



The application of a resource logic to the non-temporal analysis of processes acting on resources

Jonathan Hayman
HP Laboratories Bristol
HPL-2003-194
October 2nd, 2003*

E-mail: jmh00@doc.ic.ac.uk

business process
analysis, bunched
implications,
concurrency,
demos

Many systems exist that can be modelled as processes interacting through shared resources -people, factories and computer systems to name but a few. Building upon a logic capable of reasoning about *static* resource models, the Logic of Bunched Implications, we tentatively define the semantics of a logic capable of reasoning about *dynamic* resource models. The logic shall not consider the influence of timing on process interaction. We give a brief overview of other process logics- Hennessy-Milner logic, Computation Tree Logic and modal- μ - which indicate how we may automate the system. We illustrate how our dynamic resource logic may be applied to the analysis of Petri nets. In order to conduct an analysis of processes, we must have a formal representation for them: we choose Demos, and present its salient characteristics. We give a formal definition of its non-time semantics, and prove that this is, in some sense, correct. Following consideration of how we apply our process logic to Demos models, we propose a method for reducing the state-space of models generated by considering a (partially) synchronous execution. We show that such a system at least correctly detects deadlock.

Acknowledgements

Firstly, I would like to thank the Enterprise Solutions Department, and HP Labs as a whole, for providing such a pleasant working environment.

Thanks also to David Pym for a great deal of useful discussion on the logical side of this work; his contribution to this cannot be over-estimated.

Last, but by no means least, I would like to thank Chris Tofts for proposing this project, guiding me into the subject area, providing a wealth of knowledge and, generally, making the time I spent producing this so enjoyable.

Contents

1	Logical Framework	11
1.1	The (static) Logic of Bunched Implications	11
1.2	Transition systems	13
1.2.1	Labelled transition systems	14
1.3	Modal and Temporal Logics	15
1.3.1	Hennesy-Milner Logic	15
1.3.2	General Hennesy-Milner Logic	16
1.3.3	CTL	17
1.3.4	CTL ⁻	18
1.3.5	Modal- μ	19
1.4	Modal BI	22
1.4.1	Additive modalities	23
1.4.2	Multiplicative modalities	23
1.4.3	Observable vs. silent actions	23
1.5	Petri nets in PBI	25
2	Demos	29
2.1	Introduction	29
2.2	Transition rules	31
2.3	Transition system σ Demos	32
2.3.1	Resource actions	32
2.3.2	Hold	33
2.3.3	Initialisation	33
2.3.4	Sequential composition	34
2.3.5	Parallel composition	34
2.3.6	Process synchronisation	35
2.3.7	Non-blocking requests	36
2.3.8	Try, while and req	37
2.3.9	Repeat n -times	37
2.4	Example derivations	37
3	Correctness of transition system	49
3.1	Introduction	49
3.2	Simulation	50
3.3	Soundness of transitions	54
3.4	Weak bisimulation	57
3.5	Kernel of π Demos states	59
3.6	Correspondence between Π_k and σ Demos states	62

3.7	Correctness of translations	64
3.7.1	Correctness of g	64
3.7.2	Correctness of \bar{f}	65
4	Demos and PBI	66
4.1	PBI and σ Demos	66
4.1.1	Intuitive account	68
4.2	Deadlock in PBI	69
4.3	Examples	70
4.3.1	Sub-deadlock	72
5	Synchronous parallel rule	74
5.1	Motivation	74
5.2	Soundness: try and resource splitting	75
5.3	Soundness: Synchrony	78
5.4	Deadlock freedom equivalence in asynchronous and synchronous systems	82
5.4.1	Soundness of deadlock detection	82
5.4.2	Completeness of deadlock detection	84
6	Conclusions	87
6.1	Further work	87

List of Figures

1	Intelligence agency network and process design	9
2	Intelligence agency process execution	10
1.1	Approximants of $X = \sigma \vee [-]X$ to form $\mu X.\sigma \vee [-]X$	20
1.2	Approximants of $X = \sigma \wedge \langle - \rangle X$ to form $\nu X.\sigma \wedge \langle - \rangle X$	21
1.3	Differing semantics of least and greatest fixed points for an example state space	21
1.4	1.4(a) goes to 1.4(b). 1.4(c): no action — blocked by post-condition. 1.4(d): no action — not all pre-conditions met. 1.4(e) goes to 1.4(f)	27
2.1	BNF syntax of restricted Demos considered	30
2.2	Demos program for the intelligence agency example	30
2.3	Jobber source code	39
2.4	Semaphore source code	45
2.5	Deadlock source code	47
3.1	The correspondence between states of the transition system and the of operational semantics.	62
3.2	Bijection between kernel and transition system	64
5.1	Reachable state space for kitchen example derived by asynchronous transition system	75
5.2	Derivation made with the unsound parallelisation rule	77
5.3	Path derived, deterministically, using synchronous parallel rule	80
5.4	State space for kitchen example derived by the synchronous transition system	81

List of Tables

1.1	Kripke semantics of (static) BI	12
1.2	Semantics of CTL formulæ	18
1.3	Kripke semantics of PBI (excluding modalities)	24
1.4	Kripke semantics of additive modalities of PBI	24
1.5	Kripke semantics of multiplicative modalities of PBI	25
1.6	Kripke semantics of observable modalities of PBI	26
4.1	Demos semantics of PBI	67

Introduction

When we design complex systems, we should ask questions of them before they are implemented — will it run fast enough; would it be more efficient for a human to do a particular sub-task; *etc.* To answer such questions will require automation.

In the general case, then, how are we to automate such analyses of systems? Perhaps the most obvious manner (at least to the computer scientist with a disposition towards the practical) is to simulate the design, probably observing several ‘runs’ of the simulation, thereby repeating an experiment. Though, in principle, this approach generalises to all types of system by Turing completeness (and normally requires little background knowledge other than an understanding of the system and the ability to program), in some ways it is grossly unsatisfactory. What if, for example, we are simulating a network and no run of the simulation has a very large file being transmitted (or any other *rare-event*)? Clearly, our results may be overly optimistic, which could lead, for example, to us breaching the terms of a service-level agreement. Conversely, what if our simulation has a few too many rare-events (large files)? This time, our results may be overly pessimistic, resulting in the over-specification of the network at financial cost to us. One answer to this is to run the simulation for longer, so as to increase the probability that rare-events are encountered. Of course, though, given a finite amount of time, our results will always be approximate.

An alternative approach is to construct an analytical model of the system to be analysed. Returning to the network example, a model would probably result in a set of equations to solve in order to obtain some measurements of the system’s performance (*metrics*). The advantage of this is that one analysis gives exact results, taking rare-events into fair consideration, and, very often, can be obtained much faster than by simulation. The flip-side is that, generally, we are only able to do this for particular classes of system. For example, the simplest analyses of queuing networks (*e.g.*, computer networks that transmit packets) assume a (negative) exponential probability distribution of service time and arrival rate; that is,

$$P(\text{Service time} \leq t) \propto e^{-\mu t}.$$

Fortuitously, many systems approximate very closely the assumptions made to build such models. Proving that the assumptions behind such models are met (for example, that the time between packet arrivals follows a negative exponential distribution), though, is somewhat more complicated and subjective;

simulation has the same problem, manifested in the problem of how to choose the distributions or sample data that we use to parameterise the simulation.

We must now move on to decide upon the type of system that we are to model. Naturally, we cannot hope to provide a mathematical model for *all* types of system, able to answer *all* types of question! We shall consider a notion again familiar to computer scientists, that of *processes*, and give the definition found in [OED73]:

Process . . . A continuous and regular action or succession of actions, taking place or carried on in a definite manner; a continuous (natural or artificial) operation or series of operations. . .

Central to our analysis of processes, of course, shall be the environment in which they operate. In particular, a process may be acting in an environment with other processes; we shall say that the processes are acting in *parallel* or *concurrently*. It has long been known — before, even, the first parallel computer was built — that it is difficult to reason about parallel processes, as noted in the very first item to appear in the *Communications of the ACM* [Per58]. Our work to provide some degree of automation should, therefore, be of some value.

In order to further restrict the the scope of our work, we shall analyse the behaviour of processes when they interact through the acquisition and release of resources. To borrow Robin Milner’s famous jobshop example [Mil89], consider two workers, called *jobbers*, a hammer, a mallet and a stream of jobs to be done. Considering the workers to be processes and the hammer and mallet to be resources¹, we might want to ask all sorts of questions. For example, would we see a gain in productivity if we gave each jobber their own hammer and mallet?

A different class of question relates to classical concurrency: do the jobbers deadlock (not be able to take any action without having terminated)? Let us suppose that the jobbers are simple folk; if they work in isolation, they will work happily until the end of their day and will not deadlock. They are also tidy folk, leaving the resources as they found them at the start of the day (so they return the tools to their correct place). If the jobbers work at different times of the day, we may conclude that they will not deadlock. Should, though, our analysis of the activity of the jobbers take into account the times that they actually work? C.A.R. Hoare [Hoa85] argues not: our analysis would not take into account the outcome of a jobber arriving at work late and then working late (a not inconceivable scenario!). Thus, perhaps, we should not argue that the system never deadlocks.

To give a more concrete example and further justification, consider the communication network of Figure 1(a). Nodes represent intelligence agencies in the UK and the US, and arcs represent dedicated, secure communication channels (perhaps based upon quantum encryption) that can only be ‘owned’ by one node at a time; in a sense, they are *resources*. The hub in the centre of the diagram is capable of serving many lines at once and represents no limitation on the system. It is important to abstract away the irrelevant when modelling, so we

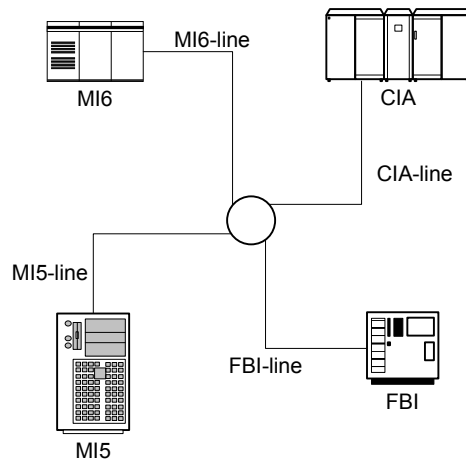
¹We could also consider the hammer and mallet to be (very simple) processes, which alludes to the consideration of resources as processes under certain circumstances

shall ignore its behaviour. Suppose that MI6 has a process that starts every day at 12:43GMT which sends some data from MI5 to the FBI. To do this, it claims ownership of the MI6 line, followed by the MI5 line, then the FBI line. Once a process (in this case, an intelligence agency) starts to get a line, no other process can start to get the same line. Between each acquisition, a 2-minute pause occurs to allow some handshaking protocol to occur. Similarly, the CIA has a process that starts every day at 12:40GMT which sends information from the FBI to MI5: in doing this, it claims the CIA line, the FBI line and then the MI5 line. These processes are represented in Figure 1(b). What normally happens is that the CIA process has control of the MI5 line before MI6 tries to acquire it. Consequently, the MI6 process waits until the MI5 line is released thereby avoiding deadlock. This execution is illustrated in 2(a).

Now, suppose that the CIA process takes an extra two minutes to gain control over the FBI line — possibly because of interference during the handshaking period. At 12:46, the CIA process will try to gain control over the MI5 line but will fail — it is now owned by MI6 — and so will wait for it to become available. At 12:47, the MI6 process will try to acquire the FBI line but will also fail — it is owned by the CIA — and so will wait. Both the CIA and the MI6 processes, then, will wait indefinitely: they are in deadlock. This execution is shown in 2(b).

Unfortunately, if we had taken time into account when deciding upon deadlock freedom, from our model we would never have known that deadlock could occur, and would fail to see the inherent flaw in the system. This could be dangerous for MI6 (and, similarly, the CIA) — MI5, which has responsibility for counterespionage, might notice that its secure link with the outside world was never available and investigate why this was so.

We have considered a *discrete-event* system — we do not consider actions on resources to happen over a period of time. Further, our consideration shall be of *discrete-resource* systems. So, rather than having 2.4 litres of oil that we can partition to allocate to processes as we wish, we shall consider n units of oil where $n \in \mathbb{N}$ (n is a non-negative integer). However, our theoretical approach could well be extended to a notion of continuous resource. Thus the system developed herein shall operate on *discrete-event*, *discrete-resource* processes. Not coincidentally, we shall consider a simulation language, Demos, that has both these properties.




(a) Communication network


MI6	CIA
Get MI6 line	Get CIA line
Get MI5 line	Get FBI line
Get FBI line	Get MI5 line
(transfer data)	(transfer data)
Release all	Release all

(b) Communication processes

Figure 1: Intelligence agency network and process design



(a) Normal execution



(b) Deadlocking execution

Figure 2: Intelligence agency process execution

Chapter 1

Logical Framework

1.1 The (static) Logic of Bunched Implications

The Logic of Bunched Implications [OP99], **BI**, allows us to reason about resources. Given a particular collection of resources (a *model*), we ask whether propositional formulæ hold.

Technically, **BI** is a logic based around *bunches*, tree-like structures of propositions (and, in the predicate case, constants). It seeks to combine the additive semantics of intuitionistic logic (giving standard meaning to the operators $\wedge, \vee, \rightarrow$) with the multiplicative semantics of Girard’s Linear Logic [Gir87] (giving the operators $*, \multimap$). The multiplicative connectives, essentially, allow us to make the kind of assertions that we would like to make about systems with resources. When we write formulæ in **BI**, the multiplicatives shall bind more strongly than the additives; conjunction shall bind more strongly than disjunction, which in turn binds more strongly than implication, and negation binds strongest of all. All binary operators are taken to be right-associative (this is only of importance to the implications — the other binary operators are commutative and associative). For example,

$$\phi * \psi \wedge \psi \multimap \gamma \equiv (\phi * \psi) \wedge (\psi \multimap \gamma)$$

Logics are defined in two parts, syntax and semantics. The syntax of the logic is just a way to write down sentences that relate to the semantics. Attached to the syntax may be a deduction system, taking sentences written according in the language defined by the syntax to derive other syntactically correct sentences. The semantics of the logic defines what the formulæ written in the syntax actually *mean*. It may well be the case (as it is with **BI**) that we can define several semantics for the logic whilst preserving the ‘correctness’ of the deduction system. We choose to consider one of the simplest (and the one that can have the simplest notion of resource): the Kripke semantics.

The Kripke semantics of **BI** allows us to define the semantics through a mathematical structure called a *monoid*,

$$\mathcal{M} = \langle M, e, \circ, \sqsubseteq \rangle,$$

where M represents a domain for our monoid, e represents the identity element under composition, \circ is commutative and represents the composition of elements of M , and \sqsubseteq is a pre-order on M . The semantics of **BI** based upon this monoid is given in Table 1.1.

Let $m \in M$, p be an atomic proposition, φ, ψ range over propositions and $\|p\| \subseteq M$ be the set of states in which the proposition p holds. Then the Kripke semantics is given by:

$m \models p$	iff	$m \in \ p\ $
$m \models \varphi \wedge \psi$	iff	$m \models \varphi$ and $m \models \psi$
$m \models \varphi \vee \psi$	iff	$m \models \varphi$ or $m \models \psi$
$m \models \varphi \rightarrow \psi$	iff	for all $n \sqsubseteq m$, $n \models \varphi$ implies $n \models \psi$
$m \models \varphi * \psi$	iff	for some $n, n' \in M$ $m \sqsubseteq n \circ n'$ and $n \models \varphi$ and $n' \models \psi$
$m \models \varphi \multimap \psi$	iff	for all $n \in M$ $n \models \varphi$ implies $m \circ n \models \psi$

Table 1.1: Kripke semantics of (static) **BI**

There are two types of operator in the logic: additive and multiplicative. The additive operators, \wedge and \rightarrow , are just as in intuitionistic propositional logic: \wedge is used to indicate that more than one formula holds in the state, and \rightarrow is implication.

More interesting are the multiplicative operators. In $\varphi * \psi$, $*$ is used to indicate that we can partition (using \circ) our current state; one part satisfies φ and the other ψ . $\varphi \multimap \psi$ indicates that if we add all that we need to satisfy φ to our current state, ψ will hold.

It may be helpful to consider a monoid based upon a notion of *resource*. If we let R denote a multiset of discrete¹ resources, $\mathcal{P}(R)$ denote the set of all sub-multisets of R and multiset union be \uplus , then we can define the monoid

$$\mathcal{M} = \langle \mathcal{P}(R), \uplus, \uplus, = \rangle.$$

The consequent intuitive meaning of $m \models \varphi * \psi$ is that we can split the resource multiset m in two, one part making φ true and the other making ψ true. $m \models \varphi \multimap \psi$ intuitively means that given the resources to make φ true, if we add them to what we have in m , ψ will be true.

Example (Discrete resources) For each resource $r \in R$, define $\|r\| = \{r\}$; so, for every resource r , r holds iff r is the only available resource. Also define a proposition $move_{bolt}$, holding with $\{wrench, mallet\}$; it signifies that we can move a bolt iff we have a hammer and a mallet available.

¹As we have said, not, for example, n metres of string that can be split into m and m' metres of string where $m + m' = n$ for arbitrary $m, m' \in \mathbb{R}$.

The Kripke discrete resource semantics gives:

$$\begin{array}{ll}
hammer & \models hammer \\
hammer, mallet & \not\models hammer \\
hammer, mallet & \models hammer * mallet \\
hammer, mallet & \models hammer * \top \\
hammer, mallet & \models hammer * \top \wedge hammer * \top \\
hammer, mallet & \not\models hammer * hammer * \top \\
hammer, mallet & \not\models (hammer * \top) * (hammer * \top) \\
hammer, mallet & \models wrench \multimap move_{bolt} * \top \\
hammer, mallet & \models (wrench \multimap move_{bolt}) * \top \\
hammer, mallet & \models move_{bolt} \multimap mallet * mallet * \top \\
hammer, mallet & \not\models mallet \multimap mallet * mallet * \top
\end{array}$$

The monograph [Pym02] provides a substantially more detailed, technical and comprehensive account of **BI**. Though it is possible to encode action (Petri nets are discussed in [POY02, Pym02]) into **BI** by attaching certain semantics, the resource interpretation above does not have a notion of how resource evolves with the action of processes; we are only able to make *static* judgements. The necessary tools to describe action, transition systems and modal logics, are introduced below.

Notice that we have defined neither negation nor falsity in this system. We could have included falsity, \perp , but this makes the associated natural deduction system incomplete. With \perp , we could define intuitionistic negation: $\neg\varphi = \varphi \rightarrow \perp$. We obtain a classical logic (thence negation) if \sqsubseteq is equality.

1.2 Transition systems

Since the dissemination of Gordon Plotkin's notes on Structural Operational Semantics, [Plo81], as detailed in [Plo03], it has become commonplace to describe the activity of programming languages as a one-step relation between states defined *structurally*. That is, we give an inductive definition of the effect of each command on the state of the system. This is normally presented in the proof-tree natural deduction style:

$$(\text{RULE-NAME}) : \frac{\text{PREMISES}}{\text{CONCLUSIONS}} (\text{Condition})$$

This asserts that if we can derive by the set of rules all the premises then the conclusions will hold, providing the side-condition is met. Thus, as we have defined the transition relation structurally, we can prove properties of the relation by structural induction.

