d)}$. Obviously $P_d \subseteq E_{r(d)}$.
 Now let $d' \in E_{r(d)}$. Then $(d, d') \in E_{r(\mathbb{G})}$. Exactly as above one obtains $(d, d') \in \tilde{P}$, and hence $d' \in P_d$. Thus $P_d = E_{r(d)}$ and $d \in P$.

Next let $e \in C$. Put $C_e = \{e' \mid (e, e') \in \tilde{C}\}$ and $P_e = \{e' \mid (e, e') \in \tilde{P}\}$. I show that $(C_e, P_e) \in \mathcal{S}(r(e))$.

- $P_e \subseteq C_e$ since $\tilde{P} \subseteq \tilde{C}$.
- C_e is finite since $\tilde{C}$ is finite.
- C_e is conflict-free since $r(e)$ is conflict-free.
- Suppose $e' <_{r(e)} e'' \in C_e$. Then $(e, e') <_{r(\mathbb{G})} (e, e'') \in \tilde{C}$. Hence $(e, e') \in \tilde{P}$ and $e' \in P_e$.

Finally the third requirement is met by construction.

" $\Leftarrow$ ". Let $(C, P) \in \mathcal{S}(\mathbb{G})$ and $(C_e, P_e) \in \mathcal{S}(r(e))$ for $e \in C$. Suppose $P_e = E_{r(e)} \Leftrightarrow e \in P$ for $e \in C$. Put $\tilde{C} = \{(e, e') \mid e \in C, e' \in C_e\}$ and $\tilde{P} = \{(e, e') \mid e \in C, e' \in P_e\}$. I show that $(\tilde{C}, \tilde{P}) \in \mathcal{S}(r(\mathbb{G}))$.

- $\tilde{P} \subseteq \tilde{C}$ since $P_e \subseteq C_e$ for $e \in C$.
- $\tilde{C}$ is finite since C and C_e are finite.
- $\tilde{C}$ is conflict-free since $C = pr_1(\tilde{C})$ is conflict-free.
- Suppose $(d, d') <_{r(\mathbb{G})} (e, e') \in \tilde{C}$. Then $d <_{\mathbb{G}} e$ or $d = e \wedge d' <_{r(e)} e'$.
 If $d <_{\mathbb{G}} e$ then $d \in P$, since $e \in C$ and $(C, P) \in \mathcal{S}(\mathbb{G})$. Thus $d \in C$ and $P_d = E_{r(d)}$. Since $d' \in E_{r(d)} = P_d = \{d' \mid (d, d') \in \tilde{P}\}$ it follows that $(d, d') \in \tilde{P}$.
 If $d = e$ then $d' \in P_e = P_d$, since $d' <_{r(e)} e' \in C_e$ and $(C_e, P_e) \in \mathcal{S}(r(e))$. So also in this case one has $(d, d') \in \tilde{P}$, which had to be proved.

ii. Suppose $(\tilde{C}, \tilde{P}) \rightarrow_{r(\mathbb{G})} (\tilde{C}', \tilde{P}')$, i.e. $(\tilde{C}, \tilde{P}), (\tilde{C}', \tilde{P}') \in \mathcal{S}(r(\mathbb{G}))$, $\tilde{C} \subseteq \tilde{C}'$ and $\tilde{P} \subseteq \tilde{P}'$. Then $r^{-1}(\tilde{C}, \tilde{P}), r^{-1}(\tilde{C}', \tilde{P}') \in \mathcal{S}(\mathbb{G})$ (by i.) and $r^{-1}(\tilde{C}, \tilde{P}) \rightarrow_{\mathbb{G}} r^{-1}(\tilde{C}', \tilde{P}')$ by definition. $\square$

I will end this section with a proposition saying that the ST-configurations of an event structure $\mathfrak{E}$ describe the behaviour of the represented concurrent system in the same way as the ordinary configurations of the *split* event structure $\text{split}(\mathfrak{E})$, obtained from $\mathfrak{E}$ by splitting every action a into the sequence of actions a^+ and a^- , representing the beginning and the end of a .

DEFINITION. For Λ a set of labels, let $\mathbb{E}(\Lambda)$ denote the domain of event structures labelled over Λ . So $\mathbb{E} = \mathbb{E}(\text{Act})$.

A Λ -refinement $r: \text{Act} \rightarrow \mathbb{E}(\Lambda) - \{0\}$ is a function that takes any action $a \in \text{Act}$ into a finite, conflict-free, non-empty event structure $r(a) \in \mathbb{E}(\Lambda)$. So a refinement as defined in Section 1 is an *Act*-refinement. If $\mathfrak{E} \in \mathbb{E}$ and r is a Λ -refinement, then $r(\mathfrak{E}) \in \mathbb{E}(\Lambda)$ is defined exactly as in Section 1.

DEFINITION. Put $\text{Act}^\pm = \{a^+ \mid a \in \text{Act}\} \cup \{a^- \mid a \in \text{Act}\}$. Let $\text{split}: \text{Act} \rightarrow \mathbb{E}(\text{Act}^\pm)$ be the $\text{Act}^\pm$ -refinement defined by $E_{\text{split}(a)} = \{a^+, a^-\}$, $a^+ <_{\text{split}(a)} a^-$ and $l_{\text{split}(a)}(a^+) = a^+$, $l_{\text{split}(a)}(a^-) = a^-$. It induces a function $\text{split}: \mathbb{E}(\text{Act}) \rightarrow \mathbb{E}(\text{Act}^\pm)$. This function was introduced on Petri nets in [91], and on event structures in [172].

PROPOSITION 4: For each event structure $\mathfrak{E} \in \mathbb{E}$, there exists a bijective mapping $i_{\mathfrak{E}}: \mathfrak{S}(\mathfrak{E}) \rightarrow \mathfrak{C}(\text{split}(\mathfrak{E}))$, such that for $S \in \mathfrak{S}(\mathfrak{E})$:

$$S \longrightarrow_{\mathfrak{E}} S' \iff i_{\mathfrak{E}}(S) \longrightarrow_{\text{split}(\mathfrak{E})} i_{\mathfrak{E}}(S').$$

PROOF: $i_{\mathfrak{E}}(C, P) = \{(e, (l_{\mathfrak{E}}(e))^+) \mid e \in C\} \cup \{(e, (l_{\mathfrak{E}}(e))^-) \mid e \in P\}$. Verification of all requirements is straightforward. $\square$

5. EQUIVALENCE NOTIONS FOR CONCURRENT SYSTEMS II

In this section the remaining equivalences of Figure 2 are defined in terms of ST-configurations.

The most straightforward generalization of interleaving semantics to the setting of ST-configurations yields split semantics. Split equivalences can be defined by generalizing the single action transition relations $\xrightarrow{a}_{\mathfrak{E}} \subseteq \mathfrak{C}(\mathfrak{E}) \times \mathfrak{C}(\mathfrak{E})$ to *split transition relations* $\xrightarrow{a^+}_{\mathfrak{E}}, \xrightarrow{a^-}_{\mathfrak{E}} \subseteq \mathfrak{S}(\mathfrak{E}) \times \mathfrak{S}(\mathfrak{E})$, for $a \in \text{Act}$ and $\mathfrak{E} \in \mathbb{E}$.

DEFINITION. $(C, P) \xrightarrow{a^+}_{\mathfrak{E}} (C', P')$ iff $(C, P) \longrightarrow_{\mathfrak{E}} (C', P')$, $P' = P$ and $C' - C = \{e\}$ with $l_{\mathfrak{E}}(e) = a$.
 $(C, P) \xrightarrow{a^-}_{\mathfrak{E}} (C', P')$ iff $(C, P) \longrightarrow_{\mathfrak{E}} (C', P')$, $C' = C$ and $P' - P = \{e\}$ with $l_{\mathfrak{E}}(e) = a$.

Here $(C, P) \xrightarrow{a^+}_{\mathfrak{E}} (C', P')$ says that if the system represented by $\mathfrak{E}$ is in the state represented by (C, P) , then it may start performing an action a and reach the state represented by (C', P') .

Furthermore $(C, P) \xrightarrow{a^-}_{\mathfrak{E}} (C', P')$ says that if the system is in the state

represented by (C, P) , then it may end performing an action a and reach the state represented by (C', P') .

DEFINITION. A sequence $a_1 \cdots a_n \in (Act^\pm)^*$ is a *split trace* of an event structure $\mathfrak{E}$ if there exist ST-configurations $(C_0, P_0), \dots, (C_n, P_n)$ of $\mathfrak{E}$ such that $(C_0, P_0) = (\emptyset, \emptyset)$ and $(C_{i-1}, P_{i-1}) \xrightarrow{a_i}_{\mathfrak{E}} (C_i, P_i)$ ($i = 1, \dots, n$). $SplitTraces(\mathfrak{E})$ denotes the set of all split traces of $\mathfrak{E}$.

Two event structures $\mathfrak{E}$ and $\mathfrak{F}$ are *split trace equivalent* - $\mathfrak{E} \approx_{2t} \mathfrak{F}$ - if $SplitTraces(\mathfrak{E}) = SplitTraces(\mathfrak{F})$.

DEFINITION. Let $\mathfrak{E}, \mathfrak{F} \in \mathbf{E}$. A relation $R \subseteq \mathfrak{S}(\mathfrak{E}) \times \mathfrak{S}(\mathfrak{F})$ is called a *split bisimulation* between $\mathfrak{E}$ and $\mathfrak{F}$ if $((\emptyset, \emptyset), (\emptyset, \emptyset)) \in R$ and whenever $((C, P), (D, Q)) \in R$ then for $a \in Act^\pm$:

- $(C, P) \xrightarrow{a}_{\mathfrak{E}} (C', P') \Rightarrow \exists D', Q' \quad \text{with} \quad (D, Q) \xrightarrow{a}_{\mathfrak{F}} (D', Q') \quad \text{and} \quad ((C', P'), (D', Q')) \in R;$
- $(D, Q) \xrightarrow{a}_{\mathfrak{F}} (D', Q') \Rightarrow \exists C', P' \quad \text{with} \quad (C, P) \xrightarrow{a}_{\mathfrak{E}} (C', P') \quad \text{and} \quad ((C', P'), (D', Q')) \in R.$

$\mathfrak{E}$ and $\mathfrak{F}$ are *split bisimulation equivalent* - $\mathfrak{E} \approx_{2b} \mathfrak{F}$ - if there exists a split bisimulation between them.

Alternatively, split equivalences can be defined as ordinary interleaving equivalences on split event structures, and even as step equivalences on split event structures. The following proposition says that this yields the same trace and bisimulation equivalences as the definitions above.

PROPOSITION 5.1: $\mathfrak{E} \approx_{2t} \mathfrak{F} \Leftrightarrow split(\mathfrak{E}) \approx_{it} split(\mathfrak{F}) \Leftrightarrow split(\mathfrak{E}) \approx_{st} split(\mathfrak{F})$
 $\mathfrak{E} \approx_{2b} \mathfrak{F} \Leftrightarrow split(\mathfrak{E}) \approx_{ib} split(\mathfrak{F}) \Leftrightarrow split(\mathfrak{E}) \approx_{sb} split(\mathfrak{F}).$

PROOF: Let $i_{\mathfrak{E}}: \mathfrak{S}(\mathfrak{E}) \rightarrow \mathcal{C}(split(\mathfrak{E}))$ be the bijection from the previous proposition, then for $S \in \mathfrak{S}(\mathfrak{E})$ and $a \in Act^\pm$: $S \xrightarrow{a}_{\mathfrak{E}} S' \Leftrightarrow i_{\mathfrak{E}}(S) \xrightarrow{a}_{split(\mathfrak{E})} i_{\mathfrak{E}}(S')$. Furthermore, if $C, C' \in \mathcal{C}(split(\mathfrak{E}))$ and A is a multiset over $Act^\pm$ consisting of the actions $a_1^+, \dots, a_n^+, b_1^-, \dots, b_m^-$ then

$$C \xrightarrow{A}_{split(\mathfrak{E})} C' \Leftrightarrow C \xrightarrow{a_1^+}_{split(\mathfrak{E})} \cdots \xrightarrow{a_n^+}_{split(\mathfrak{E})} \xrightarrow{b_1^-}_{split(\mathfrak{E})} \cdots \xrightarrow{b_m^-}_{split(\mathfrak{E})} C'.$$

From this the proposition follows immediately. $\square$

Split-semantics is just interleaving semantics, but based on interleaving of beginnings and ends of action occurrences, instead of entire action occurrences. However, since different occurrences of the same action can not be distinguished, it is in general not possible to tell when an occurrence of a^+ and an occurrence of a^- originate from to the same occurrence of a . ST-semantics is a refinement of split semantics, where occurrences of a^+ and a^- are explicitly connected if they represent the beginning and end of the same occurrence of a .

DEFINITION. A *pre-interval sequence* is a triple (E, l, σ) with E a set, $l: E \rightarrow Act$ a labelling function and σ a sequence over $E^\pm = \{e^+ \mid e \in E\} \cup \{e^- \mid e \in E\}$ whose elements are all different, and which can contain e^- only after e^+ (for $e \in E$). For $l: E \rightarrow Act$ define $l^\pm: E^\pm \rightarrow Act^\pm$ by $l(e^+) = (l(e))^+$ and $l(e^-) = (l(e))^-$. Let (E, l, σ) with $\sigma = \alpha_1 \cdots \alpha_n \in (E^\pm)^*$ be a pre-interval sequence and let $1 \leq i < j \leq n$. α_i and α_j are *connected*, notation $\alpha_i < \alpha_j$, if $\alpha_i = e^+$ and $\alpha_j = e^-$ for certain $e \in E$. Now two pre-interval sequences $(E, l, \alpha_1 \cdots \alpha_n)$ and $(E', l', \beta_1 \cdots \beta_m)$ are *isomorphic* if $n = m$, $l^\pm(\alpha_i) = l'^\pm(\beta_i)$ for $1 \leq i \leq n$, and $\alpha_i < \alpha_j \Leftrightarrow \beta_i < \beta_j$ for $1 \leq i < j \leq n$. An *interval sequence* is an isomorphism class of pre-interval sequences.

EXAMPLE: Let $E = \{e_0, e_1, e_2, e_3, e_4\}$, $l(e_0) = l(e_4) = b$ and $l(e_1) = l(e_2) = l(e_3) = a$. Figure 7 shows a pre-interval sequence over E , together with its associated interval sequence. The connectedness relation $<$ is represented by arcs.




FIGURE 7. Pre-interval sequence and interval sequence

DEFINITION. $(C, P) \xrightarrow{e^+}_{\mathfrak{E}} (C', P')$ iff $(C, P) \rightarrow_{\mathfrak{E}} (C', P')$, $P' = P$ and $C' - C = \{e\}$.
 $(C, P) \xrightarrow{e^-}_{\mathfrak{E}} (C', P')$ iff $(C, P) \rightarrow_{\mathfrak{E}} (C', P')$, $C' = C$ and $P' - P = \{e\}$.

A structure $(E_{\mathfrak{E}}, l_{\mathfrak{E}}, \alpha_1 \cdots \alpha_n)$ is a *pre-ST-trace* of an event structure $\mathfrak{E}$ if there exist ST-configurations $(C_0, P_0), \dots, (C_n, P_n)$ of $\mathfrak{E}$ such that $(C_0, P_0) = (\emptyset, \emptyset)$ and $(C_{i-1}, P_{i-1}) \xrightarrow{\alpha_i}_{\mathfrak{E}} (C_i, P_i)$ ($i = 1, \dots, n$). An *ST-trace* of $\mathfrak{E}$ is an interval sequence which is the isomorphism class of a pre-ST-trace of $\mathfrak{E}$.

$ST-Traces(\mathfrak{E})$ denotes the set of all ST-traces of $\mathfrak{E}$.

Two event structures $\mathfrak{E}$ and $\mathfrak{F}$ are *ST-trace equivalent* - $\mathfrak{E} \approx_{ST} \mathfrak{F}$ - if $ST-Traces(\mathfrak{E}) = ST-Traces(\mathfrak{F})$.

Next I propose another characterization of ST-trace equivalence that will be more convenient later on.

DEFINITION. $\mathfrak{E} \lesssim_{ST} \mathfrak{F}$ iff for every chain of ST-configurations

$$(\emptyset, \emptyset) \rightarrow_{\mathfrak{E}} (C_1, P_1) \rightarrow_{\mathfrak{E}} \cdots \rightarrow_{\mathfrak{E}} (C_n, P_n)$$

in $\mathfrak{E}$ there is a chain

$$(\emptyset, \emptyset) \rightarrow_{\mathfrak{F}} (D_1, Q_1) \rightarrow_{\mathfrak{F}} \cdots \rightarrow_{\mathfrak{F}} (D_n, Q_n)$$

in $\mathfrak{F}$ and a bijection $f: C_n \rightarrow D_n$, satisfying $l_{\mathfrak{F}}(f(e)) = l_{\mathfrak{E}}(e)$, $f(C_i) = D_i$ and $f(P_i) = Q_i$ for $i = 1, \dots, n$.

PROPOSITION 5.2: $\mathfrak{E} \approx_{ST} \mathfrak{F} \Leftrightarrow (\mathfrak{E} \lesssim_{ST} \mathfrak{F} \wedge \mathfrak{F} \lesssim_{ST} \mathfrak{E})$.