Consider, as an example, an elementary imperative programming language. The one-step transition relation shall be written \Longrightarrow . A transition system shall represent state as a function from variables to values, with $s[x \mapsto c]$ representing the *system* state s but with x having the value c . If we let $\mathcal{A}[[e]]s$ represent

the result of evaluating the arithmetic expression e in state s , we could define assignment as follows:

$$\text{(ASS)} : \frac{}{\langle x := e, s \rangle \Longrightarrow s[x \mapsto c]} (\mathcal{A} \llbracket e \rrbracket s = c)$$

This expresses that the program $x := e$ executing in an arbitrary state s terminates after one step leaving the same state other than x is now $c = \mathcal{A} \llbracket e \rrbracket s$, the result of evaluating e in s .

Together with a series of other rules for the various programming language constructs, for example the sequential composition of commands,

$$\text{(COMP}^1\text{)} : \frac{\langle s, P_1 \rangle \Longrightarrow \langle s', P_1' \rangle}{\langle s, P_1; P_2 \rangle \Longrightarrow \langle s, P_1'; P_2 \rangle}$$

this would define the transition relation \Longrightarrow .

1.2.1 Labelled transition systems

It is sometimes useful to label the the transitions between states. As noted in [Pit02], when we define the semantics of languages we usually have *reduction* transitions (which describe real action within the system) and *simplification* actions (which describe, for example, modifications to syntax). Though it may be considered desirable to define the transition relation solely in terms of reduction transitions, using a number of auxilliary relations to describe simplification, labelling allows us to make explicit the rôle of each transition. Furthermore, labelling shall allow finer queries to be made in the modal logic (as shall be indicated later) though non-labelled systems are commonplace *e.g.*, [MP83].

As we have already stated, when we come to describe Demos, we shall be considering the interaction between processes working on resources. In particular, we shall consider the *available* resources within the system; that is, those that have not been acquired by any process. This shall be described by a multiset; we let $\mathbf{R}, \mathbf{R}', \dots$ range over such multisets of resource, and P, P', Q, \dots range over the code remaining to be executed. We shall call $\mathbf{R} \triangleleft \llbracket P \rrbracket$ a *state* of the transition system. Transitions are of the form $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\alpha} \mathbf{R}' \triangleleft \llbracket P' \rrbracket$. We read this as:

Initially, all the resources in \mathbf{R} are available and we have an active process P . After an action α , the resources in \mathbf{R}' are available and the active process is P' .

We distinguish two classes of transition: σ -transitions and τ -transitions. σ -transitions indicate action, and τ -transitions indicate silent or invisible actions, related to the simplification relations described above, that perform syntactic operations to expand loops and skip by commands that are not considered.

An alternative would have been to define an equivalence relation, based upon syntax, to indicate that we consider pairs of states to be the same. For example, given that our consideration is of action rather than time, we would consider $\mathbf{R} \triangleleft \llbracket \text{hold}(t); \text{putR}(r, 1) \rrbracket$ and $\mathbf{R} \triangleleft \llbracket \text{putR}(r, 1) \rrbracket$ to be equivalent. Such an equivalence relation, though, would decrease the tractability of search in a language such

as Prolog: the arbitrary introduction of hold commands, do 0 loops *etc.* (the symmetry of the equivalence relation being the key problem) would be allowed, so we choose to define the τ relation, which is not symmetric.

An excellent account of labelled transition systems is given in [Pac89], along with a method for analyzing assertions made of programs.

1.3 Modal and Temporal Logics

The Free Online Dictionary of Computing [FOL] defines a modal logic to be: *An extension of propositional calculus [propositional logic] with operators that express various “modes” of truth.* So, rather than just being able to express “ φ is true”, we might be able to express that “ φ is possibly true,” “ φ is necessarily true,” “ φ is true at some point in the future,” “ φ is true in the next state,” and so on.

Formally, adopting the Kripke semantics of modal logics, we define a set of states (or worlds), W , and a relation R between these states. This defines a *Kripke frame* (essentially, just a set of states with transitions between them). We then define an evaluation function for propositional atoms in each state, $h(s, p) : \{\top, \perp\}$ (an alternative notation to this is $s \in \llbracket p \rrbracket \iff h(s, p)$). A modal logic allows us to write formulæ over the transitions between worlds in the frame. So, for example, we can write “given that I am in world s , in every reachable world, (something) holds”.

Regarding worlds as program states, the relation between worlds as time (or program execution), we obtain an important type of modal logic: temporal logic. We see that temporal logics allow us to reason about the activity of processes. In their simplest form, they allow us to make assertions about the state of a process after one execution step; this is easily generalised to multiple-steps, as shall be outlined below.

1.3.1 Hennessy-Milner Logic

Here we shall give an account of Hennessy-Milner logic, which was introduced in [HM80, HM85]. Our notation shall follow that presented later in [Mil89]; in §1.3.2 we shall generalise the following definition to match the system of [Mil89].

Firstly, in order to create a well-founded inductive definition of modal formulæ, we shall need to define when an atomic proposition is satisfied. The definition is straightforward: a state s satisfies an atomic proposition p , written $s \models p$, iff $s \in \llbracket p \rrbracket$, where $\llbracket p \rrbracket$ is the set of states in which the proposition holds. Similarly, the definitions of conjunction and disjunction follow the additive cases of Table 1.1:

$$s \models p \wedge q \text{ iff } s \models p \text{ and } s \models q$$

and

$$s \models p \vee q \text{ iff } s \models p \text{ or } s \models q$$



Suppose that a process is in a state s . If, by some action in a set A , a transition from s can be derived to a state in which a (possibly temporal) formula φ

is true, we write $s \models \langle A \rangle \varphi$ (where $\langle A \rangle$ is pronounced “diamond- A ”). If all transitions derivable from s by actions in A lead to states in which φ holds, we write $s \models [A] \varphi$ (pronounced “box- A ”). In a more formal notation:

Definition: (Simple modalities)

$$\begin{aligned} s \models \langle A \rangle \varphi & \text{ iff } \exists \alpha, s'. \quad \alpha \in A \text{ and } s \xrightarrow{\alpha} s' \text{ and } s' \models \varphi \\ s \models [A] \varphi & \text{ iff } \forall \alpha, s'. \quad \alpha \in A \text{ and } s \xrightarrow{\alpha} s' \text{ implies } s' \models \varphi \end{aligned}$$

Pictorially, these express the following:



Assuming the existence of negation within our logic, it is not hard to see that the following equivalence exists between the modalities:

$$\langle A \rangle \varphi \iff \neg [A] \neg \varphi$$

We may choose, then, only to define one of the above modalities within the basic logic, and define the other to be a ‘derived’ modality.

In the sequel it shall be convenient to use an abbreviated syntax to define the action set. Assuming that the set of all action labels is \mathcal{A} , we shall write:

$$\begin{array}{ll} [a, a', \dots] & \text{for } [\{a, a', \dots\}]; \\ [-A] & \text{for } [\mathcal{A} \setminus A]; \\ [-a, a', \dots] & \text{for } [\mathcal{A} \setminus \{a, a', \dots\}]; \\ [] & \text{for } [\emptyset].^2 \end{array}$$

In particular, we shall make use of $[-]$, which abbreviates $[A]$. The obvious symmetric abbreviations shall be made for the diamond modality as well.

1.3.2 General Hennessy-Milner Logic

Though the inductive definition above allows formulæ to have nested modalities, so far we can only make statements about processes after a finite number of steps. Furthermore, to assert that a program terminates before, say, 1000 steps would require a very long sentence.

²There is no fixed convention for this abbreviation. Its use shall be avoided wherever possible.

With this in mind, suppose that we generalise conjunction (or its dual, disjunction) to require validity of all conjuncts indexed by a set I :

$$s \models \bigwedge_{i \in I} \varphi_i \quad \text{iff} \quad \forall i \in I. \ s \models \varphi_i$$

Given that we can express, for arbitrary n , that a program in state s terminates (or at least stops running) after exactly n transitions from some set A as

$$s \models \langle A \rangle^n [A] \perp,$$

we may express the above property (termination before 1000 steps) succinctly as

$$s \models \bigvee_{0 \leq i \leq 1000} \langle A \rangle^i [A] \perp.$$

Notice that we have abbreviated $i \in \{0, \dots, 1000\}$ by $0 \leq i \leq 1000$.

If we allow infinitary conjunction ($i \in \mathbb{N} \cup \{\omega\}$, where ω is the first ordinal, the least number greater than every number in \mathbb{N} , otherwise written ‘ \mathbb{N} ’ itself; we shall interchangeably write \mathbb{N}_ω for $\mathbb{N} \cup \{\omega\}$), we may express deadlock freedom of processes by the simple formula:

$$\mathfrak{D}_f^{hml} \stackrel{\text{def}}{\iff} \neg \bigvee_{i \in \mathbb{N} \cup \{\omega\}} \langle A \rangle^i [A] \perp$$

which is logically equivalent to

$$\bigwedge_{i \in \mathbb{N} \cup \{\omega\}} [A]^i \langle A \rangle \top$$

The reader may gain some intuition as to the meaning of this sentence by ‘unfolding’ the conjunction. It expresses that the process is not stuck in this state, that every state reachable in one transition is not stuck (and so we can perform another transition), that every state reachable in two transitions is not stuck (and so we can perform another transition), . . .

1.3.3 CTL

The Computation Tree Logic of [CES86] allows some temporal properties to be expressed more simply than the Hennessy-Milner logic described above does, as well as being more amenable to model checking. The definition of CTL in [CES86] requires that the transition relation is total, *i.e.*, every state has a next state. We regard this as unreasonable for our purposes, so we shall *not* require this.

Letting AP be the set of all atomic propositions, the following two rules inductively define the class of CTL formulæ:

1. Every proposition $p \in AP$ is a CTL formula.
2. If φ and ψ are CTL formulæ, so are $\neg\varphi$, $\varphi \wedge \psi$, $\langle - \rangle\varphi$, $[-]\varphi$, $\forall[\varphi \mathbf{U} \psi]$ and $\exists[\varphi \mathbf{U} \psi]$.

The semantics attached to the formulæ composed of first-order propositional logic operators is conventional, as for the box and diamond modalities. $\forall[\varphi\mathbf{U}\psi]$ intuitively means: *whatever path is chosen from the given state, φ will hold until ψ holds (and ψ actually does hold somewhere)*. Similarly, $\exists[\varphi\mathbf{U}\psi]$ intuitively means: *a path can be chosen where φ holds until ψ holds (and ψ actually does hold somewhere)*.

A *path* is an possibly infinite sequence of states, $[s_0, s_1, \dots]$ where $\forall i \geq 0 [s_i \xrightarrow{\sigma_i} s_{i+1}]$ and $\xrightarrow{\sigma}$ is the transition relation of the labelled transition system. If the path is finite, there must be no derivable transition from the last state.³ Using this definition, Table 1.3.3 defines the semantics of CTL.

[CES86] gives a relatively straightforward algorithm for model-checking CTL formulæ where the state space of the given process is finite.

$s_0 \models p$	iff	$s_0 \in \ p\ $
$s_0 \models \neg\varphi$	iff	$s_0 \not\models \varphi$
$s_0 \models \varphi \wedge \psi$	iff	$s_0 \models \varphi$ and $s_0 \models \psi$
$s_0 \models \varphi \vee \psi$	iff	$s_0 \models \varphi$ or $s_0 \models \psi$
$s_0 \models \langle - \rangle \varphi$	iff	for some state s' and action $\alpha \in \mathcal{A}$, $s_0 \xrightarrow{\alpha} s'$ and $s' \models \varphi$.
$s_0 \models [-]\varphi$	iff	for all states s' and actions α , $s_0 \xrightarrow{\alpha} s'$ implies $s' \models \varphi$
$s_0 \models \forall[\varphi\mathbf{U}\psi]$	iff	for all paths $[s_0, s_1, \dots]$, $\exists i \geq 0 [s_i \models \psi \wedge \forall j [0 \leq j < i \implies s_j \models \varphi]]$
$s_0 \models \exists[\varphi\mathbf{U}\psi]$	iff	for some path $[s_0, s_1, \dots]$, $\exists i \geq 0 [s_i \models \psi \wedge \forall j [0 \leq j < i \implies s_j \models \varphi]]$

Table 1.2: Semantics of CTL formulæ

Deadlock freedom may be represented in CTL as:

$$\mathfrak{D}_f^{CTL} \stackrel{\text{def}}{\iff} \neg\exists[\top\mathbf{U}\neg\langle - \rangle\perp]$$

1.3.4 CTL⁻

CTL provides a relatively clear language for expressing how propositions hold through the evolution of processes. In many cases, however, it is sufficient to use a subclass of CTL formulæ called CTL⁻.

³As we do not require that the transition relation be total, our definition of path differs from that in [CES86] in that we allow finite sequences.

CTL⁻ does not have the $\exists[-\mathbf{U}-]$ or $\forall[-\mathbf{U}-]$ modalities, but does have $\exists\mathbf{F}$, $\forall\mathbf{F}$, $\exists\mathbf{G}$ and $\forall\mathbf{G}$. These are defined as follows (through the semantics of the CTL modalities):

- $s \models \exists\mathbf{F}\varphi \stackrel{\text{def}}{\iff} s \models \exists[\top\mathbf{U}\varphi]$
This means that there is a path in which φ holds in some future state — φ *potentially* holds.
- $s \models \forall\mathbf{F}\varphi \stackrel{\text{def}}{\iff} s \models \forall[\top\mathbf{U}\varphi]$
This means that along every path from s , there is a state in which φ holds — φ is *inevitable*.
- $s \models \exists\mathbf{G}\varphi \stackrel{\text{def}}{\iff} s \models \neg\forall[\top\mathbf{U}\neg\varphi]$
This means that there is a path from s where φ holds in every state.
- $s \models \forall\mathbf{G}\varphi \stackrel{\text{def}}{\iff} s \models \neg\exists[\top\mathbf{U}\neg\varphi]$
This means that for every path from s , φ holds in every state — φ is *globally* true.

$\forall\mathbf{F}$ is frequently written \mathbf{AF} ; $\exists\mathbf{G}$ is frequently written \mathbf{EG} ; *etc.*

We can represent deadlock freedom in CTL⁻ as:

$$\mathfrak{D}_f^{CTL^-} \stackrel{\text{def}}{\iff} \forall\mathbf{G}\langle-\rangle\top$$

1.3.5 Modal- μ

The modal- μ calculus (a good introduction may be found in [BS01]; we summarize the essentials here) is a logic where formulæ may be defined recursively. Deadlock freedom, for example, may be recursively defined with respect to a set of actions A as:

$$\mathfrak{D}_f^{rec} = {}^4\langle A \rangle\top \wedge [A]\mathfrak{D}_f^{rec}$$

This states that a state is deadlock free if it is possible to go to another state and any state reached is also deadlock free.

As before, we define the set of states in which a proposition φ holds as $\|\varphi\| \in \mathcal{P}(S)$, where S is the set of all states and $\mathcal{P}(S)$ is the powerset of the set of states, $\mathcal{P}(X) = \{Y \mid Y \subseteq X\}$. In the above example, we have used \mathfrak{D}_f^{rec} as a *variable* over logical formulae. If a formula has a free variable Z , we indicate this by writing it as $\rho(Z)$. This may be regarded as a function $\rho(Z) : \mathcal{P}S \rightarrow \mathcal{P}S$. Now, by the Knaster-Tarski theorem, the function has least and greatest *fixed points*; that is, there are least and greatest sets that are solutions to the equation $Z = \rho(Z)$. These shall be denoted $\mu Z.\rho(Z)$ and $\nu Z.\rho(Z)$, respectively. For reference, the semantics of the fixpoint operators μ and ν is given by:

$$\|\mu Z.\rho(Z)\| = \bigcap \{S \mid S \supseteq \|\rho(S)\|\}$$

and

$$\|\nu Z.\rho(Z)\| = \bigcup \{S \mid S \subseteq \|\rho(S)\|\}$$

⁴We shall see later that we should have used \subseteq rather than =

One of the difficulties of writing formulæ in modal- μ is deciding which fixpoint operator to use, and, for the reader, understanding the meaning of this choice. We may gain some intuition if we consider the way in which the fixed points may be constructed. To construct the least fixed point of a formula $X = \sigma \vee [-]X$ for the state-space illustrated through Figure 1.1, we would start with the formula holding nowhere; that is, $\|X\| = \emptyset$. We shall call this the zeroth-approximant, written X^0 (Fig. 1.1(a)). Now, one application of the formula later, we have $X^1 = \|\sigma\| \cup$ the set of states with no transitions from them; this is precisely the set of states where σ holds or every transition from the state leads to a state in \emptyset (Fig. 1.1(b)). This series of approximations continues until a fixed point is reached (Fig. 1.1(d)).

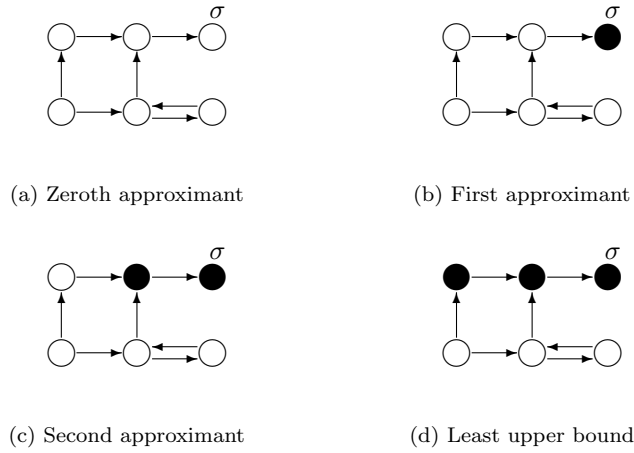


Figure 1.1: Approximants of $X = \sigma \vee [-]X$ to form $\mu X. \sigma \vee [-]X$

Similarly, the greatest fixed point may be found by iterating from the set of all states, taking the intersection of the iterates. An example of this is given in Figure 1.2.

As an alternative to this interpretation, the authors of [BS01] propose the following maxim:

ν means looping; μ means finite looping.

Consider the formula $Z = \sigma \vee \langle - \rangle Z$, meaning that a state either makes σ hold or some such state is reachable. The maxim gives different meaning to the formulæ $\mu Z. \sigma \vee \langle - \rangle Z$ and $\nu Z. \sigma \vee \langle - \rangle Z$; the former means that there is a finite-length path to a state where σ holds, whereas the latter means that there is a possibly infinite length path to a state where σ holds. Notice that this does not imply that σ holds in any particular state $s \in \mathcal{S}$! Figure 1.3 is an illustration of the semantics of this formula. Further difficulty comes from the nesting of fixpoint operators, which can be extremely difficult to interpret. The reader, again, is referred to [BS01] for a proper treatment of the logic. An exercise for the reader is to decide which fixpoint operator should be used in the definition of deadlock freedom⁵.

⁵**Answer:** ν , giving $\mathfrak{D}_f^\mu \stackrel{\text{def}}{=} \nu X. (A) \top \wedge [A]X$.

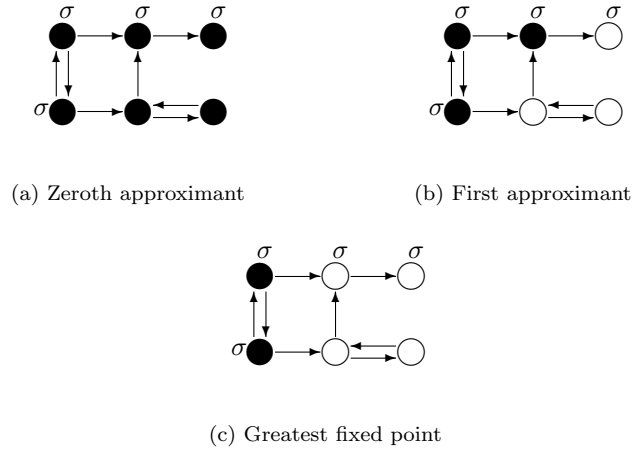



Figure 1.2: Approximants of $X = \sigma \wedge \langle - \rangle X$ to form $\nu X. \sigma \wedge \langle - \rangle X$



Filled circles represent states in the fixed point solution for Z ; hollow circles represent states not in the fixed point solution for Z . σ holds in no state.