PROOF: Write $\mathfrak{E} \lesssim_{ST}^* \mathfrak{F}$ iff for every chain

$$(\emptyset, \emptyset) \xrightarrow{\alpha_1}_{\mathfrak{E}} (C_1, P_1) \xrightarrow{\alpha_2}_{\mathfrak{E}} \cdots \xrightarrow{\alpha_n}_{\mathfrak{E}} (C_n, P_n)$$

in $\mathfrak{E}$ (with $\alpha_i \in E_{\mathfrak{E}}^{\pm}$) there is a chain

$$(\emptyset, \emptyset) \xrightarrow{\beta_1}_{\mathfrak{F}} (D_1, Q_1) \xrightarrow{\beta_2}_{\mathfrak{F}} \cdots \xrightarrow{\beta_n}_{\mathfrak{F}} (D_n, Q_n)$$

in $\mathfrak{F}$ and a bijection $f: C_n \rightarrow D_n$, satisfying $l_{\mathfrak{F}}(f(e)) = l_{\mathfrak{E}}(e)$, $f(C_i) = D_i$ and $f(P_i) = Q_i$ for $i = 1, \dots, n$.

Furthermore write $\mathfrak{E} \lesssim_{ST}^{**} \mathfrak{F}$ iff for every chain

$$(\emptyset, \emptyset) \xrightarrow{\alpha_1}_{\mathfrak{E}} (C_1, P_1) \xrightarrow{\alpha_2}_{\mathfrak{E}} \cdots \xrightarrow{\alpha_n}_{\mathfrak{E}} (C_n, P_n)$$

in $\mathfrak{E}$ (with $\alpha_i \in E_{\mathfrak{E}}^{\pm}$) there is a chain

$$(\emptyset, \emptyset) \xrightarrow{\beta_1}_{\mathfrak{F}} (D_1, Q_1) \xrightarrow{\beta_2}_{\mathfrak{F}} \cdots \xrightarrow{\beta_n}_{\mathfrak{F}} (D_n, Q_n)$$

in $\mathfrak{F}$ such that $l_{\mathfrak{E}}^{\pm}(\alpha_i) = l_{\mathfrak{F}}^{\pm}(\beta_i)$ for $1 \leq i \leq n$, and $\alpha_i < \alpha_j \Leftrightarrow \beta_i < \beta_j$ for $1 \leq i < j \leq n$.

CLAIM 1: $\mathfrak{E} \lesssim_{ST} \mathfrak{F} \Leftrightarrow \mathfrak{E} \lesssim_{ST}^* \mathfrak{F}$.

CLAIM 2: $\mathfrak{E} \lesssim_{ST}^* \mathfrak{F} \Leftrightarrow \mathfrak{E} \lesssim_{ST}^{**} \mathfrak{F}$.

CLAIM 3: $\mathfrak{E} \lesssim_{ST}^{**} \mathfrak{F} \Leftrightarrow ST\text{-Traces}(\mathfrak{E}) \subseteq ST\text{-Traces}(\mathfrak{F})$.

Now the proposition follows by combination of these claims.

Proof of claim 1: " $\Rightarrow$ ". Suppose $\mathfrak{E} \lesssim_{ST} \mathfrak{F}$ and

$$(\emptyset, \emptyset) \xrightarrow{\alpha_1}_{\mathfrak{E}} (C_1, P_1) \xrightarrow{\alpha_2}_{\mathfrak{E}} \cdots \xrightarrow{\alpha_n}_{\mathfrak{E}} (C_n, P_n) \text{ is a chain in } \mathfrak{E} \text{ with } \alpha_i \in E_{\mathfrak{E}}^{\pm}.$$

Then there must be a chain $(\emptyset, \emptyset) \rightarrow_{\mathfrak{F}} (D_1, Q_1) \rightarrow_{\mathfrak{F}} \cdots \rightarrow_{\mathfrak{F}} (D_n, Q_n)$ in $\mathfrak{F}$ and a bijection $f: C_n \rightarrow D_n$, satisfying $l_{\mathfrak{F}}(f(e)) = l_{\mathfrak{E}}(e)$, $f(C_i) = D_i$ and $f(P_i) = Q_i$ for $i = 1, \dots, n$. Because of this bijection - only considering the 'sizes' of D_i and Q_i - there must be $\beta_i \in E_{\mathfrak{F}}^{\pm}$ for $i = 1, \dots, n$ such that

$$(C_{i-1}, P_{i-1}) \xrightarrow{\beta_i}_{\mathfrak{F}} (C_i, P_i).$$

" $\Leftarrow$ ". This follows from the observation that whenever in an event structure $\mathfrak{E}$ $(C, P) \rightarrow_{\mathfrak{E}} (C', P')$, there exist ST-configurations $(C_0, P_0), \dots, (C_k, P_k)$ of $\mathfrak{E}$ and a sequence $\alpha_1 \cdots \alpha_k \in (E_{\mathfrak{E}}^{\pm})^*$ such that $(C_0, P_0) = (C, P)$, $(C_{i-1}, P_{i-1}) \xrightarrow{\alpha_i}_{\mathfrak{E}} (C_i, P_i)$ ($i = 1, \dots, k$), and $(C_k, P_k) = (C', P')$.

Proof of claim 2: " $\Rightarrow$ ". Let $(\emptyset, \emptyset) \xrightarrow{\alpha_1}_{\mathfrak{E}} \cdots \xrightarrow{\alpha_n}_{\mathfrak{E}} (C_n, P_n)$ and $(\emptyset, \emptyset) \xrightarrow{\beta_1}_{\mathfrak{F}} \cdots \xrightarrow{\beta_n}_{\mathfrak{F}} (D_n, Q_n)$ be chains of ST-configurations in $\mathfrak{E}$ and $\mathfrak{F}$ with $\alpha_i \in E_{\mathfrak{E}}^{\pm}$ and $\beta_i \in E_{\mathfrak{F}}^{\pm}$ for $i = 1, \dots, n$ and let $f: C_n \rightarrow D_n$ be a bijection, satisfying $l_{\mathfrak{F}}(f(e)) = l_{\mathfrak{E}}(e)$, $f(C_i) = D_i$ and $f(P_i) = Q_i$ for

$i = 1, \dots, n$. Since $f(C_i) = D_i$ and $f(P_i) = Q_i$ it follows that $\alpha_i = e^+ \Leftrightarrow \beta_i = f(e)^+$ and $\alpha_i = e^- \Leftrightarrow \beta_i = f(e)^-$ for $i = 1, \dots, n$. Hence $l_{\mathfrak{E}}^{\pm}(\alpha_i) = l_{\mathfrak{F}}^{\pm}(\beta_i)$ for $1 \leq i \leq n$, and $\alpha_i < \alpha_j \Leftrightarrow \beta_i < \beta_j$ for $1 \leq i < j \leq n$.
 " $\Leftarrow$ ". Let $(\emptyset, \emptyset) \xrightarrow{\alpha_1}_{\mathfrak{E}} (C_1, P_1) \xrightarrow{\alpha_2}_{\mathfrak{E}} \dots \xrightarrow{\alpha_n}_{\mathfrak{E}} (C_n, P_n)$ and $(\emptyset, \emptyset) \xrightarrow{\beta_1}_{\mathfrak{F}} \dots \xrightarrow{\beta_n}_{\mathfrak{F}} (D_n, Q_n)$ be chains of ST-configurations in $\mathfrak{E}$ and $\mathfrak{F}$ with $\alpha_i \in E_{\mathfrak{E}}^{\pm}$ and $\beta_i \in E_{\mathfrak{F}}^{\pm}$ for $i = 1, \dots, n$ such that $l_{\mathfrak{E}}^{\pm}(\alpha_i) = l_{\mathfrak{F}}^{\pm}(\beta_i)$ for $1 \leq i \leq n$, and $\alpha_i < \alpha_j \Leftrightarrow \beta_i < \beta_j$ for $1 \leq i < j \leq n$. Note that $C_i = \{e \in E_{\mathfrak{E}} \mid \exists j \leq i: \alpha_j = e^+\}$ and $P_i = \{e \in E_{\mathfrak{E}} \mid \exists j \leq i: \alpha_j = e^-\}$ and similarly for D_i and Q_i . Define $f: C_n \rightarrow D_n$ by $f(e) = d \Leftrightarrow \exists i \leq n: (\alpha_i = e^+ \wedge \beta_i = d^+)$. Since $l_{\mathfrak{E}}^{\pm}(\alpha_i) = l_{\mathfrak{F}}^{\pm}(\beta_i)$ for $1 \leq i \leq n$, f is well-defined and bijective, and satisfies $l_{\mathfrak{F}}(f(e)) = l_{\mathfrak{E}}(e)$ and $f(C_i) = D_i$ for $i = 1, \dots, n$. Finally $e \in P_i \Leftrightarrow \exists k < j \leq i: (\alpha_k = e^+ \wedge \alpha_j = e^-) \Leftrightarrow \exists k < j \leq i: (\beta_k = f(e)^+ \wedge \beta_j = f(e)^-)$ (using $\alpha_k < \alpha_j \Leftrightarrow \beta_k < \beta_j$) $\beta_j = f(e)^- \Leftrightarrow f(e) \in Q_i$ so $f(P_i) = Q_i$ for $i = 1, \dots, n$.

Finally claim 3 follows directly from the definitions. $\square$

ST-bisimulation equivalence will be defined in the same style as the alternative characterization of ST-trace equivalence. The connection of occurrences of a^+ and a^- that represent the beginning and end of the same occurrence of a is implemented by means of a bijection between related ST-configurations.

DEFINITION. Let $\mathfrak{E}, \mathfrak{F} \in \mathbf{E}$. A relation $R \subseteq \mathcal{S}(\mathfrak{E}) \times \mathcal{S}(\mathfrak{F}) \times \mathcal{P}(E_{\mathfrak{E}} \times E_{\mathfrak{F}})$ is called an *ST-bisimulation* between $\mathfrak{E}$ and $\mathfrak{F}$ if $((\emptyset, \emptyset), (\emptyset, \emptyset), \emptyset) \in R$ and whenever $((C, P), (D, Q), f) \in R$ then:

- $f: C \rightarrow D$ is a bijection, satisfying $l_{\mathfrak{F}}(f(e)) = l_{\mathfrak{E}}(e)$ and $f(P) = Q$;
- $(C, P) \rightarrow_{\mathfrak{E}} (C', P') \Rightarrow \exists D', Q', f'$ with $(D, Q) \rightarrow_{\mathfrak{F}} (D', Q')$,
 $((C', P'), (D', Q'), f') \in R$ and $f' \upharpoonright C = f$;
- $(D, Q) \rightarrow_{\mathfrak{F}} (D', Q') \Rightarrow \exists C', P', f'$ with $(C, P) \rightarrow_{\mathfrak{E}} (C', P')$,
 $((C', P'), (D', Q'), f') \in R$ and $f' \upharpoonright C = f$.

$\mathfrak{E}$ and $\mathfrak{F}$ are *ST-bisimulation equivalent* - $\mathfrak{E} \approx_{STb} \mathfrak{F}$ - if there exists an ST-bisimulation between them.

Remark that the same equivalence is obtained if in the definition above the general transition relations $\rightarrow$ are replaced by the split transition relations $\xrightarrow{a}$ for $a \in Act^{\pm}$. One direction follows from the requirements for the bijection f ; the other one follows as in the proof of Proposition 5.2. Analogously, in the previous chapter it was shown that the definition of history preserving bisimulation equivalence is invariant under replacement of the general transition relations $\rightarrow$ by the single action transition relations $\xrightarrow{a}$ for $a \in Act$. Now it is not difficult to show that if in this version of the definition of ST-bisimulation equivalence the requirement $f(P) = Q$ would be skipped, the resulting equivalence would be split bisimulation equivalence again. This

requirement ensures the connection of occurrences of a^+ and a^- originating from the same occurrence of a .

As for split equivalences, the ST-equivalences can be defined alternatively by means of split event structures. First some preliminary definitions.

DEFINITION. For $\mathfrak{E} \in \mathbb{E}(Act^\pm)$, define the connectedness relation $<_{\mathfrak{E}} \subseteq E_{\mathfrak{E}} \times E_{\mathfrak{E}}$ by

$$e' <_{\mathfrak{E}} e \text{ iff } l_{\mathfrak{E}}(e) = a^- \text{ for certain } a \in Act \text{ and for } d \in E_{\mathfrak{E}}: (d <_{\mathfrak{E}} e \Leftrightarrow d \leq_{\mathfrak{E}} e').$$

DEFINITION. Write $C \xrightarrow{e}_{\mathfrak{E}} C'$ iff $C \rightarrow_{\mathfrak{E}} C'$ and $C' - C = \{e\}$. A sequence $\alpha_1 \cdots \alpha_n \in E_{\mathfrak{E}}^*$ is a *pre-trace* of an event structure $\mathfrak{E} \in \mathbb{E}(Act^\pm)$ if there exist configurations $C_0, \dots, C_n$ of $\mathfrak{E}$ such that $C_0 = \emptyset$ and $C_{i-1} \xrightarrow{\alpha_i}_{\mathfrak{E}} C_i$ ($i = 1, \dots, n$). Two pre-traces $\alpha_1 \cdots \alpha_n$ and $\beta_1 \cdots \beta_m$ of $\mathfrak{E}$ and $\mathfrak{F}$ are *<-isomorphic* if $n = m$, $l_{\mathfrak{E}}(\alpha_i) = l_{\mathfrak{F}}(\beta_i)$ for $1 \leq i \leq n$, and $\alpha_i <_{\mathfrak{E}} \alpha_j \Leftrightarrow \beta_i <_{\mathfrak{F}} \beta_j$ for $1 \leq i < j \leq n$. A *<-trace* of $\mathfrak{E}$ is the isomorphism class of a pre-trace of $\mathfrak{E}$. $<-Traces(\mathfrak{E})$ denotes the set of all <-traces of $\mathfrak{E}$. Two event structures $\mathfrak{E}$ and $\mathfrak{F} \in \mathbb{E}(Act^\pm)$ are *<-trace equivalent* - $\mathfrak{E} \approx_{<-t} \mathfrak{F}$ - if $<-Traces(\mathfrak{E}) = <-Traces(\mathfrak{F})$.

DEFINITION. Let $\mathfrak{E}, \mathfrak{F} \in \mathbb{E}(Act^\pm)$. A relation $R \subseteq \mathcal{C}(\mathfrak{E}) \times \mathcal{C}(\mathfrak{F}) \times \mathcal{P}(E_{\mathfrak{E}} \times E_{\mathfrak{F}})$ is called a *<-bisimulation* between $\mathfrak{E}$ and $\mathfrak{F}$ if $(\emptyset, \emptyset, \emptyset) \in R$ and whenever $(C, D, f) \in R$ then:

- $f: C \rightarrow D$ is a bijection, satisfying $l_{\mathfrak{F}}(f(e)) = l_{\mathfrak{E}}(e)$ and $f(e) <_{\mathfrak{F}} f(e') \Leftrightarrow e <_{\mathfrak{E}} e'$;
 - $C \rightarrow_{\mathfrak{E}} C' \Rightarrow \exists D', f'$ with $D \rightarrow_{\mathfrak{F}} D'$, $(C', D', f') \in R$ and $f' \upharpoonright C = f$;
 - $D \rightarrow_{\mathfrak{F}} D' \Rightarrow \exists C', f'$ with $C \rightarrow_{\mathfrak{E}} C'$, $(C', D', f') \in R$ and $f' \upharpoonright C = f$.
- $\mathfrak{E}$ and $\mathfrak{F}$ are *<-bisimulation equivalent* - $\mathfrak{E} \approx_{<-b} \mathfrak{F}$ - if there exists a <-bisimulation between them.

PROPOSITION 5.3: $\mathfrak{E} \approx_{STt} \mathfrak{F} \Leftrightarrow split(\mathfrak{E}) \approx_{<-t} split(\mathfrak{F})$

$$\mathfrak{E} \approx_{STb} \mathfrak{F} \Leftrightarrow split(\mathfrak{E}) \approx_{<-b} split(\mathfrak{F}).$$

PROOF: For $\mathfrak{E} \in \mathbb{E}$ define $i: E_{\mathfrak{E}}^\pm \rightarrow E_{split(\mathfrak{E})}$ by $i(e^+) = (e, (l_{\mathfrak{E}}(e))^+)$ and $i(e^-) = (e, (l_{\mathfrak{E}}(e))^-)$. Now the bijections $i_{\mathfrak{E}}: \mathfrak{S}(\mathfrak{E}) \rightarrow \mathcal{C}(split(\mathfrak{E}))$ from Proposition 4 satisfy for $S \in \mathfrak{S}(\mathfrak{E})$ and $\alpha \in E_{\mathfrak{E}}^\pm$:

$$S \xrightarrow{\alpha}_{\mathfrak{E}} S' \Leftrightarrow i_{\mathfrak{E}}(S) \xrightarrow{i(\alpha)}_{split(\mathfrak{E})} i_{\mathfrak{E}}(S').$$

Hence $\alpha_1 \cdots \alpha_n \in (E_{\mathfrak{E}}^\pm)^*$ (actually $(E_{\mathfrak{E}}, l_{\mathfrak{E}}, \sigma)$ with $\sigma = \alpha_1 \cdots \alpha_n$) is a pre-ST-trace of an event structure $\mathfrak{E}$ iff $i(\alpha_1) \cdots i(\alpha_n) \in E_{split(\mathfrak{E})}^*$ is a pre-trace of $split(\mathfrak{E})$. Furthermore two pre-ST-traces $\alpha_1 \cdots \alpha_n$ and $\beta_1 \cdots \beta_m$ of $\mathfrak{E}$ are isomorphic iff $i(\alpha_1) \cdots i(\alpha_n)$ and $i(\beta_1) \cdots i(\beta_m)$ are <-isomorphic. Thus $<-Traces(split(\mathfrak{E}))$ is derivable from $ST-Traces(\mathfrak{E})$ and vice versa. From this the first statement of the proposition follows.