Figure 1.3: Differing semantics of least and greatest fixed points for an example state space

It is possible to encode all CTL formulæ quite simply in modal- μ .

- $\exists[\varphi \mathbf{U} \psi] \iff \mu Z. \psi \vee (\varphi \wedge \langle - \rangle Z)$.
- $\forall[\varphi \mathbf{U} \psi] \iff \mu Z. \psi \vee (\varphi \wedge [-] Z \wedge \langle - \rangle Z)$.

The definition of $\forall[-\mathbf{U}-]$ requires that we be able to enter a state in which Z holds because we do not assume that the transition relation is total.

In defining the modalities of CTL^- , we make use of the modal- μ equivalence $\neg\mu Z.(\neg\varphi) \equiv \nu Z.\varphi$.

- $\forall \mathbf{G} \varphi \stackrel{\text{def}}{\iff} \nu Z. \varphi \wedge [-] Z$;
- $\exists \mathbf{G} \varphi \stackrel{\text{def}}{\iff} \nu Z. \varphi \wedge (\langle - \rangle Z \vee [-] \perp)$;
- $\forall \mathbf{F} \varphi \stackrel{\text{def}}{\iff} \mu Z. \varphi \vee ([-] Z \wedge \langle - \rangle Z)$;
- $\exists \mathbf{F} \varphi \stackrel{\text{def}}{\iff} \mu Z. \varphi \vee \langle - \rangle Z$.

CTL may not, however, encode all formulæ written in modal- μ ; we see that modal- μ is more powerful than CTL (and also, therefore, CTL^-).

The Edinburgh Concurrency Workbench [MS] uses this fact — natively, the system checks assertions made in modal- μ , but it is possible to load macros that convert formulæ written in CTL or CTL^- to modal- μ . Users that are not fluent in modal- μ can, therefore, use a simpler language.

1.4 Modal BI

The addition of modalities to **BI** was first suggested in [POY02, Pym02]. Following our pedagogical account of modal calculi, we shall now extend these definitions to define **PBI**⁶. We shall not venture beyond the modalities of HML at this stage, though it would not be hard to extend the logic to allow recursion in the manner of modal- μ , for example. As before, process-resource terms shall be written $\mathbf{R} \triangleleft \llbracket P \rrbracket$, where \mathbf{R} represents the globally available resources and P represents the current state of the process. We define the set of all such states to be $\mathcal{S} = \text{Res} \times \text{Proc}$.

As is the case for **BI**, the semantics of the logic is defined according to a Kripke resource monoid. This time, though, we shall define the monoid over process-resource pairs:

$$\mathcal{M} = \langle \text{Res} \times \text{Proc}, (e_R, e_P), \circ, \sqsubseteq \rangle,$$

where $\text{Res} \times \text{Proc}$ is the set of all process-resource terms \mathcal{S} , \circ is a composition operator, (e_P, e_R) is the identity element under \circ , and \sqsubseteq is a pre-order on \mathcal{S} . As \mathcal{S} is a set of pairs, we shall let $\mathbf{R} \triangleleft \llbracket P \rrbracket$ range over it.

Table 1.3 gives the Kripke resource monoid semantics of the non-modal part of **PBI**. This corresponds to the semantics of **BI**, other than the multiplicative

⁶Process **BI**

units, I_P and I_R . These are introduced to allow expression of zero-resource and null-process respectively. So:

$$\|I_R\| = e_R \times Prog$$

$$\|I_P\| = Res \times e_P$$

We may see the normal multiplicative unit of **BI**, I , as $I_R \wedge I_P$. Thus I is extraneous. In the sequel, we shall take modalities to bind more strongly than anything else; between themselves, they shall be of equal precedence.

1.4.1 Additive modalities

As the static logic, **BI**, has both additive and multiplicative operators, it seems natural to introduce both additive and multiplicative modalities into our process logic.

Presented in Table 1.4, the additive modalities describe process evolution without a change in the resource component. For resource formulæ φ_R , such a definition would allow the arbitrary introduction and removal of additive modalities,

$$(DI_{res}) : \frac{\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi_R \quad \varphi_R : \mathbf{RForm}}{\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle - \rangle \varphi_R}$$

(and similar for ‘box’ and elimination rules), where $\varphi_R : \mathbf{RForm}$ indicates that φ_R is a resource formula.

1.4.2 Multiplicative modalities

We shall define four multiplicative modalities, $\langle - \rangle_{\mathbf{new}}^+$, $[-]_{\mathbf{new}}^+$, $\langle - \rangle_{\mathbf{new}}^-$ and $[-]_{\mathbf{new}}^-$, in Table 1.5. In contrast to the additive modalities, we do not have the arbitrary introduction of modalities to resource propositions. Intuitively, they refer to transitions that involve contraction or expansion of the resource component.

The distinction between the diamond and box modalities is conventional: the diamond modalities refer to the existence of a transition satisfying a property, and the box modalities refer to all transitions satisfying a property. The distinction between the ‘+’ and ‘-’ modalities is more interesting. The positive multiplicatives refer to transitions that involve ‘releasing’ a resource and process component specified by the \vdash relation. The negative modalities refer to transitions that ‘acquire’ or remove a specified resource and process component (we denote this using \vdash , though it could be the same relation).

For the moment, we shall not venture to extend the logic with *e.g.*, the modalities of CTL or modal- μ . Infinitary conjunction and disjunction shall, though, be allowed.

1.4.3 Observable vs. silent actions

Hitherto we have failed to consider how the above modalities may take into account the distinction between observable and silent actions. Two options exist:

$\mathbf{R} \triangleleft [P] \models p$	iff	$\mathbf{R} \triangleleft [P] \in \ \mathbf{p}\ $
$\mathbf{R} \triangleleft [P] \models \varphi \wedge \psi$	iff	$\mathbf{R} \triangleleft [P] \models \varphi$ and $\mathbf{R} \triangleleft [P] \models \psi$
$\mathbf{R} \triangleleft [P] \models \varphi \vee \psi$	iff	$\mathbf{R} \triangleleft [P] \models \varphi$ or $\mathbf{R} \triangleleft [P] \models \psi$
$\mathbf{R} \triangleleft [P] \models \varphi \rightarrow \psi$	iff	for all $\mathbf{R}' \triangleleft [P'] \sqsubseteq \mathbf{R} \triangleleft [P]$ $\mathbf{R}' \triangleleft [P'] \models \varphi$ implies $\mathbf{R}' \triangleleft [P'] \models \psi$
$\mathbf{R} \triangleleft [P] \models \varphi * \psi$	iff	for some $\mathbf{S} \triangleleft [Q], \mathbf{S}' \triangleleft [Q'] \in \mathcal{S}$ [$\mathbf{R} \triangleleft [P] \sqsubseteq (\mathbf{S} \triangleleft [Q]) \circ (\mathbf{S}' \triangleleft [Q'])$ and $\mathbf{S} \triangleleft [Q] \models \varphi$ and $\mathbf{S}' \triangleleft [Q'] \models \psi$]
$\mathbf{R} \triangleleft [P] \models \varphi \multimap \psi$	iff	for all $\mathbf{S} \triangleleft [Q] \in \mathcal{S}$ $\mathbf{S} \triangleleft [Q] \models \varphi$ implies $(\mathbf{R} \triangleleft [P] \circ \mathbf{S} \triangleleft [Q]) \models \psi$
$\mathbf{R} \triangleleft [P] \models \top$	for all	$\mathbf{R} \triangleleft [P] \in \mathcal{S}$
$\mathbf{R} \triangleleft [P] \models \perp$	for any	$\mathbf{R} \triangleleft [P] \in \mathcal{S}$
$\mathbf{R} \triangleleft [P] \models I_P$	iff	for some $\mathbf{R} \in Res.$ $\mathbf{R} \triangleleft [P] \sqsubseteq \mathbf{R} \triangleleft e_P$
$\mathbf{R} \triangleleft [P] \models I_R$	iff	for some $Q \in Prog.$ $\mathbf{R} \triangleleft [P] \sqsubseteq e_R \triangleleft [Q]$
$\mathbf{R} \triangleleft [P] \models I$	iff	$\mathbf{R} \triangleleft [P] \sqsubseteq e_r \triangleleft e_P$

Table 1.3: Kripke semantics of **PBI** (excluding modalities)

$\mathbf{R} \triangleleft [P] \models [A]\varphi$	iff	$\forall \alpha \in A.$ $\mathbf{R} \triangleleft [P] \xrightarrow{\alpha} \mathbf{R} \triangleleft [P']$ implies $\mathbf{R} \triangleleft [P'] \models \varphi$
$\mathbf{R} \triangleleft [P] \models \langle A \rangle \varphi$	iff	$\exists \alpha \in A.$ $\mathbf{R} \triangleleft [P] \xrightarrow{\alpha} \mathbf{R} \triangleleft [P']$ and $\mathbf{R} \triangleleft [P'] \models \varphi$

Table 1.4: Kripke semantics of additive modalities of **PBI**

$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle A \rangle_{\mathbf{new}}^+ \varphi$	iff	$\exists a \in A. \exists \mathbf{S} \triangleleft \llbracket Q \rrbracket \in \mathcal{P}_{\text{ROC}}.$ $a \upharpoonright \mathbf{S} \triangleleft \llbracket Q \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} (\mathbf{R} \triangleleft \llbracket P' \rrbracket) \circ (\mathbf{S} \triangleleft \llbracket Q \rrbracket)$ and $(\mathbf{R} \triangleleft \llbracket P' \rrbracket) \circ (\mathbf{S} \triangleleft \llbracket Q \rrbracket) \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models [A]_{\mathbf{new}}^+ \varphi$	iff	$\forall a \in A. \forall \mathbf{S} \triangleleft \llbracket Q \rrbracket \in \mathcal{P}_{\text{ROC}}.$ $a \upharpoonright \mathbf{S} \triangleleft \llbracket Q \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} (\mathbf{R} \triangleleft \llbracket P' \rrbracket) \circ (\mathbf{S} \triangleleft \llbracket Q \rrbracket)$ implies $(\mathbf{R} \triangleleft \llbracket P' \rrbracket) \circ (\mathbf{S} \triangleleft \llbracket Q \rrbracket) \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle A \rangle_{\mathbf{new}}^- \varphi$	iff	$\exists a \in A. \exists \mathbf{S} \triangleleft \llbracket Q_1 \rrbracket, \mathbf{T} \triangleleft \llbracket Q_2 \rrbracket \in \mathcal{P}_{\text{ROC}}.$ $a \upharpoonright \mathbf{S} \triangleleft \llbracket Q_1 \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \sqsubseteq (\mathbf{S} \triangleleft \llbracket Q_1 \rrbracket) \circ (\mathbf{T} \triangleleft \llbracket Q_2 \rrbracket)$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} \mathbf{T} \triangleleft \llbracket Q_2' \rrbracket$ and $\mathbf{T} \triangleleft \llbracket Q_2' \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models [A]_{\mathbf{new}}^- \varphi$	iff	$\forall a \in A. \forall \mathbf{S} \triangleleft \llbracket Q \rrbracket, \mathbf{T} \triangleleft \llbracket Q_2 \rrbracket \in \mathcal{P}_{\text{ROC}}.$ $a \upharpoonright \mathbf{S} \triangleleft \llbracket Q \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \sqsubseteq (\mathbf{S} \triangleleft \llbracket Q \rrbracket) \circ (\mathbf{T} \triangleleft \llbracket Q_2 \rrbracket)$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} \mathbf{T} \triangleleft \llbracket Q_2' \rrbracket$ implies $\mathbf{T} \triangleleft \llbracket Q_2' \rrbracket \models \varphi$

Table 1.5: Kripke semantics of multiplicative modalities of **PBI**

- The transition system could be augmented with rules to provide

$$\xrightarrow{\sigma} = \xrightarrow{\tau}^* \xrightarrow{\sigma} \xrightarrow{\tau}^*$$

for all observable σ .

- Observable transition modalities could be introduced to complement $\xrightarrow{\sigma}$. They would be just as the modalities presented here, other than requiring an observable action $\xrightarrow{\sigma}$ rather than any type.

We shall adopt the second solution: it allows us to define distinct notions of program equivalence, and otherwise results in a clearer notion of activity in the logic.

1.5 Petri nets in PBI

Petri nets can be used to represent evolving systems. State of the system is described by the distribution of *tokens* to *places*, forming a *marking*. State evolves by the pre-conditions (a set of places) of an action being met (having a token), along with all the post-conditions (a set of places) being empty. Following the action, the places forming its post-conditions are then marked. Several introductions to Petri nets exist, *e.g.*, [Mur89] and the introduction to [MBC⁺95]. Petri nets can be represented graphically: it is conventional to draw

$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle\langle A \rangle\rangle \varphi$	iff	$\exists a \in A.$ $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} \mathbf{R} \triangleleft \llbracket P' \rrbracket$ $\wedge \mathbf{R} \triangleleft \llbracket P' \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \llbracket A \rrbracket \varphi$	iff	$\forall a \in A.$ $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} \mathbf{R} \triangleleft \llbracket P' \rrbracket$ $\wedge \mathbf{R} \triangleleft \llbracket P' \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle\langle A \rangle\rangle_{\mathbf{new}}^+ \varphi$	iff	$\exists a \in A.$ $a \upharpoonright \mathbf{S} \triangleleft \llbracket Q \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} (\mathbf{R} \triangleleft \llbracket P' \rrbracket) \circ (\mathbf{S} \triangleleft \llbracket Q \rrbracket)$ and $(\mathbf{R}' \triangleleft \llbracket P' \rrbracket) \circ (\mathbf{S} \triangleleft \llbracket Q \rrbracket) \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \llbracket A \rrbracket_{\mathbf{new}}^+ \varphi$	iff	$\forall a \in A.$ $a \upharpoonright \mathbf{S} \triangleleft \llbracket Q \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} (\mathbf{R} \triangleleft \llbracket P' \rrbracket) \circ (\mathbf{S} \triangleleft \llbracket Q \rrbracket)$ implies $(\mathbf{R}' \triangleleft \llbracket P' \rrbracket) \circ (\mathbf{S} \triangleleft \llbracket Q \rrbracket) \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle\langle A \rangle\rangle_{\mathbf{new}}^- \varphi$	iff	$\exists a \in A. a \upharpoonright \mathbf{S} \triangleleft \llbracket Q_1 \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \sqsubseteq (\mathbf{S} \triangleleft \llbracket Q_1 \rrbracket) \circ (\mathbf{T} \triangleleft \llbracket Q_2 \rrbracket)$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} \mathbf{T} \triangleleft \llbracket Q'_2 \rrbracket$ and $\mathbf{T} \triangleleft \llbracket Q'_2 \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \llbracket A \rrbracket_{\mathbf{new}}^- \varphi$	iff	$\forall a \in A. a \upharpoonright \mathbf{S} \triangleleft \llbracket Q_1 \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \sqsubseteq (\mathbf{S} \triangleleft \llbracket Q_1 \rrbracket) \circ (\mathbf{T} \triangleleft \llbracket Q_2 \rrbracket)$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} \mathbf{T} \triangleleft \llbracket Q'_2 \rrbracket$ implies $\mathbf{T} \triangleleft \llbracket Q'_2 \rrbracket \models \varphi$

Table 1.6: Kripke semantics of observable modalities of **PBI**

places as circles, tokens as dots, and actions as bars. Pre-conditions of actions are drawn as arrows leading into the action, and post-conditions are drawn as arrows leading away. Figure 1.4 gives some examples.

More formally, a net is a 4-tuple, $N = \langle P, T, pre, post \rangle$. A *marking* M is a function giving the number of tokens in a given place. We shall consider $M : P \rightarrow \{0, 1\}$. Taking \mathcal{M} to be the set of all markings, P is the set of places (circles), T the set of transitions (bars), and $pre, post : T \rightarrow \mathcal{M}$ are functions that return the required markings before and after a given transition (arrows).

As described in [POY02, Pym02], it is possible to encode Petri nets in **BI**. We denote composition of markings M and N , $M + N$, and define:

$$(M + N)p = \begin{cases} 0, & Mp = 0 \wedge Np = 0 \\ 1, & \text{otherwise} \end{cases}.$$

This restricts our attention to the variant of Petri nets where at most one token occurs in any place. $[-]$ shall denote the empty marking (*i.e.*, for all p , $[-]p = 0$). We consider only finite Petri nets, so $\text{dom}(M)$ is finite.




Figure 1.4: 1.4(a) goes to 1.4(b). 1.4(c): no action — blocked by post-condition. 1.4(d): no action — not all pre-conditions met. 1.4(e) goes to 1.4(f)

Thus [POY02] defines a pre-ordered commutative monoid

$$\langle \mathcal{M}, [-], +, \sqsubseteq \rangle.$$

Petri nets, therefore, give semantics to **BI**. We represent action (transitions) between Petri net markings using \sqsubseteq . So, for example, $N \sqsubseteq M$ iff from M a number of transitions lead to a marking N . This interpretation makes \rightarrow , in a sense, a global temporal operator similar to $\forall \mathbf{G}$ in CTL: $m \models \varphi \rightarrow \psi$ is true if in all reachable states, if φ is true then ψ is true. We also lose the ability to describe only the current marking using multiplicative conjunction, $*$.

As mentioned before, we chose to allow only one token per place in each marking. Depending on the application, it would be quite easy to generalise away from this restriction: we would simply take $M, N : P \rightarrow \mathbb{N}$ and $(M+N)p = Mp + Np$, and relax the requirement that the post-condition of a transition be empty before the transition can take place.

It is also possible to encode Petri nets in **PBI**. In essence, all we do is take marked places as resources, and the graph formed by $\langle pre, post, P, T \rangle$ as the process term. Graphically, the process term would simply be the net with no places marked. Writing $N[X]$ to mean a Petri net N that has a sub-net X and \mathbf{M} for a marking (set of marked places), we can define a transition relation to describe the activity of Petri nets. As an example,

$$\text{(STRAN)} : \frac{}{\mathbf{M}, p \triangleleft \llbracket N[p \rightarrow \boxed{t} \rightarrow m] \rrbracket \xrightarrow{t} \mathbf{M}, m \triangleleft \llbracket N[p \rightarrow \boxed{t} \rightarrow m] \rrbracket} (m \notin \mathbf{M})$$

for an arbitrary net $N[p \rightarrow \boxed{t} \rightarrow m]$ with a transition t such that $pre(t) = \{p\}$ and $post(t) = \{m\}$; the pre- and post-conditions in this specific case each comprise of only one place.

With such a transition relation, we can ask questions relating to the consequences of additional places being marked and tokens being removed from places. For example, the semantics gives us that

$$\{ \} \triangleleft \llbracket N \rrbracket \models m \multimap \varphi$$

asks if φ holds of a state in which only m is marked. We may use the modalities to ask of the markings formed by transitions. The correspondence between questions posed of Demos states and Petri nets is so great that the reader is referred to §4.1 for a discussion of the logic presented.

The invariance in the process term need not necessarily be so: we could have a Petri net that changes. Such a system would describe a *dynamic* Petri net [AB96]. Unlike our presentation of additive modalities here, we could perhaps consider additive modalities as referring to transitions in which only the marking changes, and multiplicative modalities as referring to transitions in which the net may change. Further consideration of this, though, is beyond our current scope.

Chapter 2

Demos

2.1 Introduction

Demos [BT01] is a simulation language for modelling processes, executing in parallel, that interact through taking, and then relinquishing, control of resources.