As for the second statement, let $\mathfrak{E}, \mathfrak{F} \in \mathbb{E}$.

$$\begin{aligned} \mathcal{R}(\mathfrak{E}, \mathfrak{F}) &= \{((C, P), (D, Q), f) \in \mathcal{S}(\mathfrak{E}) \times \mathcal{S}(\mathfrak{F}) \times \mathcal{P}(E_{\mathfrak{E}} \times E_{\mathfrak{F}}) \mid \\ f: C \rightarrow D \text{ is a bijection, satisfying } l_{\mathfrak{F}}(f(e)) &= l_{\mathfrak{E}}(e) \text{ and } f(P) = Q\}. \end{aligned}$$

For (S, T, f) and $(S', T', f') \in \mathcal{R}(\mathfrak{E}, \mathfrak{F})$ write $(S, T, f) \rightarrow (S', T', f')$ if $S \rightarrow_{\mathfrak{E}} S'$, $T \rightarrow_{\mathfrak{F}} T'$, and $f' \upharpoonright C = f$.

$$\begin{aligned} \mathcal{R}_{split}(\mathfrak{E}, \mathfrak{F}) &= \{((C, D, f) \in \mathcal{C}(split(\mathfrak{E})) \times \mathcal{C}(split(\mathfrak{F})) \times \mathcal{P}(E_{split(\mathfrak{E})} \times E_{split(\mathfrak{F})}) \mid \\ f: C \rightarrow D \text{ is a bijection, satisfying } l_{\mathfrak{F}}(f(e)) &= l_{\mathfrak{E}}(e) \text{ and } f(e) <_{\mathfrak{F}} f(e') \\ \Leftrightarrow e <_{\mathfrak{E}} e'\}. \end{aligned}$$

For (C, D, f) and $(C', D', f') \in \mathcal{R}_{split}(\mathfrak{E}, \mathfrak{F})$ write $(C, D, f) \rightarrow (C', D', f')$ if $C \rightarrow_{split(\mathfrak{E})} C'$, $D \rightarrow_{split(\mathfrak{F})} D'$, and $f' \upharpoonright C = f$. Define $i: \mathcal{R}(\mathfrak{E}, \mathfrak{F}) \rightarrow \mathcal{R}_{split}(\mathfrak{E}, \mathfrak{F})$ by $i(S, T, f) = (i_{\mathfrak{E}}(S), i_{\mathfrak{F}}(T), i(f))$ where $i_{\mathfrak{E}}$ and $i_{\mathfrak{F}}$ are the bijections from Proposition 4 and $i(f): i_{\mathfrak{E}}(S) \rightarrow i_{\mathfrak{F}}(T)$ is defined by $i(f)(e, a^+) = (f(e), a^+)$ and $i(f)(e, a^-) = (f(e), a^-)$. Now it is not difficult to establish that i is a bijection, satisfying

$$(S, T, f) \rightarrow (S', T', f') \Leftrightarrow i(S, T, f) \rightarrow i(S', T', f').$$

From this it follows that $R \subseteq \mathcal{R}(\mathfrak{E}, \mathfrak{F})$ is an ST-bisimulation between $\mathfrak{E}$ and $\mathfrak{F}$ iff $i(R) = \{i(S, T, f) \mid (S, T, f) \in R\} \subseteq \mathcal{R}_{split}(\mathfrak{E}, \mathfrak{F})$ is a $<$ -bisimulation between $split(\mathfrak{E})$ and $split(\mathfrak{F})$. $\square$

PROPOSITION 5.4: *For all equivalences $\approx_1$ and $\approx_2$ on $\mathbb{E}$ defined so far, the formula*


$$\forall \mathfrak{E}, \mathfrak{F} \in \mathbb{E}: \mathfrak{E} \approx_1 \mathfrak{F} \Rightarrow \mathfrak{E} \approx_2 \mathfrak{F}$$

holds iff there is a path $\approx_1 \rightarrow \dots \rightarrow \approx_2$ in Figure 2.

PROOF: In order to prove the announced implications, it suffices to restrict attention to the ones corresponding with an arrow $\approx_1 \rightarrow \approx_2$ in Figure 2. Five of them are dealt with in Proposition 3 already. In order to prove the implications $\approx_{2t} \rightarrow \approx_{st}$ and $\approx_{2b} \rightarrow \approx_{sb}$, consider, for $\mathfrak{E} \in \mathbb{E}$, the mapping $j: \mathcal{C}(\mathfrak{E}) \rightarrow \mathcal{S}(\mathfrak{E})$ defined by $j(C) = (C, C)$. Note that j is a well-defined injection with $range(j) = \{(C, P) \in \mathcal{S}(\mathfrak{E}) \mid C = P\}$. Now for $C \in \mathcal{C}(\mathfrak{E})$, A a multiset over act , and $a_1 \dots a_n \in Act$ an arbitrary enumeration of A , it is easily obtained that

$$\exists C': C \xrightarrow{A}_{\mathfrak{E}} C' \wedge j(C') = (S, T) \Leftrightarrow j(C) \xrightarrow{a_1^+}_{\mathfrak{E}} \dots \xrightarrow{a_n^+}_{\mathfrak{E}} \xrightarrow{a_1^-}_{\mathfrak{E}} \dots \xrightarrow{a_n^-}_{\mathfrak{E}} (S, T).$$

From this the required implications follow immediately. In order to prove the remaining six implications, first consider the implications between equivalences on $\mathbb{E}(Act^{\pm})$ displayed in Figure 8. These implications follow immediately from the definitions. The proofs in Chapter VI that $\approx_{pt}$ and $\approx_h$ are preserved under refinement can be trivially extended to a setting with Λ -refinements for any labelling set Λ . So it follows that

FIGURE 8. Some semantic equivalences on $\mathbb{E}(\text{Act}^\pm)$

$$\mathcal{E} \approx_{pt} \mathcal{F} \Rightarrow \text{split}(\mathcal{E}) \approx_{pt} \text{split}(\mathcal{F}) \text{ and } \mathcal{E} \approx_h \mathcal{F} \Rightarrow \text{split}(\mathcal{E}) \approx_h \text{split}(\mathcal{F}).$$

Now the remaining six implications on $\mathbb{E}(\text{Act})$ follow from Propositions 5.1 and 5.3.

In order to prove the absence of other implications, it suffices to provide counterexamples against $\approx_{pt} \rightarrow \approx_{ib}$, $\approx_{ib} \rightarrow \approx_{st}$, $\approx_{sb} \rightarrow \approx_{2t}$, $\approx_{2b} \rightarrow \approx_{STt}$ and $\approx_{STb} \rightarrow \approx_{pt}$. The first two counterexamples were given already in Section 3. For the third counterexample consider the two event structures of Figure 5. In Section 3 it was established already that they are step bisimulation equivalent. Furthermore they are not split trace equivalent, since $a^+ c^+ a^- b^+ c^- b^-$ is a split trace of the first one but not of the second one.

FIGURE 9. ST-bisimulation equivalent but not pomset trace equivalent
(A variant of Example 7.1.2.a.ii of [91]).

The fourth counterexample will be provided in [94]. For the last counterexample consider the two systems represented in Figure 9. Both systems perform the actions a and b exactly once. In the first system these actions can only be independent, whereas in the second one b can be executed either dependent or independent of a . The difference between the two systems does not occur before (and unless) they reach a state where the execution of a is completed and the execution of b is not yet begun. However, in this state both systems have exactly the same future, consisting of exactly one occurrence of b . Hence they are identified in ST-bisimulation semantics. On the other hand the pomset $a \rightarrow b$ is a pomset trace of the second system, but not of the first. So the two systems are not pomset trace equivalent. This example also shows that ST-semantics does not respect causality. $\square$

6. THE REFINEMENT THEOREMS

Finally I will prove the announced refinement theorems for ST-semantics. In VAN GLABBEK & VAANDRAGER [94] it will be shown that such a theorem does not hold for split semantics.

THEOREM: *Let $\mathfrak{E}, \mathfrak{F} \in \mathbb{E}$ and r be a refinement. Then*

$$\mathfrak{E} \approx_{STb} \mathfrak{F} \Rightarrow r(\mathfrak{E}) \approx_{STb} r(\mathfrak{F}).$$

PROOF: Let $R \subseteq \mathcal{S}(E_{\mathfrak{E}}) \times \mathcal{S}(E_{\mathfrak{F}}) \times \mathcal{P}(E_{\mathfrak{E}} \times E_{\mathfrak{F}})$ be an ST-bisimulation between $\mathfrak{E}$ and $\mathfrak{F}$. Define the relation $\tilde{R}$ by:

$$\tilde{R} = \{((\tilde{C}, \tilde{P}), (\tilde{D}, \tilde{Q}), \tilde{f}) \in \mathcal{S}(E_{r(\mathfrak{E})}) \times \mathcal{S}(E_{r(\mathfrak{F})}) \times \mathcal{P}(E_{r(\mathfrak{E})} \times E_{r(\mathfrak{F})}) \mid$$

$$\exists ((C, P), (D, Q), f) \in R \text{ such that } r^{-1}(\tilde{C}, \tilde{P}) = (C, P), r^{-1}(\tilde{D}, \tilde{Q}) = (D, Q)$$

$$\text{and } \tilde{f}: \tilde{C} \rightarrow \tilde{D} \text{ is a bijection, satisfying } \tilde{f}(e, e') = (f(e), e') \text{ and } \tilde{f}(\tilde{P}) = \tilde{Q} \}.$$

I show that $\tilde{R}$ is an ST-bisimulation between $r(\mathfrak{E})$ and $r(\mathfrak{F})$.

- i. $((\emptyset, \emptyset), (\emptyset, \emptyset), \emptyset) \in \tilde{R}$ since $((\emptyset, \emptyset), (\emptyset, \emptyset), \emptyset) \in R$.
- ii. Suppose $((\tilde{C}, \tilde{P}), (\tilde{D}, \tilde{Q}), \tilde{f}) \in \tilde{R}$. Take $((C, P), (D, Q), f) \in R$ such that $r^{-1}(\tilde{C}, \tilde{P}) = (C, P)$, $r^{-1}(\tilde{D}, \tilde{Q}) = (D, Q)$ and $\tilde{f}: \tilde{C} \rightarrow \tilde{D}$ is a bijection, satisfying $\tilde{f}(e, e') = (f(e), e')$ and $\tilde{f}(\tilde{P}) = \tilde{Q}$. Now three things have to be established:

1. $\tilde{f}: \tilde{C} \rightarrow \tilde{D}$ is a bijection, satisfying $l_{r(\mathfrak{F})}(\tilde{f}(e, e')) = l_{r(\mathfrak{E})}(e, e')$ and $\tilde{f}(\tilde{P}) = \tilde{Q}$.
2. $(\tilde{C}, \tilde{P}) \rightarrow_{r(\mathfrak{E})} (\tilde{C}', \tilde{P}') \Rightarrow \exists \tilde{D}', \tilde{Q}', \tilde{f}'$ with $\tilde{f}' \upharpoonright \tilde{C} = \tilde{f}$,
 $(\tilde{D}, \tilde{Q}) \rightarrow_{r(\mathfrak{F})} (\tilde{D}', \tilde{Q}')$ and $((\tilde{C}', \tilde{P}'), (\tilde{D}', \tilde{Q}'), \tilde{f}') \in \tilde{R}$.
3. $(\tilde{D}, \tilde{Q}) \rightarrow_{r(\mathfrak{F})} (\tilde{D}', \tilde{Q}') \Rightarrow \exists \tilde{C}', \tilde{P}', \tilde{f}'$ with $\tilde{f}' \upharpoonright \tilde{C} = \tilde{f}$,
 $(\tilde{C}, \tilde{P}) \rightarrow_{r(\mathfrak{E})} (\tilde{C}', \tilde{P}')$ and $((\tilde{C}', \tilde{P}'), (\tilde{D}', \tilde{Q}'), \tilde{f}') \in \tilde{R}$

ad 1. By construction $\tilde{f}: \tilde{C} \rightarrow \tilde{D}$ is a bijection, satisfying $\tilde{f}(\tilde{P}) = \tilde{Q}$.

$$\begin{aligned} \text{Moreover } l_{r(\mathfrak{F})}(\tilde{f}(e, e')) &= l_{r(\mathfrak{F})}(f(e), e') = l_{r(l_{\mathfrak{F}}(f(e)))}(e') = \\ &= l_{r(l_{\mathfrak{E}}(e))}(e') = l_{r(\mathfrak{E})}(e, e'). \end{aligned}$$

ad 2. Suppose $(\tilde{C}, \tilde{P}) \rightarrow_{r(\mathfrak{E})} (\tilde{C}', \tilde{P}')$, i.e. $(\tilde{C}', \tilde{P}') \in \mathcal{S}(r(\mathfrak{E}))$,
 $\tilde{C} \subseteq \tilde{C}'$ and $\tilde{P} \subseteq \tilde{P}'$.

Let $(C', P') = r^{-1}(\tilde{C}', \tilde{P}')$. Using Lemma 4.ii,

$(C, P) \rightarrow_{\mathfrak{S}} (C', P')$. Since R is an ST-bisimulation, $\exists D', Q', f'$ with $(D, Q) \rightarrow_{\mathfrak{S}} (D', Q')$, $((C', P'), (D', Q'), f') \in R$ and $f' \upharpoonright C = f$.

$$\begin{aligned} \text{Let } \tilde{D}' &= \{(f'(e), e') \mid (e, e') \in \tilde{C}'\}, \\ \tilde{Q}' &= \{(f'(e), e') \mid (e, e') \in \tilde{P}'\} \text{ and} \\ \tilde{f}' &= \{((e, e'), (f'(e), e')) \mid (e, e') \in \tilde{C}'\}. \end{aligned}$$

For $e \in pr_1(\tilde{C}')$ let

$$C_e = \{e' \mid (e, e') \in \tilde{C}'\} \text{ and } P_e = \{e' \mid (e, e') \in \tilde{P}'\};$$

for $d \in pr_1(\tilde{D}')$ let

$$D_d = \{e' \mid (d, e') \in \tilde{D}'\} \text{ and } Q_d = \{e' \mid (d, e') \in \tilde{Q}'\}.$$

Remark that $Q_{f'(e)} = \{e' \mid (f'(e), e') \in \tilde{Q}'\} = \{e' \mid (e, e') \in \tilde{P}'\} = P_e$ and similarly $D_{f'(e)} = C_e$.

I prove that $(\tilde{D}, \tilde{Q}) \rightarrow_{r(\mathfrak{S})} (\tilde{D}', \tilde{Q}')$,

$$((\tilde{C}', \tilde{P}'), (\tilde{D}', \tilde{Q}'), \tilde{f}') \in \tilde{R} \text{ and } \tilde{f}' \upharpoonright \tilde{C}' = \tilde{f}.$$

I start with proving that $(\tilde{D}', \tilde{Q}') \in \mathbb{S}(r(\mathfrak{S}))$.

$$pr_1(\tilde{D}') = \{f'(e) \mid e \in pr_1(\tilde{C}')\} = f'(C') = D' \quad (1)$$

$$\text{so } \tilde{D}' = \{(d, e') \mid d \in D', e' \in D_d\},$$

$$\tilde{Q}' = \{(d, e') \mid d \in D', e' \in Q_d\}.$$

Using Lemma 4.i, it is then sufficient to show that

$$\begin{aligned} (D', Q') &\text{ is an ST-configuration of } \mathfrak{F}, \\ (D_d, Q_d) &\text{ is an ST-configuration of } r(l_{\mathfrak{S}}(d)) \text{ for } d \in D', \\ Q_d &= E_{r(l_{\mathfrak{S}}(d))} \text{ iff } d \in Q'. \end{aligned} \quad (2)$$

The first requirement is already implicit in $(D, Q) \rightarrow_{\mathfrak{S}} (D', Q')$.

Since $D' = f'(C')$ one may substitute $f'(e)$ for d and $e \in C'$ for $d \in D'$ in the remaining two requirements.

Since $D_{f'(e)} = C_e$, $Q_{f'(e)} = P_e$, $l_{\mathfrak{S}}(f'(e)) = l_{\mathfrak{S}}(e)$ and $Q' = f'(P')$ they reduce to

$$\begin{aligned} (C_e, P_e) &\text{ is an ST-configuration of } r(l_{\mathfrak{S}}(e)) \text{ for } e \in C' \text{ and} \\ P_e &= E_{r(l_{\mathfrak{S}}(e))} \text{ iff } e \in P'. \end{aligned}$$

These follow from Lemma 4.i, using that $(\tilde{C}', \tilde{P}') \in \mathbb{S}(r(\mathfrak{S}))$

$$\text{and } r^{-1}(\tilde{C}', \tilde{P}') = (C', P').$$

Hence $(\tilde{D}', \tilde{Q}') \in \mathbb{S}(r(\mathfrak{S}))$.