As indicated by the title of the manual [BT01], the value of the language is that its operational semantics is clearly defined [TB01]. It is possible, therefore, to define translations between Demos programs and other systems and then prove the translation correct. For example, Demos programs can be translated to CCS processes [TB94, TB95, BTa, BTb] and to Petri nets [Tof01] (again, if we fail to model time). Because the translation is demonstrably correct, we can have confidence that any results obtained of a Demos model by, say, translating to CCS and then using the Edinburgh Concurrency Workbench [MS], are correct.

We refer the reader to [BT01] for an introduction to the language.

This document, in order to maintain a concise notation, shall consider only a subclass of Demos programs, the syntax of which is given in Backus-Naur form in Figure 2.1. Notably, this does not include variables, or bins or synchronisations with values.

Returning to our intelligence agency example, we may model the system in Demos by the program of Figure 2.2. The trace (omitted) indicates no deadlock.

We shall use the abbreviation ε for the *null* process, a process incapable of any action. Letting \mathcal{P}_{ROG} be the set of correct processes derived from the BNF description, the set of processes is $Code = \{\varepsilon\} \cup \mathcal{P}_{\text{ROG}}$. Writing $\cdot\|\cdot$ for the parallel composition of processes and $\cdot;\cdot$ for the sequential composition of processes (see later), the following structural equivalences should be taken to

$$\begin{aligned}
Prog ::= & \text{Res}(r, n) \mid \text{Bin}(r, n) \mid \\
& \text{Entity}(class, t) \mid \text{Class name} = \{Decl\} \mid \\
& \text{close} \mid Prog; Prog \mid \text{hold}(t) \\
\\
Decl ::= & \text{getR}(r, n) \mid \text{putR}(r, n) \mid \\
& \text{getB}(r, n) \mid \text{putB}(r, n) \mid \\
& \text{sync}(x) \mid \text{getS}(x, n) \mid \text{putS}(x, n) \mid \\
& \text{try } [cond] \text{ then } Decl \text{ etry } [cond] \text{ then } Decl \mid \\
& \text{while } [cond] \text{ do } Decl \mid \text{do } n \text{ } Decl \mid \text{req}[cond] \mid \\
& \text{Entity}(class, t) \mid \text{hold}(t) \mid \\
& Decl; Decl \\
\\
cond ::= & \text{getR}(r, n) \mid \text{getB}(r, n) \mid \text{getS}(x, n) \mid cond, cond
\end{aligned}$$

Figure 2.1: BNF syntax of restriced Demos considered

```

cons ACT_TIME6 4;
cons ACT_TIMECIA 4;

class MI6 = {
  getR(MI6_line, 1); hold(2);
  getR(MI5_line, 1); hold(2);
  getR(FBI_line, 1); hold(2);

  hold(ACT_TIME6);

  putR(FBI_line, 1); putR(MI5_line, 1); putR(MI6_line, 1);
}

class CIA = {
  getR(CIA_line, 1); hold(2);
  getR(FBI_line, 1); hold(2);
  getR(MI5_line, 1); hold(2);

  hold(ACT_TIMECIA);

  putR(MI5_line, 1); putR(FBI_line, 1); putR(CIA_line, 1);
}

Res(CIA_line, 1); Res(FBI_line, 1);
Res(MI6_line, 1); Res(FBI_line, 1);

Entity(CIA, 12*60 + 40); %12:40
Entity(MI6, 12*60 + 43); %12:43

```

Figure 2.2: Demos program for the intelligence agency example

exist:

$$\begin{aligned}
[[P||Q]] &\equiv [[Q||P]] \\
[[P||(Q_1||Q_2)]] &\equiv [[(P||Q_1)||Q_2]] \\
[[P||\varepsilon]] &\equiv [[P]] \\
[[\varepsilon; P]] &\equiv [[\varepsilon]] \\
[[P; \varepsilon]] &\equiv [[P]]
\end{aligned}$$

We shall, though, define the transition system so as to make these equivalences unnecessary in the derivation of transitions.

The reader shall recall that in our definition of **PBI** we required a transition relation to describe the activity of processes. This is presented in the following section, and shall be called σ Demos.

2.2 Transition rules

Notation Let R denote the set of actual resources; r_0, r_1, \dots shall range over R . In the sequel, we shall find it convenient to think of processes as resources. We can define the set of resources (actual and processes)

$$\mathcal{R} \stackrel{\text{def}}{=} \mathcal{P}_+(R) \uplus \mathcal{P}(\mathcal{P}_{\text{ROC}})$$

where $\mathcal{P}_+(R)$ denotes the set of all sub-multisets of R . $\mathbf{R} \in \mathcal{R}$ shall denote a multiset of resources; multiset operations shall be as set operations tagged with a '+'. \mathbf{R}, r shall denote the multiset \mathbf{R} with one unit of resource r added. That is,

$$\mathbf{R}, r \stackrel{\text{def}}{=} \mathbf{R} \uplus \{r\}.$$

We shall write r^n for $\uplus_1^n \{r\}$. A *state* in the transition system shall be represented by a pair, $s \in \mathcal{R} \times \mathcal{P}_{\text{ROC}}$. We let the set of all states be $\Omega \stackrel{\text{def}}{=} \mathcal{R} \times \mathcal{P}_{\text{ROC}}$.

Let the set *Act* be the set of actions¹ that any Demos program may make, ranged over by σ . We can define a transition relation, $\longrightarrow \subseteq \Omega \times \text{Act} \times \Omega$, between programs and the resources they hold; we shall use infix notation for the relation. It is defined, inductively, according to the following set of rules, which essentially give the (exclusively) resource-sensitive, *non-temporal* structural operational semantics of Demos programs.

As well as observable actions, we shall allow *silent* (non-observable) transitions; such transitions shall be written $\mathbf{R} \triangleleft [[P]] \xrightarrow{\tau} \mathbf{R}' \triangleleft [[P']]$ ². Recall that this should decrease the level of non-determinism within the proof system compared to an equivalence relation.

It may be useful later to define a new transition relation,

$$\Longrightarrow \stackrel{\text{def}}{=} \xrightarrow{\tau}^* \xrightarrow{\sigma} \xrightarrow{\tau}^*,$$

¹Silent and observable

²The inspiration for ' τ ', obviously, is CCS

where S^* denotes the reflexive, transitive closure of the relation S . Alternatively, we may define \Longrightarrow inductively:

$$\begin{aligned}
(\text{OBSER}_0) : & \frac{\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket}{\mathbf{R} \triangleleft \llbracket P \rrbracket \Longrightarrow \mathbf{R}' \triangleleft \llbracket P' \rrbracket} \\
(\text{OBSER}_1) : & \frac{\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\tau} \mathbf{R}'' \triangleleft \llbracket P'' \rrbracket \quad \mathbf{R}'' \triangleleft \llbracket P'' \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket}{\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket} \\
(\text{OBSER}_2) : & \frac{\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}'' \triangleleft \llbracket P'' \rrbracket \quad \mathbf{R}'' \triangleleft \llbracket P'' \rrbracket \xrightarrow{\tau} \mathbf{R}' \triangleleft \llbracket P' \rrbracket}{\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket}
\end{aligned}$$

The symbol P (and $P_1, P_2, \dots, Q, \dots$) is used frequently throughout this document. In general, it ranges over processes, including parallel processes. It is also used to range over non-parallel processes — that is, bodies of code. The reader should extract the appropriate choice between these from the syntax of Demos code.

2.3 Transition system σ Demos

We shall write $\mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket$ for cleanly terminated states. The restricted subset of Demos considered here has the useful property that one can deduce what resources are held by a given process by examining the code of the process, under the assumption that the process code is *well formed* (only puts resources that it holds) — for our restricted language, this property can be checked syntactically. Consequently, we do not need to store this information, *e.g.*, in the process term.

The declaration of each process (the **Entity** command) attaches to each non-parallel process a label $\ell \in \text{Label}$, where Label is the set of all process labels. A process P labelled ℓ shall be written $P_{[\ell]}$. For example, if the jobbers in the jobshop example are called Bob and Alice, we would write $\llbracket J_{[\text{Bob}]} \parallel J_{[\text{Alice}]} \rrbracket$. The process label shall be included in the action label of any transition derived by a non-parallel process so that in the process logic we will be able to determine which process is acting. We shall illustrate this in the first rule defined below and in the definition of **Entity**; otherwise, for clarity, it shall be omitted. Moreover, in the sequel, except where this label is of particular interest to us, we shall omit the process labels. We shall assume that σ allows labels to be carried through from sub-derivations.

2.3.1 Resource actions

We commence by defining how processes acquire and release resources in σ Demos. To **get** a resource, a process simply removes the resource from the collection of available resources if it is available. If it is not, the process will not be able to pass the command, thus simulating the process waiting for the resource. Releasing a resource is just a case of replacing the resource in the collection. The rules for bins and actual resources are the same in σ Demos; in Demos itself, however, bins are used to model resources that may be generated or consumed

by processes. For example, a process that must withdraw £50 from a bank account b before executing P should be modelled $\text{getB}(b, 50); P$.

$$(\text{GET}_R^n) : \frac{}{\mathbf{R}, r^n \triangleleft \llbracket \text{getR}(r, n) \rrbracket_{[c]} \xrightarrow{\ell: \text{getR}_r^n} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{PUT}_R^n) : \frac{}{\mathbf{R} \triangleleft \llbracket \text{putR}(r, n) \rrbracket \xrightarrow{\text{putR}_r^n} \mathbf{R}, r^n \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{GET}_B^n) : \frac{}{\mathbf{R}, r^n \triangleleft \llbracket \text{getB}(r, n) \rrbracket \xrightarrow{\text{getB}_r^n} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{PUT}_B^n) : \frac{}{\mathbf{R} \triangleleft \llbracket \text{putB}(r, n) \rrbracket \xrightarrow{\text{putB}_r^n} \mathbf{R}, r^n \triangleleft \llbracket \varepsilon \rrbracket}$$

Note that in the rules (PUT), we have assumed that processes are well formed in that we have not checked that they own the resources that they are returning.

2.3.2 Hold

As we have already stated, our model of Demos programs shall not consider time. Thus the transition for $\text{hold}(t)$ shall be:

$$(\text{NOTIME}) : \frac{}{\mathbf{R} \triangleleft \llbracket \text{hold}(t) \rrbracket \xrightarrow{\tau} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket},$$

effectively skipping the command. Notice that we have used a τ -transition to indicate that this is not observable.³

2.3.3 Initialisation

It is also possible to define initialisation of the program (allowing a little vagueness by assuming that we are passed the code of *class* rather than just its name. This could be resolved by introducing a function $\Sigma : \text{classname} \rightarrow \text{code}$). Again, notice that we have defined a τ -transition.

$$(\text{ENTITY}) : \frac{}{\mathbf{R} \triangleleft \llbracket \text{Entity}(\text{class}, \text{name}, t); P \rrbracket \xrightarrow{\tau_{\text{Entity}}^{\text{class}, \text{name}}} \mathbf{R} \triangleleft \llbracket \text{class}_{[\text{name}]} \rrbracket \parallel P}$$

$$(\text{RES}_n) : \frac{}{\mathbf{R} \triangleleft \llbracket \text{Res}(r, n) \rrbracket \xrightarrow{\tau_{\text{Res}}^{r, n}} \mathbf{R}, r^n \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{BIN}_n) : \frac{}{\mathbf{R} \triangleleft \llbracket \text{Bin}(b, n) \rrbracket \xrightarrow{\tau_{\text{Bin}}^{b, n}} \mathbf{R}, b^n \triangleleft \llbracket \varepsilon \rrbracket}$$

³It may be more useful to split the hold into two transitions, $\mathbf{R} \triangleleft \llbracket \text{hold}(t) \rrbracket \xrightarrow{\text{startstop}} \mathbf{R}$, for model-checking some properties *e.g.*, mutual exclusion over a critical region.

2.3.4 Sequential composition

The sequential composition of program code is relatively straightforward to define: one command is executed, which changes the state of the system, with the rest of the program remaining. c shall range over single commands which may be compound *e.g.*, **while** loops. There are two cases to consider. Firstly, the first command executes completely:

$$(\text{COMP}_{\text{Te}}) : \frac{\mathbf{R} \triangleleft \llbracket c \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \triangleleft \llbracket c; P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_2 \rrbracket}$$

If the command at the head executes to give an intermediate step, we place this back in front:

$$(\text{COMP}_I) : \frac{\mathbf{R} \triangleleft \llbracket c \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_1 \rrbracket}{\mathbf{R} \triangleleft \llbracket c; P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_1; P_2 \rrbracket} \quad (P'_1 \neq \varepsilon)$$

2.3.5 Parallel composition

We say that processes are acting in *parallel* or *concurrently* if more than one process is ‘running’ at the same time. We shall use $\cdot \parallel \cdot$ to represent parallel composition of processes. Each parallel process, providing it is not blocked (waiting for a resource not free in \mathbf{R}), may be executed. At this stage, we choose not to define a notion of *fairness* in execution — see §6.1.

One approach that may be adopted is to only allow one process to execute one step in order to form a transition of the parallel system, forming a so-called *asynchronous* transition system. The choice of which process to execute is arbitrary, or non-deterministic.

$$(\text{PAR}_{I,1}) : \frac{\mathbf{R} \triangleleft \llbracket P_1 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_1 \rrbracket}{\mathbf{R} \triangleleft \llbracket P_1 \parallel P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_1 \parallel P_2 \rrbracket}$$

$$(\text{PAR}_{I,2}) : \frac{\mathbf{R} \triangleleft \llbracket P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_2 \rrbracket}{\mathbf{R} \triangleleft \llbracket P_1 \parallel P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_1 \parallel P'_2 \rrbracket}$$

We must also consider the cases where one of the parallel processes terminates.

$$(\text{PAR}_{\text{Te},1}) : \frac{\mathbf{R} \triangleleft \llbracket P_1 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \triangleleft \llbracket P_1 \parallel P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_2 \rrbracket}$$

$$(\text{PAR}_{\text{Te},2}) : \frac{\mathbf{R} \triangleleft \llbracket P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \triangleleft \llbracket P_1 \parallel P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_1 \rrbracket}$$

Alternatively, we may define a rule schema to allow easy access to parallel processes. One may consider this to be a derived set of rules, or consider the above to be specific instances of this; this is similar to the Expansion Law of CCS.

$$(\text{PAR}_I^i) : \frac{\mathbf{R} \triangleleft \llbracket P_i \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_i \rrbracket}{\mathbf{R} \triangleleft \llbracket \parallel_{j \in I} P_j \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \parallel_{j \in I \setminus \{i\}} P_j \rrbracket} \quad \left(\begin{array}{l} I = \{1, \dots, n\}, \\ n > 1, \\ i \in I \end{array} \right)$$

with a similar rule (PAR_{Te}).

See §5 for consideration of a rule for synchronous parallel execution.

2.3.6 Process synchronisation

A powerful feature of Demos is that it is able to model *synchronisation* between processes. For example, a process representing a train may synchronise with a process representing a station to represent loading/unloading. The train process will, once it has placed itself in the *sync pool* of arrived trains, not be able to perform any action until it is released by the station — after it has been allocated a platform and all the passengers have boarded.

To define synchronisation, we must have a notation for representing processes waiting in a pool and processes holding other processes.

Notation $\overline{\langle P \rangle}_x$ shall represent a process P that has synchronized on a pool x . This shall lie in the resource term \mathbf{R} . When this process is claimed by another process, it shall be written $\langle P \rangle_x$.

A process Q holding a multiset of processes $S = \{ \langle P_1 \rangle_{x_1}, \dots, \langle P_n \rangle_{x_n} \}$ that synchronised on pools x_1, \dots, x_n shall be written $\left[\frac{Q}{S} \right]$.

We see, then, that a process is actually a pair consisting of the remaining code and the processes synchronised on it⁴. Where the collection of processes that synchronised on a process is not affected by a rule (for example, sequential composition does not itself modify the collection of synchronized processes, though a sub-derivation may do so), in general we shall not specify this in the rule — for $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket$ it shall be implicit in $\llbracket P \rrbracket$ and $\llbracket P' \rrbracket$. Otherwise, where the collection is explicitly modified, we shall write rules of the form

$$\mathbf{R} \triangleleft \left[\frac{P}{S} \right] \xrightarrow{\sigma} \mathbf{R}' \triangleleft \left[\frac{P'}{S, s} \right].$$

The upper-case letter S denotes a multiset; the lower-case letter s denotes a single item. The above definition extends in the obvious manner to parallel processes.

$$(\text{SYNC}^x) : \frac{}{\mathbf{R} \triangleleft \llbracket \text{sync}(x); P \rrbracket \xrightarrow{\text{sync}_x} \mathbf{R}, \overline{\langle P \rangle}_x \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{GETS}^x) : \frac{}{\mathbf{R}, \overline{\langle P_i \rangle}_x \triangleleft \left[\frac{\text{getS}(x, n)}{S} \right] \xrightarrow{\text{getS}_x^n} \mathbf{R} \triangleleft \left[\frac{\varepsilon}{S, \langle P_i \rangle_x^n} \right]}$$

$$(\text{PUTS}^x) : \frac{}{\mathbf{R} \triangleleft \left[\frac{\text{putS}(x, n); Q}{S, \langle P_i \rangle_x^n} \right] \xrightarrow{\text{putS}_x^n} \mathbf{R} \triangleleft \left[\frac{Q}{S} \parallel_n P_i \right]}$$

⁴ $\mathcal{P}_{\text{ROC}} = \text{Label} \times \text{Code} \times \mathcal{P}_+(\mathcal{P}_{\text{ROC}})$, where \mathcal{P}_+ is the multiset version of powerset, *Code* is the set of syntactically correct Demos programs and *Label* is the set of process labels. The reader will, of course, recognise that this type, \mathcal{P}_{ROC} , represents the solution of a recursive domain equation.

Though we shall not write implicitly modified synchronized process collections, it is important that the reader recognizes that they exist. So, for example, sequential composition is actually:

$$(\text{COMP}_1) : \frac{\mathbf{R} \triangleleft \left[\frac{P}{S} \right] \xrightarrow{\sigma} \mathbf{R}' \triangleleft \left[\frac{P'}{S'} \right]}{\mathbf{R} \triangleleft \left[\frac{P;Q}{S} \right] \xrightarrow{\sigma} \mathbf{R}' \triangleleft \left[\frac{P';Q'}{S'} \right]}$$

$$(\text{COMP}_{\text{Te}}) : \frac{\mathbf{R} \triangleleft \left[\frac{P}{S} \right] \xrightarrow{\sigma} \mathbf{R}' \triangleleft \left[\frac{\varepsilon}{S'} \right]}{\mathbf{R} \triangleleft \left[\frac{P;Q}{S} \right] \xrightarrow{\sigma} \mathbf{R}' \triangleleft \left[\frac{Q'}{S'} \right]}$$

2.3.7 Non-blocking requests

In order to define valid transition rules for **try** and **while**, we shall need a non-blocking multi-request (a request where the list of actions is atomic — if any one of the requests fails, the whole list fails — and that we can detect failure of); **req** will be similar. We shall use the non-program label $g[\cdot]$ to denote such a request.

When we fail to complete a non-blocking request, we shall indicate this using the symbol \mathcal{FF} .

$$(\text{NBLOCK}_P^T) : \frac{\mathbf{R} \triangleleft \llbracket \text{getS}(x, 1) \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \triangleleft \llbracket g[\text{getS}(x, 1)] \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{NBLOCK}_P^F) : \frac{}{\mathcal{FF}(\mathbf{R} \triangleleft \llbracket g[\text{getS}(x, 1)] \rrbracket)} \left(\neg \exists P. \frac{x}{(P)} \in \mathbf{R} \right)$$

$$(\text{NBLOCK}_R^T) : \frac{\mathbf{R} \triangleleft \llbracket \text{getR}(r, n) \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \triangleleft \llbracket g[\text{getR}(r, n)] \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{NBLOCK}_R^F) : \frac{}{\mathcal{FF}(\mathbf{R} \triangleleft \llbracket g[\text{getR}(r, n)] \rrbracket)} (r^n \notin \mathbf{R})$$

$$(\text{NBLOCK}^T) : \frac{}{\mathbf{R} \triangleleft \llbracket g[] \rrbracket \xrightarrow{\tau} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{NBLOCK}_{1,2}^T) : \frac{\mathbf{R} \triangleleft \llbracket g[e_i] \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket \quad \mathbf{S} \triangleleft \llbracket g[e_j] \rrbracket \xrightarrow{\sigma} \mathbf{S}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket g[e_1, e_2] \rrbracket \xrightarrow{\sigma} \mathbf{R}' \uplus \mathbf{S}' \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{NBLOCK}^F) : \frac{\mathcal{FF}(\mathbf{R} \triangleleft \llbracket g[r_i] \rrbracket)}{\mathcal{FF}(\mathbf{R} \triangleleft \llbracket g[r_1, r_2] \rrbracket)} (i \in \{1, 2\})$$

where $e ::= e, e \mid \text{getR}(r, n) \mid \text{getB}(r, n) \mid \text{getS}(x, 1)$, $r \in R, n \in \mathbb{Z}^+$.

These definitions are relatively straightforward; their value is that they deal with comma-separated lists and define \mathcal{FF} .

2.3.8 Try, while and req

The **try** command introduces conditional choice into Demos. The command **try** $[e_1]$ **then** P_1 **etry** $[e_2]$ **then** P_2 attempts the atomic list of actions e_1 . If this succeeds, P_1 shall be executed. Otherwise, e_2 is tested to see if P_2 should be executed. If none of the e_i pass, the process waits (no derivable transition) until one does. If more than one pass at the same time, the earliest P_i shall be selected.

Hence we can define the rules for **try**:

$$(\text{TRY}_{\text{Te}}^{\text{T}}) : \frac{\mathbf{R} \triangleleft \llbracket g[e_1] \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \triangleleft \llbracket \text{try } [e_1] \text{ then } P \text{ etry } [e_2] \text{ then } Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P \rrbracket}$$

$$(\text{TRY}_{\text{Te}}^{\text{F}}) : \frac{\mathcal{FF}(\mathbf{R}, g[e_1]) \quad \mathbf{R} \triangleleft \llbracket g[e_2] \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \triangleleft \llbracket \text{try } [e_1] \text{ then } P \text{ etry } [e_2] \text{ then } Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket Q \rrbracket}$$

The **while** construct defines a loop: while an expression is true, a body of code will be executed. Should the expression be false at the beginning of the loop, iteration will stop. The **req** command simply defines an atomic series of actions.

$$(\text{WHILE}_{\text{Te}}^{\text{T}}) : \frac{\mathbf{R} \triangleleft \llbracket g[e] \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \triangleleft \llbracket \text{while } [e] \text{ do } Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket Q; \text{while } [e] \text{ do } Q \rrbracket}$$

$$(\text{WHILE}^{\text{F}}) : \frac{\mathcal{FF}(g[e])}{\mathbf{R} \triangleleft \llbracket \text{while } [e] \text{ do } Q \rrbracket \xrightarrow{\sigma} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket}$$

$$(\text{REQ}_{\text{Te}}) : \frac{\mathbf{R} \triangleleft \llbracket g[e] \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R} \triangleleft \llbracket \text{req}[e] \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \varepsilon \rrbracket}$$

2.3.9 Repeat n -times

Assume that n is an integer constant. We may then define the repetition command, **do**, that repeats a body n times.

$$(\text{DO}_n) : \frac{}{\mathbf{R} \triangleleft \llbracket \text{do } n \text{ } P \rrbracket \xrightarrow{\tau} \mathbf{R} \triangleleft \llbracket P; \text{do } (n-1) \text{ } P \rrbracket} \quad (n > 0)$$

$$(\text{DO}_0) : \frac{}{\mathbf{R} \triangleleft \llbracket \text{do } n \text{ } P \rrbracket \xrightarrow{\tau} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket} \quad (n \leq 0)$$

2.4 Example derivations

We now give some examples to show how transitions may be derived by the presented system.

Example 1: Jobshop

This example refers to the ‘jobshop.d2000’ program supplied with Demos. The code is given in full in Figure 2.3. We only show one of the numerous possible derivation sequences, and omit the derivations that go through the instantiation of the entities, resources and bins [see Example 3]. Hammers shall be written h and mallets m .

We shall first introduce some derived rules to make our derivations slightly more succinct. The first allows us to skip some explanation of how we execute the parallel composition of sequential processes.

$$\frac{\frac{\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket}{\mathbf{R} \triangleleft \llbracket P; Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'; Q \rrbracket} (\text{COMP}_{1,1})}{\mathbf{R} \triangleleft \llbracket P; Q \rrbracket P_2 \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'; Q \rrbracket P_2} (\text{PAR}_{1,1})}$$

$$\equiv$$

$$\frac{\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket}{\mathbf{R} \triangleleft \llbracket P; Q \rrbracket P_2 \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'; Q \rrbracket P_2} (\text{PCOMP}_1^2)$$

etc.

The second simplifies how we can write a simple ‘get’ from a bin.

$$\frac{\frac{\mathbf{R}, b \triangleleft \llbracket \text{getB}(b, 1) \rrbracket \xrightarrow{\sigma} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R}, b \triangleleft \llbracket g[\text{getB}(b, 1)] \rrbracket \xrightarrow{\sigma} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket} (\text{NBLOCK}_B^T)}{\mathbf{R}, b \triangleleft \llbracket g[\text{getB}(b, 1)] \rrbracket \xrightarrow{\sigma} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket} (\text{NBLOCK}_B^1)}$$

$$\equiv$$

$$\frac{\mathbf{R}, b \triangleleft \llbracket g[\text{getB}(b, 1)] \rrbracket \xrightarrow{\sigma} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket}{\mathbf{R}, b \triangleleft \llbracket g[\text{getB}(b, 1)] \rrbracket \xrightarrow{\sigma} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket} (\text{NBLOCK}_B^1)$$

etc.

We commence by showing how the while-loop of the worker process is expanded, and how a bin item is placed in the available resource collection.

$$\frac{\frac{\frac{\frac{\lambda m, h \int \triangleleft \llbracket g[] \rrbracket \xrightarrow{\tau} \lambda m, h \int \triangleleft \llbracket \varepsilon \rrbracket}{\lambda m, h \int \triangleleft \llbracket W \rrbracket \xrightarrow{\tau} \lambda m, h \int \triangleleft \llbracket \text{putB}(easy, 1); \dots; W \rrbracket} (\text{NBLOCK}^T)}{\lambda m, h \int \triangleleft \llbracket J \rrbracket J \rrbracket W \rrbracket \xrightarrow{\tau} \lambda m, h \int \triangleleft \llbracket J \rrbracket J \rrbracket \text{putB}(easy, 1); \dots; W \rrbracket} (\text{WHILE}_{Te}^T)}{\lambda m, h \int \triangleleft \llbracket J \rrbracket J \rrbracket W \rrbracket \xrightarrow{\tau} \lambda m, h \int \triangleleft \llbracket J \rrbracket J \rrbracket \text{putB}(easy, 1); \dots; W \rrbracket} (\text{PAR}_{1,3})}$$

$$\frac{\frac{\lambda m, h \int \triangleleft \llbracket \text{putB}(easy, 1) \rrbracket \xrightarrow{\text{putB}} \lambda m, h \int \triangleleft \llbracket \varepsilon \rrbracket}{\lambda m, h \int \triangleleft \llbracket J \rrbracket J \rrbracket \text{putB}(easy, 1); \dots; W \rrbracket \xrightarrow{\text{putB}} \lambda m, h, easy \int \triangleleft \llbracket J \rrbracket J \rrbracket \text{hold}(1); \dots; W \rrbracket} (\text{GET}_B^1)}{\lambda m, h \int \triangleleft \llbracket J \rrbracket J \rrbracket \text{putB}(easy, 1); \dots; W \rrbracket \xrightarrow{\text{putB}} \lambda m, h, easy \int \triangleleft \llbracket J \rrbracket J \rrbracket \text{hold}(1); \dots; W \rrbracket} (\text{PCOMP}_3^3)$$

where J is the code of the `Jobber` class, W is the code of the `Worker` class, and

$K \equiv \text{try } [\text{getB}(easy, 1)] \text{ then trace('Easy')} \text{ etry } [\dots] \text{ then } \dots; J.$

```

Class Jobber=
{While []
  {try [getB(easy,1)] then {Trace("Doing easy");hold(5);}
  etry [getB(mid,1)]
    then {try [getR(mallet,1)]
          then {hold(20);
                trace("Done mid");
                putR(mallet,1);
                }
          etry [getR(hammer,1)]
            then {hold(10);
                  trace("Done mid");
                  putR(hammer,1);
                  }
          }
    etry [getB(hard,1)]
      then {getR(hammer,1);
            hold(20);
            trace("Done hard");
            putR(hammer,1);
            }
      }
}

Class Work=
{while []
  {putB(easy,1);
  hold(1);
  putB(hard,1);
  hold(1);
  putB(mid,1);
  hold(25);
  }
}

Res(mallet,1);
Res(hammer,1);
Bin(easy,0);
Bin(mid,0);
Bin(hard,0);
Entity(J1,Jobber,0);
Entity(J2,Jobber,0);
Entity(W,Work,0);
Hold(200);
Hold(0);
Close;

```

Figure 2.3: Jobber source code

We can now silently remove the `hold` from the head of the worker process, and expand one of the Jobber's while loops.⁵

$$\begin{array}{c}
\frac{}{m, h, \text{easy} \triangleleft \llbracket \text{hold}(1); \dots \rrbracket} \xrightarrow{\tau} m, h, \text{easy} \triangleleft \llbracket \text{putB}(\text{hard}, 1); \dots \rrbracket} \text{(NOTIME)} \\
\frac{}{m, h, \text{easy} \triangleleft \llbracket J \llbracket \text{hold}(1); \dots \rrbracket \rrbracket} \xrightarrow{\tau} m, h, \text{easy} \triangleleft \llbracket J \llbracket \text{putB}(\text{hard}, 1); \dots \rrbracket \rrbracket} \text{(PAR}_{I,2}) \\
\frac{}{m, h, \text{easy} \triangleleft \llbracket J \llbracket J \llbracket \text{hold}(1); \dots \rrbracket \rrbracket \rrbracket} \xrightarrow{\tau} m, h, \text{easy} \triangleleft \llbracket J \llbracket J \llbracket \text{putB}(\text{hard}, 1); \dots \rrbracket \rrbracket \rrbracket} \text{(PAR}_{I,2}) \\
\\
\frac{}{m, h, \text{easy} \triangleleft \llbracket g \llbracket \rrbracket \rrbracket} \xrightarrow{\tau} m, h, \text{easy} \triangleleft \llbracket \llbracket \rrbracket \rrbracket} \text{(NBLOCK}^T) \\
\frac{}{m, h, \text{easy} \triangleleft \llbracket J \rrbracket} \xrightarrow{\tau} m, h, \text{easy} \triangleleft \llbracket K \rrbracket} \text{(WHILE}_{Te}^T) \\
\frac{}{m, h, \text{easy} \triangleleft \llbracket J \llbracket J \llbracket \text{putB}(\text{hard}, 1); \dots \rrbracket \rrbracket \rrbracket} \xrightarrow{\tau} m, h, \text{easy} \triangleleft \llbracket K \llbracket J \llbracket \text{putB}(\text{hard}, 1); \dots \rrbracket \rrbracket \rrbracket} \text{(PAR}_{I,1})
\end{array}$$

Thus we could write

$$\{ m, h \} \triangleleft \llbracket J \llbracket J \llbracket W \rrbracket \rrbracket \rrbracket \xrightarrow{\text{putB}} \{ m, h, \text{easy} \} \triangleleft \llbracket K \llbracket J \llbracket \text{putB}(\text{hard}, 1); \dots; W \rrbracket \rrbracket \rrbracket$$

⁵We shall adopt a syntactic equivalence that allows us to drop the brackets denoting multiset.

The Jobber process with its **while** loop expanded is now in a position to start a job. The only type of job that it is able to get is an easy one.

$$\begin{array}{c}
\frac{}{\frac{}{\frac{}{m, h, \text{easy} \triangleleft \llbracket g[\text{getB}(\text{easy}, 1) \rrbracket \xrightarrow{\text{getB}} m, h \triangleleft \square} \text{(NBLOCK}_B^1)}{m, h, \text{easy} \triangleleft \llbracket \text{try} [\text{getB}(\text{easy}, 1) \rrbracket \text{ then trace}(\text{'Easy'}) \text{ etry} [\dots] \text{ then } \dots \rrbracket \xrightarrow{\text{getB}} m, h \triangleleft \llbracket \text{trace}(\text{'Easy'}) \rrbracket} \text{(TRY}_{T_e}^1)}{m, h, \text{easy} \triangleleft \llbracket \text{try} [\text{getB}(\text{easy}, 1) \rrbracket \text{ then trace}(\text{'Easy'}) \text{ etry} [\dots] \text{ then } \dots \rrbracket \xrightarrow{\text{getB}} m, h \triangleleft \llbracket \text{trace}(\text{'Easy'}) \rrbracket} \text{(PCOMP}_3^1)} \\
\frac{}{m, h, \text{easy} \triangleleft \llbracket K \parallel J \parallel \text{putB}(\text{hard}, 1); \dots \rrbracket \xrightarrow{\text{getB}} m, h \triangleleft \llbracket \text{trace}(\text{'Easy'}) \rrbracket ; J \parallel \text{putB}(\text{hard}, 1); \dots \rrbracket}
\end{array}$$

We now execute the **trace** command — it performs no visible resource action, so it is silent.

$$\frac{}{\frac{}{\frac{}{m, h \triangleleft \llbracket \text{trace}(\text{'Easy'}) \rrbracket \xrightarrow{\tau} m, h \triangleleft \square} \text{(TRACE)}}{m, h \triangleleft \llbracket \text{trace}(\text{'Easy'}) \rrbracket ; J \parallel \text{putB}(\text{hard}, 1); \dots \rrbracket \xrightarrow{\tau} m, h \triangleleft \llbracket J \parallel \text{putB}(\text{hard}, 1); \dots \rrbracket} \text{(PCOMP}_3^1)}$$

Another Jobber's **while**-loop is expanded.

$$D_4 : \frac{}{\frac{}{\frac{}{m, h \triangleleft \llbracket g \rrbracket \xrightarrow{\tau} m, h \triangleleft \square} \text{(NBLOCK}^T)}{m, h \triangleleft \llbracket J \rrbracket \xrightarrow{\tau} m, h \triangleleft \llbracket K \rrbracket} \text{(WHILE}_{T_e}^T)}{m, h \triangleleft \llbracket J \parallel \text{putB}(\text{hard}, 1); \dots \rrbracket \xrightarrow{\tau} m, h \triangleleft \llbracket K \parallel \text{putB}(\text{hard}, 1); \dots \rrbracket} \text{(PAR}_1^3)}$$

The Worker process releases a hard job.

$$\frac{}{\frac{}{\frac{}{m, h \triangleleft \llbracket \text{putB}(\text{hard}, 1) \rrbracket \xrightarrow{\text{putB}} m, h, \text{hard} \triangleleft \square} \text{(PUT}_B^1)}{m, h \triangleleft \llbracket \text{putB}(\text{hard}, 1); \dots \rrbracket \xrightarrow{\text{putB}} m, h, \text{hard} \triangleleft \llbracket \text{hold}(1); \dots \rrbracket} \text{(COMP}_{T_e})}{m, h \triangleleft \llbracket K \parallel J \parallel \text{putB}(\text{hard}, 1); \dots \rrbracket \xrightarrow{\text{putB}} m, h, \text{hard} \triangleleft \llbracket K \parallel J \parallel \text{hold}(1); \dots \rrbracket} \text{(PAR}_3^3)}$$

The `hold(1)` at the head of the Worker process is now removed.

$$\frac{\frac{\frac{}{m, h, hard \triangleleft \llbracket \text{hold}(1) \rrbracket} \xrightarrow{\tau} m, h, hard \triangleleft \llbracket \square \rrbracket} \text{(NOTIME)}}{m, h, hard \triangleleft \llbracket \text{hold}(1); \dots \rrbracket} \xrightarrow{\tau} m, h, hard \triangleleft \llbracket \text{putB}(mid, 1); \dots \rrbracket} \text{(COMP}_{\text{Te}})}{m, h, hard \triangleleft \llbracket K \parallel J \parallel \text{hold}(1); \dots \rrbracket} \xrightarrow{\tau} m, h, hard \triangleleft \llbracket K \parallel J \parallel \text{putB}(mid, 1); \dots \rrbracket} \text{(PAR}_3^3)$$

The Worker process now releases a medium-difficulty job.

$$\frac{\frac{}{m, h, hard \triangleleft \llbracket \text{putB}(mid, 1) \rrbracket} \xrightarrow{\text{putB}} m, h, hard, mid \triangleleft \llbracket \square \rrbracket} \text{(PUT}_B^1)}{m, h, hard \triangleleft \llbracket K \parallel J \parallel \text{putB}(mid, 1); \dots \rrbracket} \xrightarrow{\text{putB}} m, h, hard, mid \triangleleft \llbracket K \parallel J \parallel \text{hold}(25); \dots \rrbracket} \text{(PCOMP}_3^3)$$

By a derivation similar to D_4 ,

$$m, h, hard, mid \triangleleft \llbracket K \parallel J \parallel \text{hold}(25); \dots \rrbracket \xrightarrow{\tau} m, h, hard, mid \triangleleft \llbracket K \parallel K \parallel \text{hold}(25); \dots \rrbracket$$

A process now tries to get an easy job, fails, and so gets a medium-difficulty job.

$$\frac{\frac{\frac{}{\mathcal{FF}(m, h, hard, mid \triangleleft \llbracket g[\text{getB}(easy, 1)] \rrbracket} \text{(easy)} \notin \{m, h, hard, mid\}}{m, h, hard, mid \triangleleft \llbracket \text{try} [\text{getB}(mid, 1)] \text{ then } \dots \text{ etry} [\text{getB}(m, 1)] \text{ then } \dots \rrbracket} \xrightarrow{\text{getB}} m, h, hard \triangleleft \llbracket \text{try} [\text{getR}(h, 1)] \text{ then } \dots \rrbracket} \text{(NBLOCK}_B^1)}{m, h, hard, mid \triangleleft \llbracket K \parallel K \parallel \text{hold}(25); \dots \rrbracket} \xrightarrow{\text{getB}} m, h, hard \triangleleft \llbracket \text{try} [\text{getR}(h, 1)] \text{ then } \dots \text{ etry} [\text{getR}(h, 1)] \text{ then } \dots \rrbracket} \text{(TRY}_{\text{Te}}^F)}{m, h, hard, mid \triangleleft \llbracket K \parallel K \parallel \text{hold}(25); \dots \rrbracket} \xrightarrow{\text{getB}} m, h, hard \triangleleft \llbracket \text{try} [\text{getR}(h, 1)] \text{ then } \dots \text{ etry} [\text{getR}(h, 1)] \text{ then } \dots \rrbracket} \text{(PCOMP}_3^1)$$