Now (1) and (2) above say that $D' = pr_1(\tilde{D}')$ and

$Q' = \{d \in D' \mid Q_d = E_{r(l_{\mathfrak{S}}(d))}\}$. Hence $r^{-1}(\tilde{D}', \tilde{Q}') = (D', Q')$.

It follows that $((\tilde{C}', \tilde{P}'), (\tilde{D}', \tilde{Q}'), \tilde{f}') \in \tilde{R}$.

Finally $\tilde{f}' \upharpoonright \tilde{C} = \tilde{f}$, $\tilde{D} \subseteq \tilde{D}'$ and $\tilde{Q} \subseteq \tilde{Q}'$ by construction, using that $f' \upharpoonright C = f$. With $(\tilde{D}', \tilde{Q}') \in \mathbb{S}(r(\mathfrak{F}))$, it follows that $(\tilde{D}, \tilde{Q}) \rightarrow_{r(\mathfrak{F})} (\tilde{D}', \tilde{Q}')$.

ad 3. By symmetry. □

THEOREM: Let $\mathfrak{E}, \mathfrak{F} \in \mathbb{E}$ and r be a refinement. Then

$$\mathfrak{E} \approx_{STl} \mathfrak{F} \Rightarrow r(\mathfrak{E}) \approx_{STl} r(\mathfrak{F}).$$

PROOF: It suffices to prove $\mathfrak{E} \lesssim_{STl} \mathfrak{F} \Rightarrow r(\mathfrak{E}) \lesssim_{STl} r(\mathfrak{F})$, so let $\mathfrak{E}, \mathfrak{F} \in \mathbb{E}$ with $\mathfrak{E} \lesssim_{STl} \mathfrak{F}$ and let r be a refinement. Suppose in $r(\mathfrak{E})$ there is a chain of ST-configurations

$$(\emptyset, \emptyset) \rightarrow_{r(\mathfrak{E})} (\tilde{C}_1, \tilde{P}_1) \rightarrow_{r(\mathfrak{E})} \cdots \rightarrow_{r(\mathfrak{E})} (\tilde{C}_n, \tilde{P}_n).$$

By Lemma 4.ii there is a chain of ST-configurations

$$(\emptyset, \emptyset) \rightarrow_{\mathfrak{E}} (C_1, P_1) \rightarrow_{\mathfrak{E}} \cdots \rightarrow_{\mathfrak{E}} (C_n, P_n)$$

in $\mathfrak{E}$ with $(C_i, P_i) = r^{-1}(\tilde{C}_i, \tilde{P}_i)$ for $i = 1, \dots, n$. Hence there must be a chain

$$(\emptyset, \emptyset) \rightarrow_{\mathfrak{F}} (D_1, Q_1) \rightarrow_{\mathfrak{F}} \cdots \rightarrow_{\mathfrak{F}} (D_n, Q_n)$$

in $\mathfrak{F}$ and a bijection $f: C_n \rightarrow D_n$, satisfying $l_{\mathfrak{F}}(f(e)) = l_{\mathfrak{E}}(e)$, $f(C_i) = D_i$ and $f(P_i) = Q_i$ for $i = 1, \dots, n$.

Let $\tilde{D}_i = \{(f(e), e') \mid (e, e') \in \tilde{C}_i\}$,

$\tilde{Q}_i = \{(f(e), e') \mid (e, e') \in \tilde{P}_i\}$ and

$\tilde{f} = \{((e, e'), (f(e), e')) \mid (e, e') \in \tilde{C}_n\}$.

It remains to be shown that

$$(\emptyset, \emptyset) \rightarrow_{r(\mathfrak{F})} (\tilde{D}_1, \tilde{Q}_1) \rightarrow_{r(\mathfrak{F})} \cdots \rightarrow_{r(\mathfrak{F})} (\tilde{D}_n, \tilde{Q}_n)$$

is a chain of ST-configurations in $r(\mathfrak{F})$ and $\tilde{f}: \tilde{C}_n \rightarrow \tilde{D}_n$ is a bijection satisfying $l_{r(\mathfrak{F})}(\tilde{f}(e, e')) = l_{r(\mathfrak{E})}(e, e')$, $\tilde{f}(\tilde{C}_i) = \tilde{D}_i$ and $\tilde{f}(\tilde{P}_i) = \tilde{Q}_i$ for $i = 1, \dots, n$.

The only nontrivial part of this consist of proving that $(\tilde{D}_i, \tilde{Q}_i) \in \mathbb{S}(r(\mathfrak{F}))$ for $i = 1, \dots, n$. This goes exactly as in the previous proof. □

CONCLUDING REMARKS

In this chapter ten semantic equivalences for concurrent systems are defined on a domain of labelled event structures, and their interdependencies are classified as indicated in Figure 2 of the introduction. It has been established - in [46, 84] and [94] respectively - that interleaving, step and split equivalences are strictly based on action atomicity. In particular, the owl example of [94] shows that no equivalence that can be localized between split bisimulation and interleaving trace equivalence is preserved under refinement of actions. On the other hand it has been shown - in [46] and in the previous chapter - that the two partial order equivalences of Figure 2 *are* preserved under action refinement and thus need not to be based on action atomicity. Now this chapter added that also ST-trace and ST-bisimulation equivalence are preserved under refinement. So the borderline is between split and ST-semantics.

It should be remarked that at all places where split semantics was used before it was studied for a restricted class of concurrent systems (Petri nets without autoconcurrency in [91], a subset of CCS in [5, 109] and deterministic event structures in [172]) on which it coincides with ST-semantics. The examples of [94] suggest that outside such a class, split semantics is not an interesting notion. The reason for mentioning it in this chapter is that it seems to be a natural simplification of ST-semantics and in order to indicate that for the purposes of this book this simplification should not be made.

The refinement operation considered in this chapter replaced actions by finite, conflict-free, non-empty event structures. As remarked earlier, a generalization to infinite refinements, leaving all definitions the same, is incompatible with the principle of finite causes: try to refine a in

$$a \longrightarrow b \qquad \text{by} \qquad a_1 \longrightarrow a_2 \longrightarrow a_3 \longrightarrow \cdots .$$

If one would drop this principle, there are (at least) two possibilities of interpreting event structures: events which have an infinite set of causes can happen in a finite time, or they can not. The last interpretation is slightly simpler to grasp, more common, and compatible with the view of this chapter, in which the behaviour of concurrent systems - together with all semantic equivalences - is explained in terms of finite configurations (or ST-configurations) only. Using this interpretation any 'generalized' event structure can be transformed in an ordinary prime event structure satisfying the principle of finite causes, by removing all events that have infinitely many causes. A transformed event structure and its original are equivalent with respect to all equivalences of Figure 2. On the domain of 'generalized' event structures one may drop the restriction that refinements need to be finite, and all theorems and definitions of this chapter remain valid. In fact also all proofs remain valid, since (except in the proof of Proposition 1.i) the principle of finite causes is never used. However, it can be argued that infinite refinements change the behaviour of the considered systems in a way that cannot be explained by a change in the level

of abstraction at which processes are regarded: consider a system performing the actions a and b one time each, where the occurrence of b is dependent of the occurrence of a (as depicted above); after replacement of a by an infinite event structure, b cannot happen any more; it occurs in no (finite) configuration. Finally notice that it is also possible to describe this type of refinement on the domain of prime event structures satisfying the principle of finite causes, by adding to the definition of refinement that after refinement in the sense of Section 1, events with infinitely many causes should be left out.

A generalization to refinements containing conflicts can be obtained analogously as the above generalization to infinite refinements, but is technically more complicated. On the domain of prime event structures used in this chapter, refinements with conflicts are incompatible with the principle of conflict heredity: try to replace a in

$$a \longrightarrow b \quad \text{by} \quad a_1 \dots\dots a_2.$$

This problem has been solved in Chapter IV by moving to a more general form of event structures where the axiom of conflict heredity is dropped, e.g. flow event structures [38,40]. On flow event structures we could define a refinement operator for any function $r:Act \rightarrow \mathbb{E} - \{0\}$, thus allowing both infinite refinements and refinements with conflicts. I expect that after this generalization all my theorems remain valid. Each flow event structure is equivalent to a prime event structure (with respect to any of the equivalences of Figure 2). Hence an alternative solution consists of appending to the definition of refinement some transformation that turns the refined event structure into an equivalent prime event structure.

Contrary to the previous generalization, a generalization of the refinement operator to *forgetful refinements*, where replacing actions by the empty event structure is allowed, does not seem very natural. Such refinements can drastically change the behaviour of concurrent systems and can not be explained by a change in the level of abstraction at which these systems are regarded (Chapter IV). Moreover, unlike the refinement theorems for partial order semantics (Chapter VI) the refinement theorem for ST-bisimulation semantics does not hold for forgetful refinements, as is demonstrated by the following counterexample.



The two event structures above are ST-bisimulation equivalent. However, after replacing a by the empty event structure, the resulting event structures (below)

are not ST-bisimulation equivalent.



The refinement theorems for ST-semantics show that in case preservation under refinement is required, it is not necessary to employ partial order semantics.¹ From this the natural question arises if it is necessary to employ at least ST-semantics, i.e. if any equivalence finer than a given interleaving equivalence that is preserved under refinement is also finer than some ST-equivalence. Let $\approx_x$ be an equivalence on $\mathbb{E}$. Define $\approx_{rx}$ by


$$\mathcal{E} \approx_{rx} \mathcal{F} \text{ iff for all refinements } r: Act \rightarrow \mathbb{E} - \{0\} \text{ one has } r(\mathcal{E}) \approx_x r(\mathcal{F}).$$

Then, $\approx_{rx}$ is finer than $\approx_x$ and preserved under refinement. Moreover it is coarser than any other equivalence with these properties. In other words, $\approx_{rx}$ is *fully abstract* with respect to $\approx_x$ and refinement. Of course the definition above is parametrized by the concept of refinement. Let $\approx_{rx}$ be defined under reference to general refinements $r: Act \rightarrow \mathbb{E} - \{0\}$ (using flow event structures); and let $\approx_{r'x}$ be defined under reference to refinements as defined in Section 1 of this chapter. Then I conjecture that $\approx_{STb}$ coincides with $\approx_{rib}$, i.e. ST-

1. WALTER VOGLER [176, 178, 179] argues that partial order semantics *is* necessary for action refinement. The question at stake is of course the definition of partial order semantics. In order for partial order semantics to be necessary, both ST-bisimulation and ST-trace equivalence (and all other equivalences that are preserved under action refinement) should be classified as partial order equivalences. Vogler rightly rejects the definition of partial order semantics as a semantics that makes use of partial orders - say, in order to describe system runs. Such a definition would make no sense, as the information contained in the partial orders can easily be coded in other ways. However, he argues that at least ST-trace semantics "deserve[s] to be called a partial order semantics. [ST-traces] are just sequential representations of special partial orders", namely of *interval orders*. It seems that the implicit definition requires that a partial order semantics *can* be conveniently described in terms of partial orders. This definition has two drawbacks. First of all Vogler has to amend it by stating that a "real" partial order semantics should also use partial orders that do not correspond to sequences or step sequences". The question arises why not to require the full generality of partial orders, and exclude interval orders as well. Secondly, it appears that ST-bisimulation is still more conveniently described without using partial orders, thus undermining the argument. Maybe that a semantics that is finer than a partial order semantics is automatically counted as a partial order semantics, even if it has no employ for partial orders.

In view of this I prefer the definition of partial order semantics given in this text (page 252, line 8), namely that a partial order semantics is a semantics that uses partial orders for representing causal dependence, or uses anything that can be rephrased in terms of partial orders, as long as it is for representing causal dependence. This definition is more of a conceptual than of a technical nature. Moreover, I claim that it better corresponds with the historical use of the term. Under this definition ST-trace and ST-bisimulation semantics do not count as partial order semantics.

bisimulation equivalence is fully abstract with respect to interleaving bisimulation equivalence and action refinement, and also $\approx_{STl}$ coincides with $\approx_{rit}$. To be more precise, let r_c be the refinement that replaces actions $a \in Act$ by



Then I think that $\mathcal{G} \approx_{STb} \mathcal{F} \Leftrightarrow r_c(\mathcal{G}) \approx_{ib} r_c(\mathcal{F})$ and likewise $\mathcal{G} \approx_{STl} \mathcal{F} \Leftrightarrow r_c(\mathcal{G}) \approx_{it} r_c(\mathcal{F})$, from which the conjecture follows. Furthermore, together with Walter Vogler I observed that for finite event structures $\approx_{STb}$ even coincides with $\approx_{r'ib}$. On the other hand $\approx_{r'it}$ is strictly coarser than $\approx_{rit}$, as follows from an example in LARSEN [126], see also [94].

In VOGLER [175] a ‘failures semantics based on interval semiwords’ was presented that can be regarded as the ST-version of failure semantics. He proved that this semantics is preserved under refinement of actions and also established that it is fully abstract with respect to interleaving failure semantics and refinement (allowing refinements with conflicts, but without initial and final parallism, see Section 4 of Chapter IV). The same results he obtained for ST-trace semantics. Topics for further research include

- generalizing the refinement theorems to a setting with infinite refinements and refinements with conflicts, as in Chapter IV.
- defining ‘syntactic refinement’ (replacing action symbols by terms in process expressions) on process specification languages, investigating the interaction with communication, proving syntactic refinement theorems and establishing the correspondence with ‘semantic refinement’, as employed in this chapter (cf. [5, 6, 91, 126, 141]),¹
- proving the full abstraction results conjectured above,²
- proving refinement theorems and full abstraction results for the ST-versions of decorated trace semantics,³
- and generalizing the entire theory to a setting with silent actions, or τ -moves (possibly combining the notions of branching bisimulation (for refinement of systems with silent actions) (Chapter III) and ST-bisimulation or history preserving bisimulation (for refinement of non-sequential systems)).⁴

PRATT [154] and CASTELLANO, DE MICHELIS & POMELLO [46] use the issue of action atomicity as an argument for using partial order semantics instead of interleaving semantics. This chapter shows that it is not necessary to employ partial order semantics if one does not want to assume action atomicity; ST-semantics turns out to be sufficient. In VAN GLABBEK & VAANDRAGER [91] we introduced the (related) criterion of *real-time consistency*. A semantics is real-time consistent if it does not identify systems with a different real-time behaviour. Of course interleaving semantics are not real-time consistent, but

again the criterion did not force us to consider partial order semantics: also for this purpose ST-bisimulation semantics turned out to be sufficient. Therefore the question remains whether or not there exists a convincing testing scenario, or some natural operator, that reveals the full distinguishing power of partial order semantics.⁵

1. This problem has been solved in GOLTZ, GORRIERI & RENSINK [97].
2. This has been done in VOGLER [179].
3. For failure semantics this has been done in VOGLER [175] in a setting of Petri nets, see also JATEGAONKAR & MEYER [117], and for a variant of trace semantics, in the absence of autoconcurrency, modelling a process as a set of semiwords, in NIELSEN, ENGBERG & LARSEN [141] and LARSEN [126].
4. This has been done in VOGLER [179] too, but using weak bisimulation rather than branching bisimulation. Vogler proves that weak ST-bisimulation, pomset weak ST-bisimulation and history preserving weak ST-bisimulation equivalence are preserved under action refinement. The last result is also established in DEVILLERS [72] in the setting of Petri nets. When choosing the ‘weak’ mode in dealing with τ -steps, the ‘ST’ is obviously needed in all three results: it effectively converts weak bisimulation into delay bisimulation (as defined Chapter III), which, like branching bisimulation, is preserved under action refinement. This prevents one from running into the counterexample of Section 6 in Chapter III. Independent of the choice of the ‘weak’ mode in dealing with internal actions, the ‘ST’ is necessary in the first two results; witness the counterexamples of Chapter VI. This leaves open the question whether history preserving branching bisimulation and history preserving delay bisimulation (without ‘ST’; the definitions should be obvious) are preserved under action refinement. In a joint paper with Sophie Pinchinat I will show that, surprisingly, they are not. In the presence of τ ’s ST-semantics appears to be necessary.
5. In a linear time setting (i.e. neglecting the branching structure of processes) this question has been addressed in PLOTKIN & PRATT [151].

References

References

1. S. ABRAMSKY, "Observation equivalence as a testing equivalence," *Theoretical Computer Science* **53**, pp. 225-241 (1987).
2. S. ABRAMSKY & S. VICKERS, "Quantaes, observational logic and process semantics," *Mathematical Structures in Computer Science* **3**, pp. 161-227 (1993).
3. L. ACETO, "Action refinement in process algebras," Ph.D. Thesis, University of Sussex (1991). Report No. 3/91
4. L. ACETO, R. DE NICOLA & A. FANTECHI, "Testing equivalences for event structures," pp. 1-20 in *Proceedings Advanced School on Mathematical Models for the Semantics of Parallelism, 1986*, ed. M. VENTURINI ZILLI, ED., Springer-Verlag (1987).
5. L. ACETO & M. HENNESSY, "Towards action-refinement in process algebras," *Information and Computation* **103**(2), pp. 204-269 (1993).
6. L. ACETO & M. HENNESSY, "Adding action refinement to a finite process algebra," *Information and Computation*, to appear (1994).
7. P. ACZEL, *Non-well-founded sets*, CSLI Lecture Notes No.14, Stanford University (1988).
8. G.J. AKKERMAN & J.C.M. BAETEN, "Term rewriting analysis in process algebra," Report P9006, Programming Research Group, University of Amsterdam (1990).
9. P. AMERICA, "Definition of the programming language POOL-T," ESPRIT project 415, Doc. Nr. 91, Philips Research Laboratories, Eindhoven (1985).
10. D. AUSTRY & G. BOUDOL, "Algèbre de processus et synchronisations," *Theoretical Computer Science* **30**(1), pp. 91-131 (1984).