We build a derivation tree similar to the one above (other than there is no failure) to show the same Jobber process getting the mallet it needs to complete its medium-difficulty job.

$$\begin{array}{l}
 m, h, hard \triangleleft \llbracket \text{try } [\text{getR}(m, 1)] \text{ then } \{\text{hold}(20); \text{trace}('Mid')\} \text{ etry } [\text{getR}(h, 1)] \text{ then } \dots; J \parallel K \parallel \text{hold}(25); \dots \rrbracket \\
 \xrightarrow{\text{getR}} h, hard \triangleleft \llbracket \text{hold}(20); \text{trace}('Mid'); J \parallel K \parallel \text{hold}(25); \dots \rrbracket
 \end{array}$$

from which we obtain

$$h, hard \triangleleft \llbracket \text{hold}(20); \text{trace}('Mid'); J \parallel K \parallel \text{hold}(25); \dots \rrbracket \xrightarrow{\tau} h, hard \triangleleft \llbracket J \parallel K \parallel \text{hold}(25); \dots \rrbracket$$

Example 2: Semaphore

This simple example considers a synchronisation-based semaphore and may be found in the Demos distribution as `me4.d2000`, though, for conciseness, we ignore the third user. The code is given again in Figure 2.4. Again, we omit the initial declarations of entities *etc.*

We first define, for convenience:

$$\begin{aligned}
A &\equiv \text{getS}(semS, 1); \text{hold}(5); \text{trace}('have'); \text{putS}(semS, 1) \\
S &\equiv \text{sync}(semS); \underbrace{\text{trace}('back')}_{Tr} \\
U &\equiv \text{the code of User} \\
&\equiv \text{while } [] \text{ do } A \\
Sem &\equiv \text{the code of Sem} \\
&\equiv \text{while } [] \text{ do } S
\end{aligned}$$

We shall expand the while loops that control the processes.

$$D_1 : \frac{\frac{\frac{\lambda \int \triangleleft \llbracket g \rrbracket \xrightarrow{\tau} \lambda \int \triangleleft \llbracket \varepsilon \rrbracket} \text{(NBLOCK}_T\text{)}}{\lambda \int \triangleleft \llbracket U \rrbracket \xrightarrow{\tau} \lambda \int \triangleleft \llbracket A; U \rrbracket} \text{(WHILE}_{Te}^T)}}{\lambda \int \triangleleft \llbracket U \parallel U \parallel Sem \rrbracket \xrightarrow{\tau} \lambda \int \triangleleft \llbracket A; U \parallel U \parallel Sem \rrbracket} \text{(PAR}_1^3\text{)}}$$

Very similar derivations justify

$$\lambda \int \triangleleft \llbracket A; U \parallel U \parallel Sem \rrbracket \xrightarrow{\tau} \lambda \int \triangleleft \llbracket A; U \parallel A; U \parallel Sem \rrbracket$$

and

$$\lambda \int \triangleleft \llbracket A; U \parallel A; U \parallel Sem \rrbracket \xrightarrow{\tau} \lambda \int \triangleleft \llbracket A; U \parallel A; U \parallel S; Sem \rrbracket.$$

$$D_2 : \frac{\frac{\lambda \int \triangleleft \llbracket S; Sem \rrbracket \xrightarrow{\text{sync}} \lambda \int \triangleleft \llbracket (Tr; Sem) \rrbracket \int \triangleleft \llbracket \varepsilon \rrbracket} \text{(SYNC}^x\text{)}}{\lambda \int \triangleleft \llbracket A; U \parallel A; U \parallel S; Sem \rrbracket \xrightarrow{\text{sync}} \lambda \int \triangleleft \llbracket (Tr; Sem) \rrbracket \int \triangleleft \llbracket A; U \parallel A; U \rrbracket} \text{(PAR}_3^2\text{)}}$$

We shall now allow the left user process to take the semaphore:

$$D_3 : \frac{\frac{\frac{\lambda \int \triangleleft \llbracket (Tr; Sem) \rrbracket \int \triangleleft \llbracket A; U \rrbracket \xrightarrow{\text{getS}} \lambda \int \triangleleft \llbracket \frac{\text{hold}(5); \dots; U}{(Tr; Sem)_{semS}} \rrbracket} \text{(GETS}^x\text{)}}{\lambda \int \triangleleft \llbracket (Tr; Sem) \rrbracket \int \triangleleft \llbracket A; U \parallel A; U \rrbracket \xrightarrow{\text{getS}} \lambda \int \triangleleft \llbracket \frac{\text{hold}(5); \dots; U}{(Tr; Sem)_{semS}} \parallel A; U \rrbracket} \text{(PAR}_1^1\text{)}}$$

As there are now no processes waiting in *semS*, the right user process cannot proceed, and so the critical region is not entered.

```

class sem=
{while []
  {sync(semS);
   Trace("Im back %n");
  }
}

Entity(theSem,sem,0);

class User=
{while []
  {getS(semS,1);
   hold(5);
   Trace("Have %s %n");
   putS(semS,1);
  }
}

Entity(U1,User,0);
Entity(U2,User,0);
Hold(50);
Entity(U3,User,0);
Hold(50);
HOLD(0);
CLOSE;

```

Figure 2.4: Semaphore source code

Following two τ -transitions to skip the `hold(5); trace('have')` part of the active user process, we may re-activate the semaphore process:

$$D_4 : \frac{\frac{\lambda \int \triangleleft \left[\frac{\text{putS}(semS, 1); U}{(Tr; Sem)_{semS}} \right] \xrightarrow{\text{putS}} \lambda \int \triangleleft [U || Tr; Sem]}{\lambda \int \triangleleft \left[\frac{\text{putS}(semS, 1); U}{(Tr; Sem)_{semS}} || A; U \right] \xrightarrow{\text{putS}} \lambda \int \triangleleft [U || Tr; Sem || A; U]} \text{(PAR}_1^{\downarrow})}{\text{(PUTS}^x)}$$

A τ -transition will remove the `Tr` part of the semaphore process (which is just a `trace` command).

We may derive all following states using similar derivation trees. So, starting again, we could obtain this sequence:

$$\begin{aligned} \lambda \int \triangleleft [U || U || Sem] &\xrightarrow{\tau} \lambda \int \triangleleft [A; U || U || Sem] \\ &\xrightarrow{\tau} \lambda \int \triangleleft [A; U || A; U || Sem] \\ &\xrightarrow{\text{sync}} \lambda \int \triangleleft [A; U || A; U || S; Sem] \\ &\xrightarrow{\text{getS}} \lambda \frac{semS}{(Tr; Sem)} \int \triangleleft [A; U || A; U] \\ &\xrightarrow{\text{putS}} \lambda \int \triangleleft \left[\frac{\text{hold}(5); \dots; U}{(Tr; Sem)_{semS}} || A; U \right] \\ &\xrightarrow{\tau, \tau} \lambda \int \triangleleft \left[\frac{\text{putS}(semS, 1); U}{(Tr; Sem)_{semS}} || A; U \right] \\ &\xrightarrow{\tau} \lambda \int \triangleleft [U || A; U || Tr; Sem] \\ &\xrightarrow{\tau} \lambda \int \triangleleft [U || A; U || Sem] \\ &\xrightarrow{\tau} \lambda \int \triangleleft [U || A; U || S; Sem] \\ &\xrightarrow{\text{sync}} \lambda \frac{semS}{(Tr; Sem)} \int \triangleleft [U || A; U] \\ &\xrightarrow{\text{getS}} \lambda \int \triangleleft \left[U || \frac{\text{hold}(5); \dots; U}{(Tr; Sem)_{semS}} \right] \\ &\xrightarrow{\text{putS}} \lambda \int \triangleleft \left[A; U || \frac{\text{hold}(5); \dots; U}{(Tr; Sem)_{semS}} \right] \\ &\xrightarrow{\tau, \tau} \lambda \int \triangleleft \left[A; U || \frac{\text{putS}(semS, 1); U}{(Tr; Sem)_{semS}} \right] \\ &\vdots \end{aligned}$$

Example 3: Deadlock

This example shows the transition system becoming deadlocked and gives a suggestion of how this may be expressed in the process logic. This time, we shall show processes and resources being created.

The code for this (simple) example may be found in Figure 2.5. In the following derivations, we shall abbreviate the complete code of P1 by $P_1 \equiv \text{getR}(a, 1); Q_1$ and of P2 by $P_2 \equiv \text{getR}(b, 1); Q_2$.

```

Res(a,1);
Res(b,1);

Class P1={
  getR(a,1);
  getR(b,1);
  hold(5);
  putR(a,1);
  putR(b,1);
}

class P2={
  getR(b,1);
  getR(a,1);
  hold(6);
  putR(a,1);
  putR(b,1);
}

Entity(P1,P1,0); Entity(P2,P2,0);

```

Figure 2.5: Deadlock source code

$$D_1 : \frac{\frac{\lambda \int \triangleleft \llbracket \text{Res}(a,1) \rrbracket \xrightarrow{\tau} \lambda a \int \triangleleft \llbracket \text{Res}(a,1) \rrbracket}}{\lambda \int \triangleleft \llbracket \text{Res}(a,1); \text{Res}(b,1); \dots \rrbracket \xrightarrow{\tau} \lambda a \int \triangleleft \llbracket \text{Res}(b,1); \dots \rrbracket}} (\text{COMP}_T)}{\lambda \int \triangleleft \llbracket \text{Res}(a,1) \rrbracket \xrightarrow{\tau} \lambda a \int \triangleleft \llbracket \text{Res}(a,1) \rrbracket}} (\text{RES}_1)$$

Similarly,

$$\lambda a \int \triangleleft \llbracket \text{Res}(b,1); \text{Entity}(P1,0); \dots \rrbracket \xrightarrow{\tau} \lambda a, b \int \triangleleft \llbracket \text{Entity}(P1,0); \dots \rrbracket$$

$$D_2 : \frac{\lambda a, b \int \triangleleft \llbracket \text{Entity}(P1,0); \text{Entity}(P2,0) \rrbracket \xrightarrow{\tau} \lambda a, b \int \triangleleft \llbracket P1 \parallel \text{Entity}(P2,0) \rrbracket}}{\lambda a, b \int \triangleleft \llbracket \text{Entity}(P1,0); \text{Entity}(P2,0) \rrbracket \xrightarrow{\tau} \lambda a, b \int \triangleleft \llbracket P1 \parallel \text{Entity}(P2,0) \rrbracket}} (\text{ENTITY}_I)$$

$$D_3 : \frac{\frac{\lambda a, b \int \triangleleft \llbracket \text{Entity}(P2,0) \rrbracket \xrightarrow{\tau} \lambda a, b \int \triangleleft \llbracket P2 \rrbracket}}{\lambda a, b \int \triangleleft \llbracket P1 \parallel \text{Entity}(P2,0) \rrbracket \xrightarrow{\tau} \lambda a, b \int \triangleleft \llbracket P1 \parallel P2 \rrbracket}} (\text{PAR}_I)}{\lambda a, b \int \triangleleft \llbracket \text{Entity}(P2,0) \rrbracket \xrightarrow{\tau} \lambda a, b \int \triangleleft \llbracket P2 \rrbracket}} (\text{ENTITY}_T)$$

$$D_4 : \frac{\frac{\lambda a, b \int \triangleleft \llbracket P1 \rrbracket \xrightarrow{\text{getR}} \lambda b \int \triangleleft \llbracket Q1 \rrbracket}}{\lambda a, b \int \triangleleft \llbracket P1 \parallel P2 \rrbracket \xrightarrow{\text{getR}} \lambda b \int \triangleleft \llbracket Q1 \parallel P2 \rrbracket}} (\text{COMP}_I)}{\lambda a, b \int \triangleleft \llbracket P1 \rrbracket \xrightarrow{\text{getR}} \lambda b \int \triangleleft \llbracket Q1 \rrbracket}} (\text{GETR}_1)$$

$$D_5 : \frac{\frac{\lambda b \int \triangleleft \llbracket P2 \rrbracket \xrightarrow{\text{getR}} \lambda \int \triangleleft \llbracket Q2 \rrbracket}}{\lambda b \int \triangleleft \llbracket Q1 \parallel P2 \rrbracket \xrightarrow{\text{getR}} \lambda \int \triangleleft \llbracket Q1 \parallel Q2 \rrbracket}} (\text{COMP}_I)}{\lambda b \int \triangleleft \llbracket P2 \rrbracket \xrightarrow{\text{getR}} \lambda \int \triangleleft \llbracket Q2 \rrbracket}} (\text{GETR}_1)$$

At this stage, we would not be able to derive any σ -transition. As we still have code left to execute (formally, $\llbracket Q_1 \parallel Q_2 \rrbracket \neq \llbracket \varepsilon \rrbracket$ — we have not terminated cleanly), we can deduce that the system has deadlocked.

Chapter 3

Correctness of transition system

3.1 Introduction

We now consider various properties of the transition system presented above. These properties relate to the correctness of the system with respect to the operational semantics of Demos [TB01].

The latest version of Demos, Demos 2000, has a large command set. Pragmatism leads us to consider a small subclass of Demos 2000 programs — as we did above by not considering variables or value-resources — namely, those written in π Demos with the Demos 2000 declaration format. π Demos is defined, with its operational semantics, in [BT93]. Specifically, this avoids the notational burden of considering synchronization between processes and increases the brevity of our proofs.

The properties considered shall relate to the validity (correctness) of our transition system. As shall be discussed, we shall not be able to demonstrate the soundness of the transition system with respect to the operational semantics in the familiar sense. This is to be expected: we intentionally made the transition system non-time sensitive, and so transitions derived may not be possible according to the temporal (time-based) constraints actually in place.

π Demos allows us to consider the essential aspect of Demos programs: the interaction of processes through shared resources. In essence, there are three main commands:

getR The command `getR(r , 1)` acquires one item of r from the globally available resources, thus making it unavailable to other processes. If r is not available at the time of request, the process will wait until it is available.

putR The command `putR(r , 1)` releases one item of r , making it available to all other processes. If a process is waiting for the release of r , it will be activated. Otherwise, it will be placed in the resource pool.

hold The command `hold(t)` signifies that the process waits T time units before resuming execution.

We shall assume that processes have been initialised correctly; failing to make this assumption adds little to the presentation. Therefore, we will not consider *Entity etc.*

State in the operational semantics is represented by a 3-tuple:

$$S = \langle \mathcal{EL}, \mathcal{RL}, \Sigma \rangle.$$

- Σ is a function that records the various syntactic definitions found within the system. It will contain, for example, class definitions and resource labels. As we are not going to consider any commands that manipulate Σ (e.g., `class` and `cons`), we shall frequently omit it when discussing state.
- \mathcal{RL} is the resource list, a list of 3-tuples $\langle id, avail, \mathcal{WL} \rangle$. id is the resource name or label; $avail \in \{\mathbf{tt}, \mathbf{ff}\}$ is a boolean indicating whether the resource is taken; and \mathcal{WL} is a list of processes waiting on the resource, non-empty only if not *avail*.
- \mathcal{EL} is the event list, $[(P, PD(code, attrs, t))]$. It is a list of processes not waiting on any resource. Processes are represented by a pair comprising a label and the process descriptor. The process descriptor is another 3-tuple, comprising the code remaining to be executed, the resources held by the process, and the simulation time at which the next command is to be executed.

The reader is referred to [BT93] for the formal presentation of the operational semantics.

In the sequel, to match the account presented in [BT93], we shall only consider one-item get / put, with each identifier in the resource list being unique (that is, only one of any resource shall exist).

3.2 Simulation

We shall represent the globally *available* resources in the transition system by \mathbf{R} . Process terms shall be bracketed $\llbracket P \rrbracket$. The state of a program in the transition system is of the form $\mathbf{R} \triangleleft \llbracket P \rrbracket$; an equivalent representation in the operational semantics shall be written $f(\mathbf{R} \triangleleft \llbracket P \rrbracket)$. Transitions derived according to the operational semantics shall be denoted \Longrightarrow (note that this is unlabelled, to distinguish it from observable transitions made in σ Demos); as usual, single-step transitions in the transition system shall be denoted $\xrightarrow{\sigma}$. Explicitly observable actions shall be denoted \Longrightarrow^{σ} .

Our transition system is non-temporal — it does not take into account any times inside the Demos program. So, for example, it will completely ignore synchronisation by any `hold` statement. Consequently, it will consider the execution of commands on resources in any and every order, including those that are impossible according to the operational semantics of the program, which does take into account time.

We can see, therefore, that for the full Demos system we do not have the following:

$$\text{If } \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket \text{ then } f(\mathbf{R} \triangleleft \llbracket P \rrbracket) \implies f(\mathbf{R}' \triangleleft \llbracket P' \rrbracket)$$

for some (justified — *i.e.*, not something vacuous that makes the above trivially true) function f mapping transition system states to operational semantics states.

We may, however, consider the converse: does our transition system *simulate* the operational semantics? That is, do we have the property:

$$\text{If } S \implies S' \text{ then } g(S) \xrightarrow{\sigma^*} g(S'),$$

where g maps states in the operational semantics to states in the transition system and $\xrightarrow{\sigma^*}$ is the reflexive, transitive closure of $\xrightarrow{\sigma}$ (the effect of running an arbitrary number of steps, including τ -transitions), $\sigma \in Act$.

In order to demonstrate such a relationship, we must define the map g between states in the operational semantics and states in the transition system. Firstly, though, to aid clarity of expression, we shall introduce a new syntactic abbreviation for the parallelisation operator that operates outside the semantic brackets:

$$\begin{aligned} \llbracket P_1 \rrbracket \mid \llbracket P_2 \rrbracket &\equiv \llbracket P_1 \parallel P_2 \rrbracket \\ \llbracket P_1 \rrbracket \mid \llbracket \varepsilon \rrbracket &\equiv \llbracket P_1 \rrbracket \end{aligned}$$

We shall define g through three auxilliary functions.

- g_E shall convert processes in the event list to transition system processes.

$$\begin{aligned} g_E(\llbracket \rrbracket) &\stackrel{\text{def}}{=} \llbracket \varepsilon \rrbracket \\ g_E((P, PD(\llbracket \rrbracket, \llbracket \rrbracket, t)) : \mathcal{EL}) &\stackrel{\text{def}}{=} g_E(\mathcal{EL}) \\ g_E((P, PD(\llbracket \rrbracket, attrs, t)) : \mathcal{EL}) &\stackrel{\text{def}}{=} \llbracket \varepsilon \rrbracket \\ g_E((P, PD(b:body, attrs, t)) : \mathcal{EL}) &\stackrel{\text{def}}{=} \begin{cases} \llbracket \varepsilon \rrbracket, & \text{if } b = \text{close} \\ \llbracket b; body_{[P]} \rrbracket \mid g_E(\mathcal{EL}), & \text{otherwise} \end{cases} \end{aligned}$$

- g_W shall convert lists of waiting (blocked) processes to processes of the transition system.

$$\begin{aligned} g_W(r, \llbracket \rrbracket) &\stackrel{\text{def}}{=} \llbracket \varepsilon \rrbracket \\ g_W((r, (P, PD(body, attr, t))) : \mathcal{WL}) &\stackrel{\text{def}}{=} \llbracket \text{getR}(r, 1); body_{[P]} \rrbracket \mid g_W(\mathcal{WL}) \end{aligned}$$

- To convert the resource list, $\langle \langle ident, avail, waiting \rangle \rangle : \langle \langle R, \mathbb{B}, [Proc] \rangle \rangle$, we shall define $g_R : \langle \langle R, \mathbb{B}, [Proc] \rangle \rangle \rightarrow \langle \mathcal{R}, \llbracket \mathcal{P}_{\text{ROG}} \rrbracket \rangle$; the operator \oplus shall combine these pairs:

$$\langle \mathbf{R}, \llbracket P_{[\ell_1]} \rrbracket \rangle \oplus \langle \mathbf{S}, \llbracket Q_{[\ell_2]} \rrbracket \rangle \stackrel{\text{def}}{=} \langle \mathbf{R} \uplus \mathbf{S}, \llbracket P_{[\ell_1]} \rrbracket \mid \llbracket Q_{[\ell_2]} \rrbracket \rangle$$

$$\begin{aligned}
g_R(\langle \rangle) &\stackrel{\text{def}}{=} \langle \cup, \llbracket \varepsilon \rrbracket \rangle \\
g_R(\langle r, \mathbf{t}, \langle \rangle \rangle : \mathcal{RL}) &\stackrel{\text{def}}{=} \langle \{r\}, \llbracket \varepsilon \rrbracket \rangle \oplus g_R(\mathcal{RL}) \\
g_R(\langle r, \mathbf{ff}, W \rangle : \mathcal{RL}) &\stackrel{\text{def}}{=} \langle \cup, g_W(W) \rangle \oplus g_R(\mathcal{RL})
\end{aligned}$$

We may now define our map between states in the two systems:

$$g(\mathcal{EL}, \mathcal{RL}, \Sigma) = \mathbf{R} \triangleleft g_E(\mathcal{EL}) | W,$$

where $g_R(\mathcal{RL}) = \langle \mathbf{R}, W \rangle$.