11. J.C.M. BAETEN & J.A. BERGSTRA, "Global renaming operators in concrete process algebra," *I&C* **78**(3), pp. 205-245 (1988).
12. J.C.M. BAETEN, J.A. BERGSTRA & J.W. KLOP, "Conditional axioms and α/β calculus in process algebra," pp. 53-75 in *Formal Description of Programming Concepts - III, Proceedings of the third IFIP WG 2.2 working conference, Ebberup 1986*, ed. M. WIRSING, ED., North-Holland, Amsterdam (1987).
13. J.C.M. BAETEN, J.A. BERGSTRA & J.W. KLOP, "On the consistency of Koomen's fair abstraction rule," *Theoretical Computer Science* **51**(1/2), pp. 129-176 (1987).
14. J.C.M. BAETEN, J.A. BERGSTRA & J.W. KLOP, "Ready-trace semantics for concrete process algebra with the priority operator," *The Computer Journal* **30**(6), pp. 498-506 (1987).
15. J.C.M. BAETEN & R.J. VAN GLABBEEK, "Another look at abstraction in process algebra," pp. 84-94 in *Proceedings ICALP 87, Karlsruhe*, ed. TH. OTTMAN, ED., Springer-Verlag (1987).
16. J.C.M. BAETEN & R.J. VAN GLABBEEK, "Merge and termination in process algebra," pp. 153-172 in *Proceedings 7th Conference on Foundations of Software Technology & Theoretical Computer Science, Pune, India*, ed. K.V. NORI, ED., Springer-Verlag (1987).
17. J.C.M. BAETEN & W.P. WEIJLAND, *Process algebra*, Cambridge University Press (1990).
18. J.W. DE BAKKER, J.A. BERGSTRA, J.W. KLOP & J.-J.CH. MEYER, "Linear time and branching time semantics for recursion with merge," *Theoretical Computer Science* **34**, pp. 135-156 (1984).
19. J.W. DE BAKKER, J.N. KOK, J.-J.CH. MEYER, E.-R. OLDEROG & J.I. ZUCKER, "Contrasting themes in the semantics of imperative concurrency," pp. 51-121 in *Current trends in concurrency*, ed. J.W. DE BAKKER, W.-P. DE ROEVER & G. ROZENBERG, EDS., Springer-Verlag (1986).
20. J.W. DE BAKKER & J.I. ZUCKER, "Processes and the denotational semantics of concurrency," *I&C* **54**(1/2), pp. 70-120 (1982).
21. J. VAN BENTHEM, J. VAN EIJCK & V. STEBLETSOVA, "Modal logic, transition systems and processes," Report CR-R9321, CWI, Amsterdam (1993).
22. J.A. BERGSTRA, "A process creation mechanism in process algebra," Logic Group Preprint Series Nr. 2, CIF, State University of Utrecht (1985). in: *Applications of process algebra*, (J.C.M. Baeten, ed.), Cambridge University Press, 1990, pp. 81-88
23. J.A. BERGSTRA, J. HEERING & P. KLINT, "Module algebra," *JACM* **37**(2), pp. 335-372 (1990).
24. J.A. BERGSTRA & J.W. KLOP, "The algebra of recursively defined processes and the algebra of regular processes," pp. 82-95 in *Proceedings ICALP 84, Antwerp*, ed. J. PAREDAENS, ED., Springer-Verlag (1984).
25. J.A. BERGSTRA & J.W. KLOP, "Process algebra for synchronous communication," *I&C* **60**(1/3), pp. 109-137 (1984).

26. J.A. BERGSTRA & J.W. KLOP, "Algebra of communicating processes with abstraction," *Theoretical Computer Science* 37(1), pp. 77-121 (1985).
27. J.A. BERGSTRA & J.W. KLOP, "A complete inference system for regular processes with silent moves," pp. 21-81 in *Proceedings Logic Colloquium 1986*, ed. F.R. DRAKE & J.K. TRUSS, EDS., North Holland, Hull (1988). also appeared as: Report CS-R8420, Centrum voor Wiskunde en Informatica, Amsterdam 1984
28. J.A. BERGSTRA & J.W. KLOP, "Process theory based on bisimulation semantics," pp. 50-122 in *Proceedings REX School/Workshop on Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency*, Noordwijkerhout, ed. J.W. DE BAKKER, W.-P. DE ROEVER & G. ROZENBERG, EDS., Springer-Verlag (1989).
29. J.A. BERGSTRA, J.W. KLOP & E.-R. OLDEROG, "Failures without chaos: a new process semantics for fair abstraction," pp. 77-103 in *Formal Description of Programming Concepts - III, Proceedings of the third IFIP WG 2.2 working conference, Ebberup 1986*, ed. M. WIRSING, ED., North-Holland, Amsterdam (1987).
30. J.A. BERGSTRA, J.W. KLOP & E.-R. OLDEROG, "Readies and failures in the algebra of communicating processes," *SIAM Journal on Computing* 17(6), pp. 1134-1177 (1988).
31. J.A. BERGSTRA & J. TIURYN, "Process algebra semantics for queues," *Fund. Inf.* X, pp. 213-224, also appeared as: MC Report IW 241, Amsterdam 1983 (1987).
32. J.A. BERGSTRA & J.V. TUCKER, "Top down design and the algebra of communicating processes," *SCP* 5(2), pp. 171-199 (1985).
33. E. BEST, R. DEVILLERS, A. KIEHN & L. POMELLO, "Concurrent Bisimulations in Petri nets," *Acta Informatica* 28, pp. 231-264 (1991).
34. B. BLOOM, "Structural operational semantics for weak bisimulations," Technical Report TR 93-1373, Cornell University, Ithaca, New York (1993).
35. B. BLOOM, S. ISTRAIL & A.R. MEYER, "Bisimulation can't be traced: preliminary report," pp. 229-239 in *Conference Record of the 15th Annual ACM Symposium on Principles of Programming Languages (POPL), San Diego, California* (1988).
36. A. BOUALI, "Weak and branching bisimulation in FCTOOL," RAPPORTS DE RECHERCHE N° 1575, INRIA, CENTRE SOPHIA-ANTIPOLIS, VALBONNE CEDEX (1992).
37. G. BOUDOL, "Atomic actions (note)," *Bulletin of the EATCS* 38, pp. 136-144 (1989).
38. G. BOUDOL, "Flow event structures and flow nets," pp. 62-95 in *Semantics of Systems of Concurrent Processes (I. Guessarian, ed.), Proceedings of the LITP Spring School on Theoretical Computer Science, La Roche Posay, France 1990*, Springer-Verlag (1990).
39. G. BOUDOL & I. CASTELLANI, "On the semantics of concurrency: partial orders and transition systems," pp. 123-137 in *Proceedings TAPSOFT 87, Vol. I*, ed. H. EHRLIG, R. KOWALSKI, G. LEVI & U. MONTANARI, EDS.,

- Springer-Verlag (1987).
40. G. BOUDOL & I. CASTELLANI, "Permutation of transitions: an event structure semantics for CCS & SCCS," pp. 411-427 in *Proceedings REX School/Workshop on Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency*, Noordwijkerhout, ed. J.W. DE BAKKER, W.-P. DE ROEVER & G. ROZENBERG, EDS., Springer-Verlag (1989).
 41. W. BRAUWER, R. GOLD & W. VOGLER, "A survey of behaviour and equivalence preserving refinements of Petri nets," pp. 1-46 in *Advances in Petri nets 1990*, ed. G. ROZENBERG, ED., Springer-Verlag (1991).
 42. S.D. BROOKES, C.A.R. HOARE & A.W. ROSCOE, "A theory of communicating sequential processes," *JACM* **31**(3), pp. 560-599 (1984).
 43. M.C. BROWNE, E.M. CLARKE & O. GRUMBERG, "Characterizing finite Kripke structures in propositional temporal logic," *Theoretical Computer Science* **59**(1,2), pp. 115-131 (1988).
 44. M. BROY, "Views on queues," *Sci. Comp. Prog.* **11**, pp. 65-86 (1988).
 45. N. BUSI, R.J. VAN GLABBEK & R. GORRIERI, *Axiomatising ST-Bisimulation Equivalence*, Available by anonymous ftp from Boole.stanford.edu, 1993.
 46. L. CASTELLANO, G. DE MICHELIS & L. POMELLO, "Concurrency vs Interleaving: an instructive example," *Bulletin of the EATCS* **31**, pp. 12-15 (1987).
 47. F. CHERIEF, "Contributions à la sémantique du parallélisme: Bisimulations pour le raffinement et le vrai parallélisme," Pd.D. Thesis, Univ. Grenoble (1992).
 48. F. CHERIEF & PH. SCHNOEBELEN, " τ -Bisimulations and full abstraction for refinement of actions," *Information processing letters* **40**, pp. 219-222 (1991).
 49. S. CHRISTENSEN, H. HÜTTEL & C. STIRLING, "Bisimulation equivalence is decidable for all context-free processes," pp. 138-147 in *Proceedings CONCUR '92, Stony Brook, NY, USA, August 1992*, ed. W.R. CLEAVELAND, ED., Springer-Verlag (1992).
 50. E.M. CLARKE & E.A. EMERSON, "Design and synthesis of synchronization skeletons using branching-time temporal logic," pp. 52-71 in *Proceedings of the Workshop on Logic of Programs, Yorktown Heights*, ed. D. KOZEN, ED., Springer-Verlag (1981).
 51. I. CZAJA, R.J. VAN GLABBEK & U. GOLTZ, "Interleaving Semantics and Action Refinement with Atomic Choice," pp. 89-107 in *Advances in Petri Nets 1992*, ed. G. ROZENBERG, ED., Springer-Verlag (1992).
 52. PH. DARONDEAU, "An enlarged definition and complete axiomatisation of observational congruence of finite processes," pp. 47-62 in *Proceedings international symposium on programming: 5th colloquium, Aarhus*, ed. M. DEZANI-CIANCAGLINI & U. MONTANARI, EDS., Springer-Verlag (1982).
 53. P. DARONDEAU & P. DEGANO, "Causal trees," in *Proceedings ICALP 89, Stresa, July 1988*, ed. G. AUSIELLO, M. DEZANI-CIANCAGLINI & S. RONCHI DELLA ROCCA, EDS. (1989).

54. PH. DARONDEAU & P. DEGRANO, "Causal trees: Interleaving + Causality," pp. 239-255 in *Semantics of Systems of Concurrent Processes (I. Guessarian, ed.), Proceedings of the LITP Spring School on Theoretical Computer Science, La Roche Posay, France 1990*, Springer-Verlag (1990).
55. PH. DARONDEAU & P. DEGRANO, "Event structures, causal trees, and refinements," in *Proceedings Mathematical Foundations of Computer Science 1990 (MFCS)*, ed. B. ROVAN, ED., Springer-Verlag (1990).
56. PH. DARONDEAU & P. DEGRANO, "About semantic action refinement," *Fundamenta Informaticae* **XIV**, pp. 221-234 (1991).
57. PH. DARONDEAU & P. DEGRANO, "Refinement of actions in event structures and causal trees," *Theoretical Computer Science* **118**, pp. 21-48 (1993).
58. P. DEGRANO, R. DE NICOLA & U. MONTANARI, "Observational equivalences for concurrency models," pp. 105-129 in *Formal Description of Programming Concepts - III, Proceedings of the third IFIP WG 2.2 working conference, Ebberup 1986*, ed. M. WIRSING, ED., Elsevier Science Publishers B.V. (North Holland) (1987).
59. P. DEGRANO, R. DE NICOLA & U. MONTANARI, "A distributed operational semantics for CCS based on condition/event systems," *Acta Informatica* **26**(1/2), pp. 59-91 (1988).
60. P. DEGRANO, R. DE NICOLA & U. MONTANARI, "Partial orderings descriptions and observations of nondeterministic concurrent processes," pp. 438-466 in *Proceedings REX School/Workshop on Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency, Noordwijkerhout, The Netherlands, May/June 1988*, ed. J.W. DE BAKKER, W.-P. DE ROEVER & G. ROZENBERG, EDS., Springer-Verlag (1989).
61. P. DEGRANO & R. GORRIERI, "Atomic refinement for process description languages," pp. 121-130 in *Proceedings Mathematical Foundations of Computer Science 1991 (MFCS)*, ed. A. TARLECKI, ED., Springer-Verlag (1991).
62. P. DEGRANO & R. GORRIERI, "An operational definition of action refinement," Technical Report 28, Dipartimento di Informatica, Università di Pisa (1992).
63. R. DE NICOLA, "Extensional equivalences for transition systems," *Acta Informatica* **24**, pp. 211-237 (1987).
64. R. DE NICOLA & M. HENNESSY, "Testing equivalences for processes," *Theoretical Computer Science* **34**, pp. 83-133 (1984).
65. R. DE NICOLA, P. INVERARDI & M. NESI, "Using axiomatic presentation of behavioural equivalences for manipulating CCS specifications," pp. 54-67 in *Automatic Verification Methods for Finite State Systems*, ed. J. SIFAKIS, ED., Springer-Verlag (1990).
66. R. DE NICOLA, U. MONTANARI & F.W. VAANDRAGER, "Back and forth bisimulations," pp. 152-165 in *Proceedings CONCUR 90, Amsterdam*, ed. J.C.M. BAETEN & J.W. KLOP, EDS., Springer-Verlag (1990). also appeared as CWI Report CS-R9021
67. R. DE NICOLA & F.W. VAANDRAGER, "Three logics for branching

- bisimulation (extended abstract)," pp. 118-129 in *Proceedings 5th Annual Symposium on Logic in Computer Science (LICS 90)*, Philadelphia, USA, IEEE Computer Society Press, Los Alamitos, CA (1990). full version appeared as CWI Report CS-R9012, Amsterdam 1990
68. R. DE NICOLA & F.W. VAANDRAGER, "Action versus state based logics for transition systems," pp. 407-419 in *Semantics of Systems of Concurrent Processes (I. Guessarian, ed.)*, *Proceedings of the LITP Spring School on Theoretical Computer Science*, La Roche Posay, France 1990, Springer-Verlag (1990).
 69. T. DENVIR, W. HARWOOD, M. JACKSON & M. RAY, *The analysis of concurrent systems*, *Proceedings of a Tutorial and Workshop*, Cambridge University 1983, Springer-Verlag (1985).
 70. J. DESEL & A. MERCERON, "Vicinity respecting net morphisms," pp. 115-138 in *Proceedings of the 10th International Conference on Petri Nets*, Bonn (1989).
 71. R. DEVILLERS, "On the definition of a bisimulation notion based on partial words," *Petri Net Newsletter* **29**, pp. 16-19, Gesellschaft für Informatik, Bonn (1988).
 72. R. DEVILLERS, "Maximality preserving bisimulation," *Theoretical Computer Science* **102**, pp. 165-183 (1992).
 73. R. DEVILLERS, "Maximality preservation and the ST-idea for action refinement," pp. 108-151 in *Advances in Petri Nets 1992*, ed. G. ROZENBERG, ED., Springer-Verlag (1992).
 74. E.A. EMERSON & J.Y. HALPERN, "'Sometimes' and 'Not Never' revisited: on branching time versus linear time temporal logic," *JACM* **33**(1), pp. 151-178 (1986).
 75. H.J. GENRICH & E. STANKIEWICZ-WIECHNO, "A dictionary of some basic notions of Petri nets," in *Advanced course on general net theory of processes and systems*, Hamburg 1979, ed. W. BRAUER, ED., Springer-Verlag (1980).
 76. J.L. GISCHER, "The equational theory of pomsets," *Theoretical Computer Science* **61**, pp. 199-224 (1988).
 77. R.J. VAN GLABBEEK, "Bounded nondeterminism and the approximation induction principle in process algebra," pp. 336-347 in *Proceedings STACS 87*, ed. F.J. BRANDENBURG, G. VIDAL-NAQUET & M. WIRSING, EDS., Springer-Verlag (1987).
 78. R.J. VAN GLABBEEK, "The refinement theorem for ST-bisimulation semantics," pp. 27-52 in *Programming Concepts and Methods (M. Broy & C.B. Jones, eds.)*, *Proceedings IFIP Working Group 2.2/2.3 Working Conference, Sea of Galilee, Israel 1990*, Elsevier Science Publishers B.V. (North Holland) (1990).
 79. R.J. VAN GLABBEEK, "The linear time - branching time spectrum," pp. 278-297 in *Proceedings CONCUR 90*, Amsterdam, ed. J.C.M. BAETEN & J.W. KLOP, EDS., Springer-Verlag (1990).
 80. R.J. VAN GLABBEEK, "A complete axiomatization for branching bisimulation congruence of finite-state behaviours," pp. 473-484 in *Proceedings*

- Mathematical Foundations of Computer Science 1993 (MFCS)*, Gdansk, Poland, August/September 1993, ed. A.M. BORZYSZKOWSKI & S. SOKOŁOWSKI, EDS., Springer-Verlag (1993).
81. R.J. VAN GLABBEK, *The linear time - branching time spectrum II; the semantics of sequential systems with silent moves*, Preliminary version available from boole.stanford.edu; Extended abstract in Proceedings CONCUR'93, Hildesheim, Germany, August 1993 (E. Best, ed), LNCS 715, Springer-Verlag, pp. 66-81 1993.
 82. R.J. VAN GLABBEK, "What is branching time and why to use it?," Report No. STAN-CS-93-1486, Stanford University (1993). Available by anonymous ftp from Boole.stanford.edu
 83. R.J. VAN GLABBEK, "Full Abstraction in Structural Operational Semantics (extended abstract)," pp. 77-84 in *Proceedings of the Third International Conference on Algebraic Methodology and Software Technology (AMAST'93)*, The Netherlands, June 1993, ed. M. NIVAT, C. RATTRAY, T. RUS & G. SCOLLO, EDS., Springer-Verlag (1993). Available by anonymous ftp from Boole.stanford.edu
 84. R.J. VAN GLABBEK & U. GOLTZ, "Equivalence notions for concurrent systems and refinement of actions," Arbeitspapiere der GMD 366, Sankt Augustin (1989). extended abstract appeared in: Proceedings 14th Symposium on Mathematical Foundations of Computer Science (MFCS), Porąbka-Kozubnik, Poland, August/September 1989 (A. Kreczmar & G. Mirkowska, eds.), LNCS 379, Springer-Verlag, pp. 237-248
 85. R.J. VAN GLABBEK & U. GOLTZ, "Partial order semantics for refinement of actions - neither necessary nor always sufficient but appropriate when used with care," *Bulletin of the EATCS* **38**, pp. 154-163 (1989).
 86. R.J. VAN GLABBEK & U. GOLTZ, "Refinement of actions in causality based models," pp. 267-300 in *Proceedings of the REX Workshop on Stepwise Refinement of Distributed Systems: Models, Formalism, Correctness* (J.W. de Bakker, W.-P. de Roever & G. Rozenberg, eds.), Mook, The Netherlands 1989, Springer-Verlag (1990).
 87. R.J. VAN GLABBEK & U. GOLTZ, "Equivalences and refinement," pp. 309-333 in *Semantics of Systems of Concurrent Processes* (I. Guessarian, ed.), *Proceedings of the LITP Spring School on Theoretical Computer Science*, La Roche Posay, France 1990, Springer-Verlag (1990). This is an extended and updated version of [84]
 88. R.J. VAN GLABBEK & U. GOLTZ, "A Deadlock-sensitive Congruence for Action Refinement," SFB-Bericht Nr. 342/23/90 A, Institut für Informatik der Technischen Universität München (1990).
 89. R.J. VAN GLABBEK & J.J.M.M. RUTTEN, "The processes of De Bakker and Zucker represent bisimulation equivalence classes," pp. 243-246 in *J.W. de Bakker, 25 jaar semantiek, liber amicorum* (1989).
 90. R.J. VAN GLABBEK, S.A. SMOLKA, B. STEFFEN & C.M.N. TOFTS, "Reactive, generative, and stratified models of probabilistic processes," pp. 130-141 in *Proceedings 5th Annual Symposium on Logic in Computer*

- Science (LICS 90)*, Philadelphia, USA, IEEE Computer Society Press, Los Alamitos, CA (1990).
91. R.J. VAN GLABBEK & F.W. VAANDRAGER, "Petri net models for algebraic theories of concurrency," pp. 224-242 in *Proceedings PARLE conference, Eindhoven, Vol. II (Parallel Languages)*, ed. J.W. DE BAKKER, A.J. NIJMAN & P.C. TRELEAVEN, EDS., Springer-Verlag (1987).
 92. R.J. VAN GLABBEK & F.W. VAANDRAGER, "Modular specifications in process algebra - with curious queues," Report CS-R8821, Centrum voor Wiskunde en Informatica, Amsterdam (1988). extended abstract in: *Algebraic Methods: Theory, Tools and Applications* (M. Wirsing & J.A. Bergstra, eds.), LNCS 394, Springer-Verlag, 1989, pp. 465-506. A revision of part of this paper is:
 93. R.J. VAN GLABBEK & F.W. VAANDRAGER, "Modular specification of process algebras," *TCS* 113(2), pp. 293-348 (1993).
 94. R.J. VAN GLABBEK & F.W. VAANDRAGER, *The difference between splitting in n and $n + 1$* , in preparation, 1994.
 95. R.J. VAN GLABBEK & W.P. WEIJLAND, "Refinement in branching time semantics," Report CS-R8922, Centrum voor Wiskunde en Informatica, Amsterdam (1989). appeared in: J.W. de Bakker, 25 jaar semantiek, *liber amicorum*, Centrum voor Wiskunde en Informatica, Amsterdam 1989, pp. 247-252, and also in: *Proceedings AMAST Conference*, May 1989, Iowa, USA, pp. 197-201, and finally in:
 96. R.J. VAN GLABBEK & W.P. WEIJLAND, "Branching time and abstraction in bisimulation semantics," Report TUM-I9052, SFB-Bericht Nr. 342/29/90 A, Institut für Informatik, Technische Universität München (1990). under revision for JACM. Extended abstract in: *Information Processing 89* (G.X. Ritter, ed.), *Proceedings of the IFIP 11th World Computer Congress*, San Fransisco 1989, Elsevier Science Publishers B.V. (North Holland), 1989, pp. 613-618
 97. U. GOLTZ, R. GORRIERI & A. RENSINK, "On syntactic and semantic action refinement," in *Proc. Theoretical Aspects of Computer Science, Sendai, Japan*, Springer-Verlag (1994). to appear
 98. R. GORRIERI, "Refinement, atomicity and transactions for process description languages," Ph.D. Thesis, TD-2/91, University of Pisa (1991).
 99. R. GORRIERI & C. LANEVE, "Split and ST Bisimulation Semantics," *Information and Computation*, to appear (1994).
 100. R. GORRIERI, S. MARCHETTI & U. MONTANARI, " A^2 CCS: a simple extension of CCS for handling atomic actions," pp. 258-270 in *Proceedings CAAP 88, Nancy, France*, ed. M. DAUGHET & M. NIVAT, EDS., Springer-Verlag (1988).
 101. J. GRABOWSKI, "On partial languages," *Fundamenta Informaticae* IV(2), pp. 427-498 (1981).
 102. S. GRAF & J. SIFAKIS, "Readiness semantics for regular processes with silent actions," pp. 115-125 in *Proceedings ICALP 87, Karlsruhe*, ed. TH. OTTMAN, ED., Springer-Verlag (1987).

103. E.P. GRIBOMONT, "Stepwise refinement and concurrency: a small exercise," pp. 219-238 in *Mathematics of program construction*, ed. J.L.A. VAN DE SNEPSCHOUT, ED., Springer-Verlag (1989).
104. J.F. GROOTE, "A new strategy for proving ω -completeness with applications in process algebra," pp. 314-331 in *Proceedings CONCUR 90, Amsterdam*, ed. J.C.M. BAETEN & J.W. KLOP, EDS., Springer-Verlag (1990).
105. J.F. GROOTE & H. HÜTTEL, "Undecidable equivalences for basic process algebra," *Information and Computation*, to appear (1994).
106. J.F. GROOTE & F.W. VAANDRAGER, "An efficient algorithm for branching bisimulation and stuttering equivalence," pp. 626-638 in *Proceedings ICALP 90, Warwick*, ed. M.S. PATERSON, ED., Springer-Verlag (1990).
107. J.F. GROOTE & F.W. VAANDRAGER, "Structured operational semantics and bisimulation as a congruence," *I&C* **100**(2), pp. 202-260 (1992).
108. M. HENNESSY, "Acceptance trees," *JACM* **32**(4), pp. 896-928 (1985).
109. M. HENNESSY, "Axiomatising finite concurrent processes," *SIAM Journal on Computing* **17**(5), pp. 997-1017 (1988).
110. M. HENNESSY, "Concurrent testing of processes (extended abstract)," pp. 94-107 in *Proceedings CONCUR '92, Stony Brook, NY, USA, August 1992*, ed. W.R. CLEAVELAND, ED., Springer-Verlag (1992).
111. M. HENNESSY & R. MILNER, "On observing nondeterminism and concurrency," pp. 299-309 in *Proceedings ICALP 80*, ed. J.W. DE BAKKER & J. VAN LEEUWEN, EDS., Springer-Verlag (1980). this is a preliminary version of:
112. M. HENNESSY & R. MILNER, "Algebraic laws for nondeterminism and concurrency," *JACM* **32**(1), pp. 137-161 (1985).
113. M. HENNESSY & G.D. PLOTKIN, "A term model for CCS," pp. 261-274 in *Proceedings 9th Symposium on Mathematical Foundations of Computer Science (MFCS)*, ed. P. DEMBIŃSKI, ED., Springer-Verlag (1980).
114. C.A.R. HOARE, "Communicating sequential processes," *Communications of the ACM* **21**(8), pp. 666-677 (1978).
115. C.A.R. HOARE, "Communicating sequential processes," pp. 229-254 in *On the construction of programs - an advanced course*, ed. R.M. MCKEAG & A.M. MACNAGHTEN, EDS., Cambridge University Press (1980).
116. C.A.R. HOARE, *Communicating sequential processes*, Prentice-Hall International (1985).
117. L. JATEGAONKAR & A. MEYER, "Testing equivalence for Petri nets with action refinement," pp. 17-31 in *Proceedings CONCUR '92, Stony Brook, NY, USA, August 1992*, ed. W.R. CLEAVELAND, ED., Springer-Verlag (1992).
118. HE JIFENG & C.A.R. HOARE, "Algebraic specification and proof of a distributed recovery algorithm," *Distributed Computing* **2**(1), pp. 1-12 (1987).
119. B. JONSSON & J. PARROW, "Deciding bisimulation equivalences for a class of non-finite-state programs," *Information and Computation* **107**(2), pp. 272-302 (1993).

120. P.C. KANELLAKIS & S.A. SMOLKA, "CCS expressions, finite state processes, and three problems of equivalence," *Information and Computation* **86**, pp. 43-68 (1990).
121. J.K. KENNAWAY, "Formal semantics of nondeterminism and parallelism," Ph.D. Thesis, University of Oxford (1981).
122. W. KORCZYŃSKI, "An algebraic characterization of concurrent systems," *Fundamenta Informaticae* **11**(2), pp. 171-194 (1988).
123. C.P.J. KOYMANS & J.C. MULDER, "A modular approach to protocol verification using process algebra," Logic Group Preprint Series Nr. 6, CIF, State University of Utrecht (1986). in: Applications of process algebra, (J.C.M. Baeten, ed.), Cambridge University Press, 1990, pp. 261-306
124. L. LAMPORT, "What good is temporal logic?," pp. 657-668 in *Information Processing 83*, ed. R.E. MASON, ED., Elsevier Science Publishers B.V. (North Holland) (1983).
125. L. LAMPORT, "On interprocess communication," *Distributed Computing* **1**, pp. 77-101 (1986).
126. K.S. LARSEN, "A fully abstract model for a process algebra with refinements," Master Thesis, Aarhus University, Denmark (1988).
127. K.G. LARSEN & R. MILNER, "A complete protocol verification using relativized bisimulation," pp. 126-135 in *Proceedings ICALP 87, Karlsruhe*, ed. TH. OTTMANN, ED., Springer-Verlag (1987).
128. K.G. LARSEN & A. SKOU, "Bisimulation through probabilistic testing," *I&C* **94**(1), pp. 1-28, a preliminary report appeared in: Conference Record of the 16th Annual ACM Symposium on Principles of Programming Languages (POPL), Austin, Texas, ACM Press, New York 1989, pp. 344-352 (1991).
129. S. MAUW, "An algebraic specification of process algebra, including two examples," Report FVI 87-06, Dept. of Computer Science, University of Amsterdam (1987). extended abstract in: Algebraic Methods: Theory, Tools and Applications (M. Wirsing & J.A. Bergstra, eds.), LNCS 394, Springer-Verlag, pp. 507-554
130. S. MAUW & G.J. VELTINK, "A process specification formalism," *Fundamenta Informaticae* **XIII**, pp. 85-139 (1990).
131. J. MESEGUER & U. MONTANARI, "Petri nets are monoids: a new algebraic foundation for net theory," pp. 155-164 in *Proceedings 3th Annual Symposium on Logic in Computer Science (LICS 88)*, Edinburgh, IEEE Computer Society Press, Washington (1988).
132. A.R. MEYER, "Report on the 5th international workshop on the semantics of programming languages in Bad Honnef," *Bulletin of the EATCS* **27**, pp. 83-84 (1985).
133. G.J. MILNE, "CIRCAL and the representation of communication, concurrency, and time," *TOPLAS* **7**(2), pp. 270-298 (1985).
134. R. MILNER, *A Calculus of Communicating Systems*, Springer-Verlag (1980).
135. R. MILNER, "Modal characterisation of observable machine behaviour,"

- pp. 25-34 in *Proceedings CAAP 81*, ed. G. ASTESIANO & C. BOHM, EDS., Springer-Verlag (1981).
136. R. MILNER, "Calculi for synchrony and asynchrony," *Theoretical Computer Science* **25**, pp. 267-310 (1983).
 137. R. MILNER, "Lectures on a Calculus for Communicating Systems," pp. 197-220 in *Seminar on Concurrency*, ed. S.D. BROOKES, A.W. ROSCOE & G. WINSKEL, EDS., Springer-Verlag (1985).
 138. R. MILNER, *Communication and concurrency*, Prentice-Hall International (1989).
 139. F. MOLLER, "The importance of the left merge operator in process algebras," pp. 752-764 in *Proceedings ICALP 90, Warwick*, ed. M.S. PATERSON, ED., Springer-Verlag (1990).
 140. J.D. MONK, *Mathematical logic*, Springer-Verlag (1976).
 141. M. NIELSEN, U. ENGBERG & K.S. LARSEN, "Fully abstract models for a process language with refinement," pp. 523-548 in *Proceedings REX School/Workshop on Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency, Noordwijkerhout*, ed. J.W. DE BAKKER, W.-P. DE ROEVER & G. ROZENBERG, EDS., Springer-Verlag (1989).
 142. M. NIELSEN, G.D. PLOTKIN & G. WINSKEL, "Petri nets, event structures and domains, part I," *Theoretical Computer Science* **13**(1), pp. 85-108 (1981).
 143. E.-R. OLDEROG, U. GOLTZ & R.J. VAN GLABBEK, "Combining compositionality and concurrency, summary of a GMD-workshop, Königswinter, March 1988," *Arbeitspapiere der GMD* 320, Sankt Augustin (1988).
 144. E.-R. OLDEROG & C.A.R. HOARE, "Specification-oriented semantics for communicating processes," *Acta Informatica* **23**, pp. 9-66 (1986).
 145. D.M.R. PARK, "Concurrency and automata on infinite sequences," pp. 167-183 in *Proceedings 5th GI Conference*, ed. P. DEUSSEN, ED., Springer-Verlag (1981).
 146. J. PARROW & P. SJÖDIN, "Multiway synchronization verified with coupled simulation," pp. 518-533 in *Proceedings CONCUR '92, Stony Brook, NY, USA, August 1992*, ed. W.R. CLEVELAND, ED., Springer-Verlag (1992).
 147. G.L. PETERSON, "Myths about the mutual exclusion problem," *Information Processing Letters* **12**(3), pp. 115-116 (1981).
 148. J. PETERSON, "Petri nets," *ACM Computing Surveys* **9**(3), pp. 223-252 (1977).
 149. C.A. PETRI, "Non-sequential processes," *Interner Bericht ISF-77-05*, Gesellschaft für Mathematik und Datenverarbeitung, Sankt Augustin (1977).
 150. I.C.C. PHILLIPS, "Refusal testing," *Theoretical Computer Science* **50**, pp. 241-284 (1987).
 151. G.D. PLOTKIN & V.R. PRATT, *Teams can see pomsets*, Manuscript available from Boole.stanford.edu, 1990.