Theorem 3.2.1 Let $S = \langle \mathcal{EL}, \mathcal{RL} \rangle$ be a reachable state in a π Demos program which evolves in one step of the operational semantics to $S' = \langle \mathcal{EL}', \mathcal{RL}' \rangle$. Then this is simulated in the transition system. That is, for some series of σ ,

$$\text{If } S \Longrightarrow S' \text{ then } g(S) \xrightarrow{\sigma^*} g(S')$$

PROOF *By analysis of the cases of next command to be executed (the command at the head of the event list).*

We may assume, without loss of generality, that the choice of the transition to make is decided solely by the command at the head of the event list. Generality is not lost as this provides an arbitrary selection from the processes with equal (but still the least) event time, assuming insertion order is arbitrary where equal event times occur.

The event list is empty:

We have $S \equiv \langle \langle \rangle, \mathcal{RL} \rangle$. According to the operational semantics, the program halts with an error. As there is no next state, the statement is vacuously true.

Lead object has no actions left:

We consider two cases:

$attrs \neq \langle \rangle$:

Then the program halts with an error, so, again, the statement is vacuously true.

$attrs = \langle \rangle$:

Then $S \equiv ((P, PD(\langle \rangle, \langle \rangle, t)) : \mathcal{EL}, \mathcal{RL})$. The operational semantics gives $S' \equiv (\mathcal{EL}', \mathcal{RL}')$. By the definition of g_E , $g(S) = g(S')$, and, as $\xrightarrow{\sigma^*}$ is reflexive, we have $g(S) \xrightarrow{\sigma^*} g(S')$ as required.

close:

Then $S \equiv ((P, PD(\text{close} : \text{body}, attrs, t)) : \mathcal{EL}, \mathcal{RL})$. The operational semantics tells us that the program halts and that there is no next state, so the statement is vacuously true.

hold(t):

Then we have $S \equiv ((P, PD(\text{hold}(t) : \text{body}, attrs, evt)) : \mathcal{EL}, \mathcal{RL})$. By applying the rule for **hold**, we obtain $S' \equiv (\mathcal{EL}', \mathcal{RL}')$ where \mathcal{EL}' is the same

as the original event list with t added to evt (and, consequently, causing a re-order of the event list) and $\text{hold}(t)$ removed.

Now, $g(S) = \mathbf{R} \triangleleft \llbracket \text{hold}(t); \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid W$ where $g_R(\mathcal{R}\mathcal{L}) = \langle \mathbf{R}, W \rangle$. Similarly, we have $g(S') = \mathbf{R} \triangleleft g_E(\mathcal{E}\mathcal{L}) \mid \llbracket \text{body} \rrbracket \mid W$, where $g_R(\mathcal{R}\mathcal{L}) = \langle \mathbf{R}, W \rangle$. By the associativity and commutativity of parallelisation, this is just the same as the state derived by executing the hold in the transition system,

$$\mathbf{R} \triangleleft \llbracket \text{hold}(t); \text{body} \rrbracket \xrightarrow{\tau} \mathbf{R} \triangleleft \llbracket \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid W.$$

putR(r,1):

There are two subcases for this instruction:

Nothing is waiting on the resource:

Then $S \equiv ((P, \text{PD}(\text{putR}(r, 1); \text{body}, \text{attrs}, t)) : \mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L})$. After we have executed putR , we have state $S' \equiv ((P, \text{PD}(\text{body}, \text{attrs}, t)) : \mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L}')$. We know that r was previously unavailable, so (without loss of generality thanks to the order of $\mathcal{R}\mathcal{L}$ being arbitrary), $\mathcal{R}\mathcal{L} = (r, \text{ff}, []) : \mathcal{R}\mathcal{L}''$, and $\mathcal{R}\mathcal{L}' = (r, \text{tt}, []) : \mathcal{R}\mathcal{L}''$.

Now, $g(S) = \mathbf{R} \triangleleft \llbracket \text{putR}(r, 1); \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid W$ where $g_R(\mathcal{R}\mathcal{L}) = \langle W, \mathbf{R} \rangle$. With appropriate selection of parallelisation and sequential composition rules, we can derive the transition

$$\mathbf{R} \triangleleft \llbracket \text{putR}(r, 1); \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid W \xrightarrow{\sigma} \mathbf{R}, r \triangleleft \llbracket \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid W;$$

the actual label, $\ell_1 : \text{putR}_r^1$, is unimportant, so shall just call it σ . Consider $g(S') = \mathbf{R}' \triangleleft \llbracket \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid W'$: by the definition of g_R , because no processes were waiting on r , $W = W'$ and $\mathbf{R}' = \mathbf{R}, r$, so we have the the state obtained by the transition system.

There are processes waiting on the resource:

S is as before. After putR , we have $S' \equiv ((P, \text{PD}(\text{body}, \text{attrs}', t)) : \mathcal{E}\mathcal{L}', \mathcal{R}\mathcal{L}')$. Assume, without loss of generality, that $\mathcal{R}\mathcal{L} = (r, \text{ff}, (P_1, \text{PD}(\text{body}_1, \text{attrs}_1, t_1)) : W) : \mathcal{R}\mathcal{L}''$. Then $\mathcal{R}\mathcal{L}' = (r, \text{ff}, W) : \mathcal{R}\mathcal{L}''$. The operational semantics tells us that $\mathcal{E}\mathcal{L}' = \mathcal{E}\mathcal{L}$ with $(P_1, \text{PD}(\text{body}_1, \text{attrs}_1, t_1))$ inserted at some point.

Now, $g(S) = \mathbf{R} \triangleleft \llbracket \text{putR}(r, 1); \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid W$, where $g_R(\mathcal{R}\mathcal{L}) = \langle \mathbf{R}, W \rangle$. By consideration of the definition of g_R , we see that $W = \llbracket \text{getR}(r, 1); \text{body}_1 \rrbracket \mid W'$, and so we have $g(S) = \mathbf{R} \triangleleft \llbracket \text{putR}(r, 1); \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid \llbracket \text{getR}(r, 1); \text{body}_1 \rrbracket \mid W'$.

$$\begin{aligned} g(S) &\xrightarrow{\sigma} \mathbf{R}, r \triangleleft \llbracket \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid \llbracket \text{getR}(r, 1); \text{body}_1 \rrbracket \mid W' \\ &\xrightarrow{\sigma} \mathbf{R} \triangleleft \llbracket \text{body} \rrbracket \mid g_E(\mathcal{E}\mathcal{L}) \mid \llbracket \text{body}_1 \rrbracket \mid W' \end{aligned}$$

which is exactly (up to the associativity and commutativity of parallel process composition) $g(S')$.

getR(r,1) — resource available:

We have $S \equiv ((P, \text{PD}(\text{getR}(r, 1); \text{body}, \text{attrs}, t)) : \mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L})$. The resource is available, so $S' \equiv ((P, \text{PD}(\text{body}, \text{attrs}', t)) : \mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L}')$. Without loss of generality (thanks to the arbitrary order of $\mathcal{R}\mathcal{L}$ — it is a multiset rather

than a list), let $\mathcal{RL} = (r, \mathbf{t}, []) : \mathcal{RL}''$, and so $\mathcal{RL}' = (r, \mathbf{ff}, []) : \mathcal{RL}''$ by the operational semantics.

Now, $g(S) = \mathbf{R} \triangleleft \llbracket \text{body} \rrbracket \mid g_E(\mathcal{EL}) \mid W$, where $g_R(\mathcal{RL}) = \langle \mathbf{R}, W \rangle$. Considering $g_R(\mathcal{RL})$, we notice that $\mathbf{R} = \mathbf{R}'', r$, and so $g(S) = \mathbf{R}'', r \triangleleft \llbracket \text{body} \rrbracket \mid g_E(\mathcal{EL}) \mid W$. Finally,

$$\mathbf{R}'', r \triangleleft \llbracket \text{getR}(r, 1); \text{body} \rrbracket \mid g_E(\mathcal{EL}) \mid W \xrightarrow{\sigma} \mathbf{R}'' \triangleleft \llbracket \text{body} \rrbracket \mid g_E(\mathcal{EL}) \mid W,$$

which is exactly $g(S')$.

getR(r,1) — the resource is unavailable:

We have $S \equiv ((P, \text{PD}(\text{getR}(r, 1); \text{body}, \text{attrs}, t)) : \mathcal{EL}, \mathcal{RL})$. After executing $\text{getR}(r, 1)$, we have $S' \equiv \langle \mathcal{EL}, \mathcal{RL}' \rangle$. Without loss of generality, let $\mathcal{RL} = (r, \mathbf{ff}, W) : \mathcal{RL}''$; hence, $\mathcal{RL}' = (r, \mathbf{ff}, W') : \mathcal{RL}''$ where W' is W with $(P, \text{PD}(\text{body}, \text{attrs}, t))$ inserted at the appropriate point.

$g(S) = \mathbf{R} \triangleleft \llbracket \text{getR}(r, 1); \text{body} \rrbracket \mid g_E(\mathcal{EL}) \mid V$, where $g_R(\mathcal{RL}) = \langle \mathbf{R}, V \rangle$. Notice that, by the definition of g_R , $g_R(\mathcal{RL}') = \langle \mathbf{R}, V \mid \llbracket \text{getR}(r, 1); \text{body} \rrbracket \rangle$. Consequently, $g(S') = \mathbf{R} \triangleleft g_E(\mathcal{EL}) \mid V \mid \llbracket \text{getR}(r, 1); \text{body} \rrbracket$, which, thanks to the associativity and commutativity of the parallel composition operator, is $g(S)$. Now, as $\xrightarrow{\sigma}^*$ by definition is reflexive, we have $g(S) \xrightarrow{\sigma}^* g(S)$, or, equivalently, $g(S) \xrightarrow{\sigma}^* g(S')$, as required. \square

As we stated earlier, we must also show that g faithfully translates states belonging to the operational semantics. A number of preliminaries must be covered first, so this is presented in §3.7.1.

3.3 Soundness of transitions

As we have already stated, we do not have the obvious soundness property stating that any transition derived by our transition system could have been derived according to the operational semantics.

The transition system was designed to take no heed of synchronisation by time, nor does it take into account the priority of processes. We should, though, be able to prove soundness with respect to the operational semantics with the following assumptions:

1. All event times are ignored; *i.e.*, $t = 0$, or some other arbitrary constant, for all PD entries.
2. All priorities are ignored.
3. The orders of the event list and waiting lists, consequently, are arbitrary.
4. When a resource is released, it is not necessarily a waiting process that acquires the resource — it could be either a process in a waiting list or a process about to enter a waiting list (*i.e.*, with getR at its head).

In order to prove the property, we firstly need to define the map from the transition system state to state in the operational semantics, \bar{f} . We use a number of auxiliary functions:

- h_P defines the resources held by a non-parallel process (note here that we assume that the programs passed are *well-formed*; that is, they only put resources that they own *etc.*). ‘\’ denotes subtraction from a multiset.

$$\begin{aligned} h_P(\llbracket \varepsilon \rrbracket) &\stackrel{\text{def}}{=} \emptyset \\ h_P(\llbracket \text{getR}(r, 1); P \rrbracket) &\stackrel{\text{def}}{=} h_P(\llbracket P \rrbracket) \setminus \{r\} \\ h_P(\llbracket \text{putR}(r, 1); P \rrbracket) &\stackrel{\text{def}}{=} h_P(\llbracket P \rrbracket) \uplus \{r\} \end{aligned}$$

- h shall be the multiset of all held resources:

$$\begin{aligned} h(\llbracket P_1 \parallel P_2 \rrbracket) &\stackrel{\text{def}}{=} h(\llbracket P_1 \rrbracket) \uplus h(\llbracket P_2 \rrbracket) \\ h(\llbracket P \rrbracket) &\stackrel{\text{def}}{=} h_P(\llbracket P \rrbracket), P \text{ not parallel} \end{aligned}$$

- w_R defines the multiset of $\langle \text{process}, \text{resource} \rangle$ pairs where the process is waiting on the resource

$$\begin{aligned} w_R(\llbracket \varepsilon \rrbracket, \mathbf{R}) &\stackrel{\text{def}}{=} \emptyset \\ w_R(\llbracket P \parallel Q \rrbracket, \mathbf{R}) &\stackrel{\text{def}}{=} w_R(P) \uplus w_R(Q) \\ w_R(\llbracket c; P \rrbracket, \mathbf{R}) &\stackrel{\text{def}}{=} \begin{cases} \{ \langle \llbracket P \rrbracket, r \rangle \}, & \text{if } c = \text{getR}(r, 1) \text{ and } r \notin \mathbf{R} \\ \emptyset, & \text{otherwise} \end{cases} \end{aligned}$$

- w is just the multiset of waiting processes (without the resource that they are waiting on):

$$w(\llbracket P \rrbracket, \mathbf{R}) \stackrel{\text{def}}{=} \{ \llbracket P \rrbracket \text{ s.t. } \langle \cdot, \llbracket P \rrbracket \rangle \in w_R(\llbracket P \rrbracket, \mathbf{R}) \}$$

- $mset$ simply splits parallel processes into a multiset of non-parallel processes:

$$\begin{aligned} mset(\llbracket \varepsilon \rrbracket) &\stackrel{\text{def}}{=} \emptyset \\ mset(\llbracket P \rrbracket) &\stackrel{\text{def}}{=} \{ \llbracket P \rrbracket \}, \text{ where } \llbracket P \rrbracket \text{ not parallel} \\ mset(\llbracket P_1 \parallel P_2 \rrbracket) &\stackrel{\text{def}}{=} mset(P_1) \uplus mset(P_2) \end{aligned}$$

- The multiset of active processes is just the multiset of all processes minus the multiset of waiting processes.

$$a(\llbracket P \rrbracket, \mathbf{R}) \stackrel{\text{def}}{=} mset(\llbracket P \rrbracket) \setminus w(\llbracket P \rrbracket, \mathbf{R})$$

- The set of processes waiting on a resource r is given by:

$$w_P(r) \stackrel{\text{def}}{=} \{ \llbracket P \rrbracket \mid \langle \llbracket P \rrbracket, r \rangle \in w_R(\llbracket P \rrbracket, \mathbf{R}) \}$$

We are now in a position to define our translation \bar{f} :

$$\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket) \stackrel{\text{def}}{=} \langle \mathcal{EL}, \mathcal{RL} \rangle$$

where

$$\begin{aligned} \mathcal{EL} &\stackrel{\text{def}}{=} \left[(\ell, \text{PD}(\text{body}, \text{attrs}, t)) \mid \begin{array}{l} \llbracket \text{body}_{[\ell]} \rrbracket \in a(\llbracket P \rrbracket, \mathbf{R}) \wedge \\ \text{attrs} = h_P(\llbracket \text{body} \rrbracket) \wedge t = 0 \end{array} \right] \\ \mathcal{RL} &\stackrel{\text{def}}{=} \{ \langle r, \mathbf{t}, [] \rangle \text{ s.t. } r \in \mathbf{R} \} \uplus \{ \langle r, \mathbf{ff}, \mathcal{W}_r \rangle \text{ s.t. } r \in h(\llbracket P \rrbracket) \} \\ \mathcal{W}_r &= [(\ell, \text{PD}(\llbracket \text{body} \rrbracket, h_P(\llbracket \text{body} \rrbracket), 0)) \text{ s.t. } \llbracket \text{body}_{[\ell]} \rrbracket \in w_P(r)] \end{aligned}$$

We shall now proceed with the body of the proof, which shall be by induction on the structure of σ Demos processes.

Theorem 3.3.1 The transition system is sound with respect to a non-time and non-priority sensitive operational semantics. That is,

$$\text{If } \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket \text{ then } \bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket) \Longrightarrow^* \bar{f}(\mathbf{R}' \triangleleft \llbracket P' \rrbracket).$$

PROOF *By induction on the structure of σ Demos processes P .*

$P \equiv \text{hold}(t)$:

TS	We may only derive the transition $\mathbf{R} \triangleleft \llbracket \text{hold}(t) \rrbracket \xrightarrow{\tau} \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket$
Op sem	$\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket) = \langle [P, \text{PD}(\llbracket \text{hold}(t) \rrbracket, \text{attrs}, 0)], \mathcal{RL} \rangle$. Executing the hold yields $\mathcal{EL}' = \langle [P, \text{PD}([], \text{attrs}, 0)] \rangle$. Given the well-formedness of states, $\text{attrs} = []$, so another execution gives $\mathcal{EL}'' = [], \mathcal{RL}'' = \mathcal{RL}$. Notice that $\bar{f}(\mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket) = \langle [], \mathcal{RL} \rangle = \langle \mathcal{EL}'', \mathcal{RL}'' \rangle$.

$P \equiv \text{getR}(r, 1)$:

TS	Assuming that we are able to make a transition, the resource must be available, so $\mathbf{R} \equiv \mathbf{R}'', r$, for some \mathbf{R}'' . We have the transition $\mathbf{R}'', r \triangleleft \llbracket \text{getR}(r, 1) \rrbracket \xrightarrow{\sigma} \mathbf{R}'' \triangleleft \llbracket \varepsilon \rrbracket$.
Op sem	Without loss of generality thanks to the arbitrary order of the resource list, because r is available we may assume that $\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket) = \langle [(P, \text{PD}(\llbracket \text{getR}(r, 1) \rrbracket, \text{attrs}, 0)), [(r, \mathbf{t}, [])] : \mathcal{RL}'' \rangle$. The operational semantics allows us to take the resource, yielding $\langle [(P, \text{PD}([], [], 0)), [(r, \mathbf{ff}, [])] : \mathcal{RL}'' \rangle$. The ‘exec’ command allows us to take another step to $\langle [], [(r, \mathbf{ff}, [])] : \mathcal{RL}'' \rangle$. Assuming the well-formedness of the program (this atomic command in itself is not well-formed — we could, alternatively, have made this a special case in the COMP case below), this is $\bar{f}(\mathbf{R}'' \triangleleft \llbracket P'' \rrbracket)$.

$P \equiv \text{putR}(r, 1)$:

Similar — but we use assumption (4) to eliminate the possibility of a waiting process taking the resource straight away.