152. A. PNUELI, "Linear and branching structures in the semantics and logics of reactive systems," pp. 15-32 in *Proceedings ICALP 85, Nafplion*, ed. W. BRAUER, ED., Springer-Verlag (1985).
153. L. POMELLO, "Some equivalence notions for concurrent systems. An overview," pp. 381-400 in *Advances in Petri Nets 1985*, ed. G. ROZENBERG, ED., Springer-Verlag (1986).
154. V.R. PRATT, "Modelling concurrency with partial orders," *International Journal of Parallel Programming* **15**(1), pp. 33-71 (1986).
155. A. RABINOVICH & B.A. TRAKHTENBROT, "Behavior Structures and Nets," *Fundamenta Informaticae* **11**(4), pp. 357-404 (1988).
156. W. REISIG, *Petri nets - an introduction*, Springer-Verlag (1985).
157. W. REISIG, "Petri nets in software engineering," pp. 63-96 in *Petri Nets: Applications and Relationships to Other Models of Concurrency, Advances in Petri Nets 1986, Part II; Proceedings of an Advanced Course, Bad Honnef, September 1986*, ed. W. BRAUER, W. REISIG & G. ROZENBERG, EDS., Springer-Verlag (1987).
158. A. RENSINK, "Models and methods for action refinement," Ph.D. Thesis, University of Twente (1993).
159. W.C. ROUNDS & S.D. BROOKES, "Possible futures, acceptances, refusals and communicating processes," pp. 140-149 in *Proceedings 22nd Annual Symposium on Foundations of Computer Science, Nashville, USA 1981*, IEEE, New York (1981).
160. M. SANDERSON, "," Ph.D. Thesis, University of Edinburgh (1982).
161. D.T. SANNELLA & A. TARLECKI, "Toward formal development of programs from algebraic specifications: implementations revisited," *Acta Informatica* **25**, pp. 233-281 (1988).
162. D.T. SANNELLA & M. WIRSING, "A kernel language for algebraic specification and implementation (extended abstract)," pp. 413-427 in *Proceedings International Conference on Foundations of Computation Theory, Borgholm*, ed. M. KARPINSKI, ED. (1983). long version: Report CSR-131-83, Dept. of Computer Science, Univ. of Edinburgh, 1983
163. PH. SCHNOEBELEN, "Experiments on Processes with Backtracking," pp. 80-94 in *Proceedings CONCUR 91, Amsterdam*, ed. J.C.M. BAETEN & J.F. GROOTE, EDS., Springer-Verlag (1991).
164. I. SUZUKI & T. MURATA, "A method for stepwise refinement and abstraction of Petri nets," *Journal of Computer and System Sciences* **27**(1), pp. 51-76 (1983).
165. D.A. TAUBNER & W. VOGLER, "Step failure semantics and a complete proof system," *Acta Informatica* **27**, pp. 125-156 (1989).
166. I. ULIDOWSKI, "Equivalences on observable processes," pp. 148-159 in *Proceedings 7th Annual Symposium on Logic in Computer Science (LICS 92), Santa Cruz, California*, IEEE Computer Society Press, Los Alamitos, CA.
167. F.W. VAANDRAGER, "Verification of two communication protocols by means of process algebra," Report CS-R8608, Centrum voor Wiskunde en Informatica, Amsterdam (1986).

168. F.W. VAANDRAGER, "Process algebra semantics of POOL," Report CS-R8629, Centrum voor Wiskunde en Informatica, Amsterdam (1986). in: Applications of process algebra, (J.C.M. Baeten, ed.), Cambridge University Press, 1990, pp. 173-236
169. F.W. VAANDRAGER, "Some observations on redundancy in a context," Report CS-R8812, Centrum voor Wiskunde en Informatica, Amsterdam (1988). in: Applications of process algebra, (J.C.M. Baeten, ed.), Cambridge University Press, 1990, pp. 237-260
170. F.W. VAANDRAGER, "An explicit representation of equivalence classes of the history preserving bisimulation," Unpublished manuscript, quoted in [54] (1989).
171. F.W. VAANDRAGER, "Algebraic techniques for concurrency and their application," Ph.D. Thesis, University of Amsterdam (1990).
172. F.W. VAANDRAGER, "Determinism $\rightarrow$ (event structure isomorphism = step sequence equivalence)," *Theoretical Computer Science* **79**, pp. 275-294 (1991).
173. R. VALETTE, "Analysis of Petri nets by stepwise refinements," *Journal of Computer and System Sciences* **18**, pp. 35-46 (1979).
174. W. VOGLER, "Behaviour preserving refinements of Petri nets," pp. 82-93 in *Proceedings 12th International Workshop on Graph-Theoretic Concepts in Computer Science, Bernried, 1986*, ed. G.TINHOFFER & G. SCHMIDT, EDS., Springer-Verlag (1987).
175. W. VOGLER, "Failures semantics based on interval semiwords is a congruence for refinement," *Distributed Computing* **4**, pp. 139-162 (1991).
176. W. VOGLER, "Is partial order semantics necessary for action refinement?," SFB-Bericht Nr. 342/1/91 A, Institut für Informatik der Technischen Universität München (1991).
177. W. VOGLER, "Deciding history preserving bisimilarity," pp. 495-505 in *Proceedings ICALP 91*, ed. J. LEACH ALBERT, B. MONIEN & M. RODRÍGUEZ ARTALEJO, EDS., Springer-Verlag (1991).
178. W. VOGLER, *Modular construction and partial order semantics of Petri nets*, Springer-Verlag (1992).
179. W. VOGLER, "Bisimulation and action refinement," *Theoretical Computer Science* **114**, pp. 173-200 (1993).
180. W. VOGLER, "Generalized OM-bisimulation," *Information and Computation*, to appear (1995).
181. D.J. WALKER, "Bisimulation and divergence," *I&C* **85**(2), pp. 202-241 (1990).
182. W.P. WEIJLAND, "Synchrony and asynchrony in process algebra," Ph.D. Thesis, University of Amsterdam (1989).
183. G. WINSKEL, "Event structures," pp. 325-392 in *Petri Nets: Applications and Relationships to Other Models of Concurrency, Advances in Petri Nets 1986, Part II; Proceedings of an Advanced Course, Bad Honnef, September 1986*, ed. W. BRAUER, W. REISIG & G. ROZENBERG, EDS., Springer-Verlag (1987).

CWI TRACTS

- 1 D.H.J. Epema. *Surfaces with canonical hyperplane sections*. 1984.
- 2 J.J. Dijkstra. *Fake topological Hilbert spaces and characterizations of dimension in terms of negligibility*. 1984.
- 3 A.J. van der Schaft. *System theoretic descriptions of physical systems*. 1984.
- 4 J. Koene. *Minimal cost flow in processing networks, a primal approach*. 1984.
- 5 B. Hoogenboom. *Intertwining functions on compact Lie groups*. 1984.
- 6 A.P.W. Böhm. *Dataflow computation*. 1984.
- 7 A. Blokhuis. *Few-distance sets*. 1984.
- 8 M.H. van Hoorn. *Algorithms and approximations for queueing systems*. 1984.
- 9 C.P.J. Koymans. *Models of the lambda calculus*. 1984.
- 10 C.G. van der Laan, N.M. Temme. *Calculation of special functions: the gamma function, the exponential integrals and error-like functions*. 1984.
- 11 N.M. van Dijk. *Controlled Markov processes; time-discretization*. 1984.
- 12 W.H. Hundsdorfer. *The numerical solution of nonlinear stiff initial value problems: an analysis of one step methods*. 1985.
- 13 D. Grune. *On the design of ALEPH*. 1985.
- 14 J.G.F. Thiemann. *Analytic spaces and dynamic programming: a measure theoretic approach*. 1985.
- 15 F.J. van der Linden. *Euclidean rings with two infinite primes*. 1985.
- 16 R.J.P. Groothuizen. *Mixed elliptic-hyperbolic partial differential operators: a case-study in Fourier integral operators*. 1985.
- 17 H.M.M. ten Eikelder. *Symmetries for dynamical and Hamiltonian systems*. 1985.
- 18 A.D.M. Kester. *Some large deviation results in statistics*. 1985.
- 19 T.M.V. Janssen. *Foundations and applications of Montague grammar, part 1: Philosophy, framework, computer science*. 1986.
- 20 B.F. Schriever. *Order dependence*. 1986.
- 21 D.P. van der Vecht. *Inequalities for stopped Brownian motion*. 1986.
- 22 J.C.S.P. van der Woude. *Topological dynamix*. 1986.
- 23 A.F. Monna. *Methods, concepts and ideas in mathematics: aspects of an evolution*. 1986.
- 24 J.C.M. Baeten. *Filters and ultrafilters over definable subsets of admissible ordinals*. 1986.
- 25 A.W.J. Kolen. *Tree network and planar rectilinear location theory*. 1986.
- 26 A.H. Veen. *The misconstrued semicolon: Reconciling imperative languages and dataflow machines*. 1986.
- 27 A.J.M. van Engelen. *Homogeneous zero-dimensional absolute Borel sets*. 1986.
- 28 T.M.V. Janssen. *Foundations and applications of Montague grammar, part 2: Applications to natural language*. 1986.
- 29 H.L. Trentelman. *Almost invariant subspaces and high gain feedback*. 1986.
- 30 A.G. de Kok. *Production-inventory control models: approximations and algorithms*. 1987.
- 31 E.E.M. van Berkum. *Optimal paired comparison designs for factorial experiments*. 1987.
- 32 J.H.J. Einmahl. *Multivariate empirical processes*. 1987.
- 33 O.J. Vrieze. *Stochastic games with finite state and action spaces*. 1987.
- 34 P.H.M. Kersten. *Infinitesimal symmetries: a computational approach*. 1987.
- 35 M.L. Eaton. *Lectures on topics in probability inequalities*. 1987.
- 36 A.H.P. van der Burgh, R.M.M. Mattheij (eds.). *Proceedings of the first international conference on industrial and applied mathematics (ICIAM 87)*. 1987.
- 37 L. Stougie. *Design and analysis of algorithms for stochastic integer programming*. 1987.
- 38 J.B.G. Frenk. *On Banach algebras, renewal measures and regenerative processes*. 1987.
- 39 H.J.M. Peters, O.J. Vrieze (eds.). *Surveys in game theory and related topics*. 1987.
- 40 J.L. Geluk, L. de Haan. *Regular variation, extensions and Tauberian theorems*. 1987.
- 41 Sape J. Mullender (ed.). *The Amoeba distributed operating system: Selected papers 1984-1987*. 1987.
- 42 P.R.J. Asveld, A. Nijholt (eds.). *Essays on concepts, formalisms, and tools*. 1987.
- 43 H.L. Bodlaender. *Distributed computing: structure and complexity*. 1987.
- 44 A.W. van der Vaart. *Statistical estimation in large parameter spaces*. 1988.
- 45 S.A. van de Geer. *Regression analysis and empirical processes*. 1988.
- 46 S.P. Spekrijse. *Multigrid solution of the steady Euler equations*. 1988.
- 47 J.B. Dijkstra. *Analysis of means in some non-standard situations*. 1988.
- 48 F.C. Drost. *Asymptotics for generalized chi-square goodness-of-fit tests*. 1988.
- 49 F.W. Wubs. *Numerical solution of the shallow-water equations*. 1988.
- 50 F. de Kerf. *Asymptotic analysis of a class of perturbed Korteweg-de Vries initial value problems*. 1988.
- 51 P.J.M. van Laarhoven. *Theoretical and computational aspects of simulated annealing*. 1988.
- 52 P.M. van Loon. *Continuous decoupling transformations for linear boundary value problems*. 1988.
- 53 K.C.P. Machielsen. *Numerical solution of optimal control problems with state constraints by sequential quadratic programming in function space*. 1988.
- 54 L.C.R.J. Willenborg. *Computational aspects of survey data processing*. 1988.
- 55 G.J. van der Steen. *A program generator for recognition, parsing and transduction with syntactic patterns*. 1988.
- 56 J.C. Ebergen. *Translating programs into delay-insensitive circuits*. 1989.
- 57 S.M. Verduyn Lunel. *Exponential type calculus for linear delay equations*. 1989.
- 58 M.C.M. de Gunst. *A random model for plant cell population growth*. 1989.
- 59 D. van Dulst. *Characterizations of Banach spaces not containing l^1* . 1989.
- 60 H.E. de Swart. *Vacillation and predictability properties of low-order atmospheric spectral models*. 1989.

- 61 P. de Jong. *Central limit theorems for generalized multilinear forms*. 1989.
- 62 V.J. de Jong. *A specification system for statistical software*. 1989.
- 63 B. Hanzon. *Identifiability, recursive identification and spaces of linear dynamical systems, part I*. 1989.
- 64 B. Hanzon. *Identifiability, recursive identification and spaces of linear dynamical systems, part II*. 1989.
- 65 B.M.M. de Weger. *Algorithms for diophantine equations*. 1989.
- 66 A. Jung. *Cartesian closed categories of domains*. 1989.
- 67 J.W. Polderman. *Adaptive control & identification: Conflict or conflux?*. 1989.
- 68 H.J. Woerdeman. *Matrix and operator extensions*. 1989.
- 69 B.G. Hansen. *Monotonicity properties of infinitely divisible distributions*. 1989.
- 70 J.K. Lenstra, H.C. Tijms, A. Volgenant (eds.). *Twenty-five years of operations research in the Netherlands: Papers dedicated to Gijs de Leve*. 1990.
- 71 P.J.C. Spreij. *Counting process systems. Identification and stochastic realization*. 1990.
- 72 J.F. Kaashoek. *Modeling one dimensional pattern formation by anti-diffusion*. 1990.
- 73 A.M.H. Gerards. *Graphs and polyhedra. Binary spaces and cutting planes*. 1990.
- 74 B. Koren. *Multigrid and defect correction for the steady Navier-Stokes equations. Application to aerodynamics*. 1991.
- 75 M.W.P. Savelsbergh. *Computer aided routing*. 1992.
- 76 O.E. Flippo. *Stability, duality and decomposition in general mathematical programming*. 1991.
- 77 A.J. van Es. *Aspects of nonparametric density estimation*. 1991.
- 78 G.A.P. Kindervater. *Exercises in parallel combinatorial computing*. 1992.
- 79 J.J. Lodder. *Towards a symmetrical theory of generalized functions*. 1991.
- 80 S.A. Smulders. *Control of freeway traffic flow*. 1996.
- 81 P.H.M. America, J.J.M.M. Rutten. *A parallel object-oriented language: design and semantic foundations*. 1992.
- 82 F. Thuijsman. *Optimality and equilibria in stochastic games*. 1992.
- 83 R.J. Kooman. *Convergence properties of recurrence sequences*. 1992.
- 84 A.M. Cohen (ed.). *Computational aspects of Lie group representations and related topics. Proceedings of the 1990 Computational Algebra Seminar at CWI, Amsterdam*. 1991.
- 85 V. de Valk. *One-dependent processes*. 1994.
- 86 J.A. Baars, J.A.M. de Groot. *On topological and linear equivalence of certain function spaces*. 1992.
- 87 A.F. Monna. *The way of mathematics and mathematicians*. 1992.
- 88 E.D. de Goede. *Numerical methods for the three-dimensional shallow water equations*. 1993.
- 89 M. Zwaan. *Moment problems in Hilbert space with applications to magnetic resonance imaging*. 1993.
- 90 C. Vuik. *The solution of a one-dimensional Stefan problem*. 1993.
- 91 E.R. Verheul. *Multimedians in metric and normed spaces*. 1993.
- 92 J.L.M. Maubach. *Iterative methods for non-linear partial differential equations*. 1994.
- 93 A.W. Ambergen. *Statistical uncertainties in posterior probabilities*. 1993.
- 94 P.A. Zegeling. *Moving-grid methods for time-dependent partial differential equations*. 1993.
- 95 M.J.C. van Pul. *Statistical analysis of software reliability models*. 1993.
- 96 J.K. Scholma. *A Lie algebraic study of some integrable systems associated with root systems*. 1993.
- 97 J.L. van den Berg. *Sojourn times in feedback and processor sharing queues*. 1993.
- 98 A.J. Koning. *Stochastic integrals and goodness-of-fit tests*. 1993.
- 99 B.P. Sommeijer. *Parallelism in the numerical integration of initial value problems*. 1993.
- 100 J. Molenaar. *Multigrid methods for semiconductor device simulation*. 1993.
- 101 H.J.C. Huijberts. *Dynamic feedback in nonlinear synthesis problems*. 1994.
- 102 J.A.M. van der Weide. *Stochastic processes and point processes of excursions*. 1994.
- 103 P.W. Hemker, P. Wesseling (eds.). *Contributions to multigrid*. 1994.
- 104 I.J.B.F. Adan. *A compensation approach for queueing problems*. 1994.
- 105 O.J. Boxma, G.M. Koole (eds.). *Performance evaluation of parallel and distributed systems - solution methods. Part 1*. 1994.
- 106 O.J. Boxma, G.M. Koole (eds.). *Performance evaluation of parallel and distributed systems - solution methods. Part 2*. 1994.
- 107 R.A. Trompert. *Local uniform grid refinement for time-dependent partial differential equations*. 1995.
- 108 M.N.M. van Lieshout. *Stochastic geometry models in image analysis and spatial statistics*. 1995.
- 109 R.J. van Glabbeek. *Comparative concurrency semantics and refinement of actions*. 1996.
- 110 W. Vervaat, H. Holwerda (ed.). *Probability and lattices*. 1997.
- 111 I. Helsloot. *Covariant formal group theory and some applications*. 1995.
- 112 R.N. Bol. *Loop checking in logic programming*. 1995.
- 113 G.J.M. Koole. *Stochastic scheduling and dynamic programming*. 1995.
- 114 M.J. van der Laan. *Efficient and inefficient estimation in semiparametric models*. 1995.
- 115 S.C. Borst. *Polling models*. 1996.
- 116 G.D. Otten. *Statistical test limits in quality control*. 1996.
- 117 K.G. Langendoen. *Graph reduction on shared-memory multiprocessors*. 1996.
- 118 W.C.A. Maas. *Nonlinear $\mathcal{H}_\infty$ control: the singular case*. 1996.
- 119 A. Di Bucchanico. *Probabilistic and analytical aspects of the umbral calculus*. 1997.
- 120 M. van Loon. *Numerical methods in smog prediction*. 1997.
- 121 B.J. Wijers. *Nonparametric estimation for a windowed line-segment process*. 1997.
- 122 W.K. Klein Hanefeld, O.J. Vrieze, L.C.M. Kallenberg (editors). *Ten years LNMB - Ph.D. research and graduate courses of the Dutch Network of Operations Research*. 1997.

MATHEMATICAL CENTRE TRACTS

- 1 T. van der Walt. *Fixed and almost fixed points*. 1963.
- 2 A.R. Bloemena. *Sampling from a graph*. 1964.
- 3 G. de Leve. *Generalized Markovian decision processes, part I: model and method*. 1964.
- 4 G. de Leve. *Generalized Markovian decision processes, part II: probabilistic background*. 1964.
- 5 G. de Leve, H.C. Tijms, P.J. Weeda. *Generalized Markovian decision processes, applications*. 1970.
- 6 M.A. Maurice. *Compact ordered spaces*. 1964.
- 7 W.R. van Zwet. *Convex transformations of random variables*. 1964.
- 8 J.A. Zonneveld. *Automatic numerical integration*. 1964.
- 9 P.C. Baayen. *Universal morphisms*. 1964.
- 10 E.M. de Jager. *Applications of distributions in mathematical physics*. 1964.
- 11 A.B. Paalman-de Miranda. *Topological semigroups*. 1964.
- 12 J.A.Th.M. van Berckel, H. Brandt Corstius, R.J. Mokken, A. van Wijngaarden. *Formal properties of newspaper Dutch*. 1965.
- 13 H.A. Lauwerier. *Asymptotic expansions*. 1966, out of print: replaced by MCT 54.
- 14 H.A. Lauwerier. *Calculus of variations in mathematical physics*. 1966.
- 15 R. Doornbos. *Slippage tests*. 1966.
- 16 J.W. de Bakker. *Formal definition of programming languages with an application to the definition of ALGOL 60*. 1967.
- 17 R.P. van de Riet. *Formula manipulation in ALGOL 60, part 1*. 1968.
- 18 R.P. van de Riet. *Formula manipulation in ALGOL 60, part 2*. 1968.
- 19 J. van der Slot. *Some properties related to compactness*. 1968.
- 20 P.J. van der Houwen. *Finite difference methods for solving partial differential equations*. 1968.
- 21 E. Wattel. *The compactness operator in set theory and topology*. 1968.
- 22 T.J. Dekker. *ALGOL 60 procedures in numerical algebra, part 1*. 1968.
- 23 T.J. Dekker, W. Hoffmann. *ALGOL 60 procedures in numerical algebra, part 2*. 1968.
- 24 J.W. de Bakker. *Recursive procedures*. 1971.
- 25 E.R. Paërl. *Representations of the Lorentz group and projective geometry*. 1969.
- 26 European Meeting 1968. *Selected statistical papers, part I*. 1968.
- 27 European Meeting 1968. *Selected statistical papers, part II*. 1968.
- 28 J. Oosterhoff. *Combination of one-sided statistical tests*. 1969.
- 29 J. Verhoeff. *Error detecting decimal codes*. 1969.
- 30 H. Brandt Corstius. *Exercises in computational linguistics*. 1970.
- 31 W. Molenaar. *Approximations to the Poisson, binomial and hypergeometric distribution functions*. 1970.
- 32 L. de Haan. *On regular variation and its application to the weak convergence of sample extremes*. 1970.
- 33 F.W. Steutel. *Preservations of infinite divisibility under mixing and related topics*. 1970.
- 34 I. Juhász, A. Verbeek, N.S. Kroonenberg. *Cardinal functions in topology*. 1971.
- 35 M.H. van Emden. *An analysis of complexity*. 1971.
- 36 J. Grasman. *On the birth of boundary layers*. 1971.
- 37 J.W. de Bakker, G.A. Blaauw, A.J.W. Duijvestijn, E.W. Dijkstra, P.J. van der Houwen, G.A.M. Kamsteeg-Kemper, F.E.J. Kruseman Aretz, W.L. van der Poel, J.P. Schaap-Kruseman, M.V. Wilkes, G. Zoutendijk. *MC-25 Informatica Symposium*. 1971.
- 38 W.A. Verloren van Themaat. *Automatic analysis of Dutch compound words*. 1972.
- 39 H. Bavinck. *Jacobi series and approximation*. 1972.
- 40 H.C. Tijms. *Analysis of (s,S) inventory models*. 1972.
- 41 A. Verbeek. *Superextensions of topological spaces*. 1972.
- 42 W. Vervaat. *Success epochs in Bernoulli trials (with applications in number theory)*. 1972.
- 43 F.H. Ruymgaart. *Asymptotic theory of rank tests for independence*. 1973.
- 44 H. Bart. *Meromorphic operator valued functions*. 1973.
- 45 A.A. Balkema. *Monotone transformations and limit laws*. 1973.
- 46 R.P. van de Riet. *ABC ALGOL, a portable language for formula manipulation systems, part 1: the language*. 1973.
- 47 R.P. van de Riet. *ABC ALGOL, a portable language for formula manipulation systems, part 2: the compiler*. 1973.
- 48 F.E.J. Kruseman Aretz, P.J.W. ten Hagen, H.L. Oudshoorn. *An ALGOL 60 compiler in ALGOL 60, text of the MC-compiler for the EL-X8*. 1973.
- 49 H. Kok. *Connected orderable spaces*. 1974.
- 50 A. van Wijngaarden, B.J. Mailloux, J.E.L. Peck, C.H.A. Koster, M. Sintzoff, C.H. Lindsey, L.G.L.T. Meertens, R.G. Fisker (eds.). *Revised report on the algorithmic language ALGOL 68*. 1976.
- 51 A. Hordijk. *Dynamic programming and Markov potential theory*. 1974.
- 52 P.C. Baayen (ed.). *Topological structures*. 1974.
- 53 M.J. Faber. *Metrizability in generalized ordered spaces*. 1974.
- 54 H.A. Lauwerier. *Asymptotic analysis, part 1*. 1974.
- 55 M. Hall, Jr., J.H. van Lint (eds.). *Combinatorics, part 1: theory of designs, finite geometry and coding theory*. 1974.
- 56 M. Hall, Jr., J.H. van Lint (eds.). *Combinatorics, part 2: graph theory, foundations, partitions and combinatorial geometry*. 1974.
- 57 M. Hall, Jr., J.H. van Lint (eds.). *Combinatorics, part 3: combinatorial group theory*. 1974.
- 58 W. Albers. *Asymptotic expansions and the deficiency concept in statistics*. 1975.
- 59 J.L. Mijnheer. *Sample path properties of stable processes*. 1975.
- 60 F. Göbel. *Queueing models involving buffers*. 1975.
- 63 J.W. de Bakker (ed.). *Foundations of computer science*. 1975.
- 64 W.J. de Schipper. *Symmetric closed categories*. 1975.
- 65 J. de Vries. *Topological transformation groups, I: a categorical approach*. 1975.
- 66 H.G.J. Pijls. *Logically convex algebras in spectral theory and eigenfunction expansions*. 1976.
- 68 P.P.N. de Groen. *Singularly perturbed differential operators of second order*. 1976.
- 69 J.K. Lenstra. *Sequencing by enumerative methods*. 1977.
- 70 W.P. de Roeper, Jr. *Recursive program schemes: semantics and proof theory*. 1976.
- 71 J.A.E.E. van Nunen. *Contracting Markov decision processes*. 1976.
- 72 J.K.M. Jansen. *Simple periodic and non-periodic Lamé functions and their applications in the theory of conical waveguides*. 1977.
- 73 D.M.R. Leivant. *Absoluteness of intuitionistic logic*. 1979.
- 74 H.J.J. te Riele. *A theoretical and computational study of generalized aliquot sequences*. 1976.
- 75 A.E. Brouwer. *Treelike spaces and related connected topological spaces*. 1977.
- 76 M. Rem. *Associations and the closure statements*. 1976.
- 77 W.C.M. Kallenberg. *Asymptotic optimality of likelihood ratio tests in exponential families*. 1978.
- 78 E. de Jonge, A.C.M. van Rooij. *Introduction to Riesz spaces*. 1977.
- 79 M.C.A. van Zijl. *Empirical distributions and rank statistics*. 1977.
- 80 P.W. Hemker. *A numerical study of stiff two-point boundary problems*. 1977.
- 81 K.R. Apt, J.W. de Bakker (eds.). *Foundations of computer science II, part 1*. 1976.
- 82 K.R. Apt, J.W. de Bakker (eds.). *Foundations of computer science II, part 2*. 1976.
- 83 L.S. van Benthem Jutting. *Checking Landau's "Grundlagen" in the AUTOMATH system*. 1979.
- 84 H.L.L. Busard. *The translation of the elements of Euclid from the Arabic into Latin by Hermann of Carinthia (?), books vii-xii*. 1977.
- 85 J. van Mill. *Supercompactness and Wallmann spaces*. 1977.
- 86 S.G. van der Meulen, M. Veldhorst. *Torric I, a programming system for operations on vectors and matrices over arbitrary fields and of variable size*. 1978.
- 88 A. Schrijver. *Matroids and linking systems*. 1977.
- 89 J.W. de Roeper. *Complex Fourier transformation and analytic functionals with unbounded carriers*. 1978.
- 90 L.P.J. Groenewegen. *Characterization of optimal strategies in dynamic games*. 1981.

- 91 J.M. Geysel. *Transcendence in fields of positive characteristic*. 1979.
- 92 P.J. Weeda. *Finite generalized Markov programming*. 1979.
- 93 H.C. Tijms, J. Wessels (eds.). *Markov decision theory*. 1977.
- 94 A. Bijsma. *Simultaneous approximations in transcendental number theory*. 1978.
- 95 K.M. van Hee. *Bayesian control of Markov chains*. 1978.
- 96 P.M.B. Vitányi. *Lindemayer systems: structure, languages, and growth functions*. 1980.
- 97 A. Federgruen. *Markovian control problems; functional equations and algorithms*. 1984.
- 98 R. Geel. *Singular perturbations of hyperbolic type*. 1978.
- 99 J.K. Lenstra, A.H.G. Rinnooy Kan, P. van Emde Boas (eds.). *Interfaces between computer science and operations research*. 1978.
- 100 P.C. Baayen, D. van Dulst, J. Oosterhoff (eds.). *Proceedings bicentennial congress of the Wiskundig Genootschap, part 1*. 1979.
- 101 P.C. Baayen, D. van Dulst, J. Oosterhoff (eds.). *Proceedings bicentennial congress of the Wiskundig Genootschap, part 2*. 1979.
- 102 D. van Dulst. *Reflexive and superreflexive Banach spaces*. 1978.
- 103 K. van Harn. *Classifying infinitely divisible distributions by functional equations*. 1978.
- 104 J.M. van Wouwe. *GO-spaces and generalizations of metrizability*. 1979.
- 105 R. Helmers. *Edgeworth expansions for linear combinations of order statistics*. 1982.
- 106 A. Schrijver (ed.). *Packing and covering in combinatorics*. 1979.
- 107 C. den Heijer. *The numerical solution of nonlinear operator equations by imbedding methods*. 1979.
- 108 J.W. de Bakker, J. van Leeuwen (eds.). *Foundations of computer science III, part 1*. 1979.
- 109 J.W. de Bakker, J. van Leeuwen (eds.). *Foundations of computer science III, part 2*. 1979.
- 110 J.C. van Vliet. *ALGOL 68 transput, part I: historical review and discussion of the implementation model*. 1979.
- 111 J.C. van Vliet. *ALGOL 68 transput, part II: an implementation model*. 1979.
- 112 H.C.P. Berbee. *Random walks with stationary increments and renewal theory*. 1979.
- 113 T.A.B. Snijders. *Asymptotic optimality theory for testing problems with restricted alternatives*. 1979.
- 114 A.J.E.M. Janssen. *Application of the Wigner distribution to harmonic analysis of generalized stochastic processes*. 1979.
- 115 P.C. Baayen, J. van Mill (eds.). *Topological structures II, part 1*. 1979.
- 116 P.C. Baayen, J. van Mill (eds.). *Topological structures II, part 2*. 1979.
- 117 P.J.M. Kallenberg. *Branching processes with continuous state space*. 1979.
- 118 P. Groeneboom. *Large deviations and asymptotic efficiencies*. 1980.
- 119 F.J. Peters. *Sparse matrices and substructures, with a novel implementation of finite element algorithms*. 1980.
- 120 W.P.M. de Ruyter. *On the asymptotic analysis of large-scale ocean circulation*. 1980.
- 121 W.H. Haemers. *Eigenvalue techniques in design and graph theory*. 1980.
- 122 J.C.P. Bus. *Numerical solution of systems of nonlinear equations*. 1980.
- 123 I. Yuhász. *Cardinal functions in topology - ten years later*. 1980.
- 124 R.D. Gill. *Censoring and stochastic integrals*. 1980.
- 125 R. Eising. *2-D systems, an algebraic approach*. 1980.
- 126 G. van der Hoek. *Reduction methods in nonlinear programming*. 1980.
- 127 J.W. Klop. *Combinatory reduction systems*. 1980.
- 128 A.J.J. Talman. *Variable dimension fixed point algorithms and triangulations*. 1980.
- 129 G. van der Laan. *Simplicial fixed point algorithms*. 1980.
- 130 P.J.W. ten Hagen, T. Hagen, P. Klint, H. Noot, H.J. Sint, A.H. Veen. *ILP: intermediate language for pictures*. 1980.
- 131 R.J.R. Back. *Correctness preserving program refinements: proof theory and applications*. 1980.
- 132 H.M. Mulder. *The interval function of a graph*. 1980.
- 133 C.A.J. Klaassen. *Statistical performance of location estimators*. 1981.
- 134 J.C. van Vliet, H. Wupper (eds.). *Proceedings international conference on ALGOL 68*. 1981.
- 135 J.A.G. Groenendijk, T.M.V. Janssen, M.J.B. Stokhof (eds.). *Formal methods in the study of language, part I*. 1981.
- 136 J.A.G. Groenendijk, T.M.V. Janssen, M.J.B. Stokhof (eds.). *Formal methods in the study of language, part II*. 1981.
- 137 J. Telgen. *Redundancy and linear programs*. 1981.
- 138 H.A. Lauwerier. *Mathematical models of epidemics*. 1981.
- 139 J. van der Wal. *Stochastic dynamic programming, successive approximations and nearly optimal strategies for Markov decision processes and Markov games*. 1981.
- 140 J.H. van Geldrop. *A mathematical theory of pure exchange economies without the no-critical-point hypothesis*. 1981.
- 141 G.E. Welters. *Abel-Jacobi isogenies for certain types of Fano threefolds*. 1981.
- 142 H.R. Bennett, D.J. Lutzer (eds.). *Topology and order structures, part 1*. 1981.
- 143 J.M. Schumacher. *Dynamic feedback in finite- and infinite-dimensional linear systems*. 1981.
- 144 P. Eijgenraam. *The solution of initial value problems using interval arithmetic; formulation and analysis of an algorithm*. 1981.
- 145 A.J. Brentjes. *Multi-dimensional continued fraction algorithms*. 1981.
- 146 C.V.M. van der Mee. *Semigroup and factorization methods in transport theory*. 1981.
- 147 H.H. Tigelaar. *Identification and informative sample size*. 1982.
- 148 L.C.M. Kallenberg. *Linear programming and finite Markovian control problems*. 1983.
- 149 C.B. Huijsmans, M.A. Kaashoek, W.A.J. Luxemburg, W.K. Vietsch (eds.). *From A to Z, proceedings of a symposium in honour of A.C. Zaanen*. 1982.
- 150 M. Veldhorst. *An analysis of sparse matrix storage schemes*. 1982.
- 151 R.J.M.M. Does. *Higher order asymptotics for simple linear rank statistics*. 1982.
- 152 G.F. van der Hoeven. *Projections of lawless sequences*. 1982.
- 153 J.P.C. Blanc. *Application of the theory of boundary value problems in the analysis of a queueing model with paired services*. 1982.
- 154 H.W. Lenstra, Jr., R. Tijdeman (eds.). *Computational methods in number theory, part I*. 1982.
- 155 H.W. Lenstra, Jr., R. Tijdeman (eds.). *Computational methods in number theory, part II*. 1982.
- 156 P.M.G. Apers. *Query processing and data allocation in distributed database systems*. 1983.
- 157 H.A.W.M. Kneppers. *The covariant classification of two-dimensional smooth commutative formal groups over an algebraically closed field of positive characteristic*. 1983.
- 158 J.W. de Bakker, J. van Leeuwen (eds.). *Foundations of computer science IV, distributed systems, part 1*. 1983.
- 159 J.W. de Bakker, J. van Leeuwen (eds.). *Foundations of computer science IV, distributed systems, part 2*. 1983.
- 160 A. Rezus. *Abstract AUTOMATH*. 1983.
- 161 G.F. Helminck. *Eisenstein series on the metaplectic group, an algebraic approach*. 1983.
- 162 J.J. Dik. *Tests for preference*. 1983.
- 163 H. Schippers. *Multiple grid methods for equations of the second kind with applications in fluid mechanics*. 1983.
- 164 F.A. van der Duyn Schouten. *Markov decision processes with continuous time parameter*. 1983.
- 165 P.C.T. van der Hoeven. *On point processes*. 1983.
- 166 H.B.M. Jonkers. *Abstraction, specification and implementation techniques, with an application to garbage collection*. 1983.
- 167 W.H.M. Zijm. *Nonnegative matrices in dynamic programming*. 1983.
- 168 J.H. Evertse. *Upper bounds for the numbers of solutions of diophantine equations*. 1983.
- 169 H.R. Bennett, D.J. Lutzer (eds.). *Topology and order structures, part 2*. 1983.