$P \equiv c; \text{body}$:

c is a command and body some non-parallel process. Consider the cases for c .

$c \equiv \text{hold}(t)$:

TS	$P \equiv \text{hold}(t); \text{body}.$ We may only derive the transition $\mathbf{R} \triangleleft \llbracket \text{hold}(t); \text{body} \rrbracket \xrightarrow{\tau} \mathbf{R} \triangleleft \llbracket \text{body} \rrbracket.$
Op sem	$\bar{f}(\mathbf{R} \triangleleft \llbracket \text{hold}(t); \text{body} \rrbracket) = \langle \langle (P, \text{PD}(\llbracket \text{hold}(t); \text{body} \rrbracket, \text{attrs}, 0)) \rangle, \mathbf{R} \rangle.$ The operational semantics yields, upon execution of the hold, $\langle \langle (P, \text{PD}(\llbracket \text{body} \rrbracket, \text{attrs}, 0)) \rangle, \mathbf{R} \rangle$, which is exactly $\bar{f}(\mathbf{R} \triangleleft \llbracket \text{body} \rrbracket)$

$c \equiv \text{getR}(r, 1)$ or $\text{putR}(r, 1)$:

As before, but we will not need to use the ‘exec’ rule to reduce down to an empty event list. (The consideration of getR above could be placed here instead of above — it was placed there to show how a proof that considered bins would be constructed)

$P \equiv P_1 \parallel P_2$:

Assume, without loss of generality, that $\mathbf{R} \triangleleft \llbracket P_1 \parallel P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_1' \parallel Q \rrbracket$. We therefore have $\mathbf{R} \triangleleft \llbracket P_1 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_1' \rrbracket$, so, by induction, $\bar{f}(\mathbf{R} \triangleleft \llbracket P_1 \rrbracket) \Longrightarrow^* \bar{f}(\mathbf{R}' \triangleleft \llbracket P_1' \rrbracket)$. We must show that this execution matches the execution of $\bar{f}(\mathbf{R} \triangleleft \llbracket P_1 \parallel P_2 \rrbracket)$. P_2 is either waiting on a resource or it is not: if it is, it will be placed in the waiting list of some resource. It, therefore, will not affect the execution of this command. If P_2 is placed in the event list, thanks to the arbitrary order of the event list, it is possible that P_1 ’s execution may be followed again. Moreover, this property is maintained through the execution of any intermediate steps made by P due to all times being equal. □

3.4 Weak bisimulation

To consider some further properties, it will be useful to have a notion of equivalence between states; we shall adopt *observational equivalence* similar to that developed in [Mil89]. Observational equivalence between states, informally, tells us that any *visible* (in our case, ignoring τ -transitions) series of transitions from one of the states can be made if, and only if, it can be made from the other.

More formally, writing Act for the set of all actions and Act^* for the set of observable actions (*i.e.*, $Act^* = Act \setminus \{\tau\}$):

Definition (Weak bisimulation)

Let \star represent a relation between states in the transition system, *i.e.*, $\star \subseteq \Omega^2$. We shall use infix notation, writing $\mathbf{R} \triangleleft \llbracket P \rrbracket \star \mathbf{R}' \triangleleft \llbracket Q \rrbracket$ ¹ rather than $(\mathbf{R} \triangleleft \llbracket P \rrbracket, \mathbf{R}' \triangleleft \llbracket Q \rrbracket) \in \star$. \star is a *weak bisimulation* between states if:

1. for any transition $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket$ where $\sigma \in Act^*$, we have $\mathbf{R} \triangleleft \llbracket Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket Q' \rrbracket$ and $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \star \mathbf{R}' \triangleleft \llbracket Q' \rrbracket$; and
2. for any transition $\mathbf{R} \triangleleft \llbracket Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket Q' \rrbracket$ where $\sigma \in Act^*$, we have $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket$ and $\mathbf{R}' \triangleleft \llbracket Q' \rrbracket \star \mathbf{R}' \triangleleft \llbracket P' \rrbracket$.

¹We shall not allow bisimulation between states with different resource sets at this stage, so if $\mathbf{R} \neq \mathbf{S}$, $\mathbf{R} \triangleleft \llbracket P \rrbracket \not\star \mathbf{S} \triangleleft \llbracket Q \rrbracket$ for any P, Q .

Definition (Observational equivalence relation)

We shall use \approx to denote observable equivalence. We write $\mathbf{R} \triangleleft \llbracket P \rrbracket \approx \mathbf{R} \triangleleft \llbracket Q \rrbracket$ if there is a weak bisimulation between these states. That is,

$$\approx \stackrel{\text{def}}{=} \bigcup \{ \star \mid \star \text{ is a bisimulation} \}$$

It may be that \approx is a *congruence relation*. That is, for any process $\Psi[P]$ containing P as a sub-term, if $P \approx Q$ then $\Psi[P] \approx \Psi[Q]$ (where $\Psi[Q]$ is $\Psi[P]$ with P substituted for Q). We do not need to prove this property, but do require that if two states are observationally equivalent and we create the parallel composition of each of them with a third state, observational equivalence is preserved.

Lemma 3.4.1 If $\mathbf{R} \triangleleft \llbracket P \rrbracket \approx \mathbf{R} \triangleleft \llbracket P' \rrbracket$, then $\mathbf{R} \triangleleft \llbracket P \parallel Q \rrbracket \approx \mathbf{R} \triangleleft \llbracket P' \parallel Q \rrbracket$ for any process Q .

PROOF *Very similar to CCS, [Mil89] pp. 98*

We shall demonstrate that $\star = \{ (\mathbf{R} \triangleleft \llbracket P_1 \parallel Q \rrbracket, \mathbf{R} \triangleleft \llbracket P_2 \parallel Q \rrbracket) \mid \mathbf{R} \triangleleft \llbracket P_1 \rrbracket \approx \mathbf{R} \triangleleft \llbracket P_2 \rrbracket \}$ is a bisimulation.

To prove this we can show that, for $\sigma \in \text{Act}^*$ and supposing $\mathbf{R} \triangleleft \llbracket P_1 \parallel Q \rrbracket \star \mathbf{R} \triangleleft \llbracket P_2 \parallel Q \rrbracket$:

1. (a) $\mathbf{R} \triangleleft \llbracket P_1 \parallel Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_1 \parallel Q \rrbracket$ implies $\mathbf{R} \triangleleft \llbracket P_2 \parallel Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_2 \parallel Q \rrbracket$ and $\mathbf{R}' \triangleleft \llbracket P'_1 \parallel Q \rrbracket \star \mathbf{R}' \triangleleft \llbracket P'_2 \parallel Q \rrbracket$.
- (b) $\mathbf{R} \triangleleft \llbracket P_1 \parallel Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_1 \parallel Q' \rrbracket$ implies $\mathbf{R} \triangleleft \llbracket P_2 \parallel Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_2 \parallel Q' \rrbracket$ and $\mathbf{R}' \triangleleft \llbracket P_1 \parallel Q' \rrbracket \star \mathbf{R}' \triangleleft \llbracket P_2 \parallel Q' \rrbracket$.
2. The same, swapping P_2 and P_1 . This shall follow immediately from the symmetry of the parallelisation operator.
1. (a) $\mathbf{R} \triangleleft \llbracket P_1 \parallel Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_1 \parallel Q \rrbracket$, so must be able to derive the transition $\mathbf{R} \triangleleft \llbracket P_1 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_1 \rrbracket$. As $\mathbf{R} \triangleleft \llbracket P_1 \rrbracket \approx \mathbf{R} \triangleleft \llbracket P_2 \rrbracket$, we can also derive $\mathbf{R} \triangleleft \llbracket P_2 \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_2 \rrbracket$ with $\mathbf{R}' \triangleleft \llbracket P'_1 \rrbracket \approx \mathbf{R}' \triangleleft \llbracket P'_2 \rrbracket$. We may then derive $\mathbf{R} \triangleleft \llbracket P_2 \parallel Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P'_2 \parallel Q \rrbracket$, and $\mathbf{R}' \triangleleft \llbracket P'_1 \parallel Q \rrbracket \star \mathbf{R}' \triangleleft \llbracket P'_2 \parallel Q \rrbracket$.
- (b) $\mathbf{R} \triangleleft \llbracket P_1 \parallel Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_1 \parallel Q' \rrbracket$, so $\mathbf{R} \triangleleft \llbracket Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket Q' \rrbracket$. Hence, $\mathbf{R} \triangleleft \llbracket P_2 \parallel Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P_2 \parallel Q' \rrbracket$, and $\mathbf{R}' \triangleleft \llbracket P_1 \parallel Q' \rrbracket \star \mathbf{R}' \triangleleft \llbracket P_2 \parallel Q' \rrbracket$. □

Another useful property of observational equivalence is that it is an equivalence relation: it is *reflexive*, *symmetric* and *transitive*.

Theorem 3.4.2 Observational equivalence is an equivalence relation.

PROOF We split the proof into proving the three sub-properties. That is, for all $\mathbf{R}_1 \triangleleft \llbracket P_1 \rrbracket$, $\mathbf{R}_2 \triangleleft \llbracket P_2 \rrbracket$ and $\mathbf{R}_3 \triangleleft \llbracket P_3 \rrbracket$,

1. Reflexivity: $\mathbf{R}_1 \triangleleft \llbracket P_1 \rrbracket \approx \mathbf{R}_1 \triangleleft \llbracket P_1 \rrbracket$.
2. Symmetry: $\mathbf{R}_1 \triangleleft \llbracket P_1 \rrbracket \approx \mathbf{R}_2 \triangleleft \llbracket P_2 \rrbracket$ implies $\mathbf{R}_2 \triangleleft \llbracket P_2 \rrbracket \approx \mathbf{R}_1 \triangleleft \llbracket P_1 \rrbracket$.

3. Transitivity: $\mathbf{R}_1 \triangleleft \llbracket P_1 \rrbracket \approx \mathbf{R}_2 \triangleleft \llbracket P_2 \rrbracket$ and $\mathbf{R}_2 \triangleleft \llbracket P_2 \rrbracket \approx \mathbf{R}_3 \triangleleft \llbracket P_3 \rrbracket$ implies $\mathbf{R}_1 \triangleleft \llbracket P_1 \rrbracket \approx \mathbf{R}_3 \triangleleft \llbracket P_3 \rrbracket$.

The proofs of (1) and (2) are immediate from the definitions of \approx and bisimulation. (3) follows easily, too. \square

3.5 Kernel of π Demos states

Let Π be the set of states in the operational semantics of π Demos, ranged over by tuples $D = \langle \mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L}, \Sigma \rangle$. The equivalence class of a state D with respect to g ,

$$\|D\| \stackrel{\text{def}}{=} \{E \in \Pi \mid g(D) \approx g(E)\},$$

is the set of states (containing programs) that are “equal”, in some sense relating to the execution of the states, to the state under consideration. In our case this shall be observational equivalence in the transition system.

Suppose we want to find a member of the equivalence class for a particular state, D . We cannot simply construct $\|D\|$ by checking every Demos state for equivalence: the set of states Π is infinite. Instead, we must obtain such a program (which we will say belongs to the *kernel* of Π , Π_K) from its syntax.

Definition The π Demos state $D_k \in \Pi_K$ shall be the π Demos state $D = (\mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L})$ with all the `hold(t)` commands removed and all times (both in the event list and the resource lists) set to zero. (If priorities existed in π Demos — and they do not — we would equalise these too.) Where the first command of a process in the event list is `getR(r , 1)` and r is unavailable in $\mathcal{R}\mathcal{L}$, the remaining body of code of the process is placed (with the other kernelised process information) in the waiting list of r . The resource list of the kernel state has no extraneous resource entries; that is, there are no entries $(r, \mathbf{ff}, [])$ where r is not referred to by any process. We shall assume that the operational semantics are extended to maintain this property.

The following lemma asserts that the the corresponding kernel state for a non-parallel process belongs to the equivalence class. We use the notation $\mathcal{R}\mathcal{L}[(r, b, W)]$ to mean the resource list $\mathcal{R}\mathcal{L}$ with the entry for r being (r, b, W) . We use $\mathcal{E}\mathcal{L}[P/Q]$ to mean the event list $\mathcal{E}\mathcal{L}$ having P in place of Q (thus possibly causing the order to be different if the times of P and Q are different). The reader should be cautious as we use the subscript k to indicate kernelisation following other subscripts.

Lemma 3.5.1 For all Π states $(\mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L})$, if $g(\mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L}) = \mathbf{R} \triangleleft \llbracket P \rrbracket$ for some $\mathbf{R} \triangleleft \llbracket P \rrbracket$ where P is non-parallel, then $g(\mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L}) = \mathbf{R} \triangleleft \llbracket P \rrbracket \approx \mathbf{R}' \triangleleft \llbracket P_k \rrbracket = g((\mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L})_k)$

PROOF By the definition of g , we have $g(\mathcal{E}\mathcal{L}, \mathcal{R}\mathcal{L}) = \mathbf{R} \triangleleft g_E(\mathcal{E}\mathcal{L}) \mid W$ where $g_R(\mathcal{R}\mathcal{L}) = \langle \mathbf{R}, W \rangle$. Since the result is not the parallel composition of processes, there is only one body of code in $(\mathcal{R}\mathcal{L}, \mathcal{E}\mathcal{L})$. Irrespective of where it came from, we may induce over the body of code (*i.e.*, $P \in \Pi$) found in some process descriptor.

No code in body:

Notice that this would imply that the process came from the event list; otherwise, if it had been a process waiting on a resource, we would have `getR` at the head of the body.

Kernelisation yields exactly the same process (albeit possibly with a different event time). Consequently, $g(D) = g(D_k)$. As \approx is reflexive, $g(D) \approx g(D_k)$

$P \equiv \text{putR}(r, 1)$ or `close`:

Again,² $D = D_k$ implies that $g(D) = g(D_k)$, so $g(D) \approx g(D_k)$.

$P \equiv \text{getR}(r, 1)$:

There are two cases to consider: either $(r, \mathbf{t}, [])$ is in \mathcal{RL} or (r, \mathbf{ff}, W) is in \mathcal{RL} .

If $(r, \mathbf{t}, [])$ is in \mathcal{RL} , we have $D = D_k$ which implies $g(D) = g(D_k)$, so $g(D) \approx g(D_k)$.

Otherwise, r is not in \mathbf{R} , the available resource collection of $g(D)$ or in \mathbf{R}' , the available resource collection of $g(D_k)$. Consequently, irrespective of whether the process descriptor was in \mathcal{EL} or waiting on r in \mathcal{RL} , we can derive no transition from $g(D)$, nor can we from $g(D_k)$. Thus $g(D) \approx g(D_k)$.

$P \equiv \text{hold}(t)$:

Then $g(D) = \mathbf{R} \triangleleft \llbracket \text{hold}(t) \rrbracket$; the corresponding kernel state, D_k , gives $g(D_k) = \mathbf{R} \triangleleft \llbracket \varepsilon \rrbracket$. Neither of these processes may make any observable transition — $\llbracket \text{hold}(t) \rrbracket$ may only make a silent transition.

$P \equiv P_1; P_2$:

Without loss of generality, we may assume that $P \equiv a; \text{body}$, where a is a single command. Consider the cases for a :

$a \equiv \text{getR}(r, 1)$:

Then $P \equiv \text{getR}(r, 1); \text{body}$, and $g(D) = \mathbf{R} \triangleleft \llbracket \text{getR}(r, 1); \text{body} \rrbracket$. There are two cases to consider: either $r \in \mathbf{R}$ iff $\langle r, \mathbf{t}, [] \rangle \in \mathcal{RL}$ or $r \notin \mathbf{R}$ iff $\langle r, \mathbf{ff}, W \rangle \in \mathcal{RL}$.

- ($r \notin \mathbf{R}$): We can derive no transition from $\mathbf{R} \triangleleft \llbracket \text{getR}(r, 1); \text{body} \rrbracket$. In the kernel, $g(D_k) = \mathbf{R} \triangleleft \llbracket \text{getR}(r, 1); \text{body}_k \rrbracket$, irrespective of whether body was in the waiting list of r or $\text{getR}(r, 1); \text{body}$ was in the event list, which we can derive no transition from.

- ($r \in \mathbf{R}$): Then $\mathbf{R} \equiv \mathbf{R}'$, r , and we may only derive the transition $\mathbf{R}', r \triangleleft \llbracket \text{getR}(r, 1); \text{body} \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \text{body} \rrbracket$. Kernelisation yields D_k where the process is in the event list (because r is available), so $g(D_k) = \mathbf{R}', r \triangleleft \llbracket \text{getR}(r, 1); \text{body}_k \rrbracket$, from which we may only derive the transition $\mathbf{R}', r \triangleleft \llbracket \text{getR}(r, 1); \text{body}_k \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket \text{body}_k \rrbracket$.

Now, let D' be the π Demos state $D[\text{body}/\text{getR}(r, 1); \text{body}][\langle r, \mathbf{ff}, [] \rangle / \langle r, \mathbf{t}, [] \rangle]$.

Then $g(D') = \mathbf{R}' \triangleleft \llbracket \text{body} \rrbracket$, and, by induction, $g(D'_k) = \mathbf{R}' \triangleleft \llbracket \text{body}_k \rrbracket \approx \mathbf{R}' \triangleleft \llbracket \text{body} \rrbracket$. Consequently, by the definition of \approx , $g(D) = \mathbf{R}', r \triangleleft \llbracket \text{getR}(r, 1); \text{body} \rrbracket \approx \mathbf{R}', r \triangleleft \llbracket \text{getR}(r, 1); \text{body}_k \rrbracket = g(D_k)$.

²We shall write ‘=’ in the sense of two states in Π being exactly the same other than having possibly different event times. Their translations into the transition system will thus be the same.

$a \equiv \text{putR}(r, 1)$:

Similar to $(r \in \mathbf{R})$ case.

$a \equiv \text{close}$:

Similar to $(r \notin \mathbf{R})$ case.

$a \equiv \text{hold}(t)$:

Then $P \equiv \text{hold}(t); \text{body}$ and $g(D) = \mathbf{R} \triangleleft \llbracket \text{hold}(t); \text{body} \rrbracket$. Kernelisation yields $g(D_k) = \mathbf{R} \triangleleft \llbracket \text{body}_k \rrbracket$. By induction, $g(D[\text{body}/\text{hold}(t); \text{body}]) = \mathbf{R} \triangleleft \llbracket \text{body} \rrbracket \approx \mathbf{R} \triangleleft \llbracket \text{body}_k \rrbracket = g(D_k)$. Notice that the hold generates a non-observable τ -transition, so $g(D) \approx g(D_k)$ by the definition of \approx . □

We now generalise, showing that the translation of any πDemos state into the kernel generates an element of the equivalence class.

Theorem 3.5.2 Translation of a πDemos program into the kernel generates a member of the equivalence class. That is, $D_k \in \llbracket D \rrbracket$; or $g(D) \approx g(D_k)$.

PROOF *By induction on $n = |\mathcal{EL}| + |\mathcal{WL}|$ ³ where $|\mathcal{WL}|$ = the number of processes waiting on resources.*

We shall say that two πDemos states are equal even if they have differing event times or the order of lists differs if that order is arbitrary as their interpretation in the transition system will be the same. For convenience, we shall have two base cases.

$n = 0$:

There are no waiting or active processes. Consequently, $D \equiv D_k$ so $g(D) = g(D_k)$. \approx is an equivalence relation and therefore reflexive, so $g(D) \approx g(D_k)$.

$n = 1$:

Immediate by Lemma 3.5.1.

$n > 1$:

The event and resource list translated into σDemos states will be of the form $g(D) = \mathbf{R} \triangleleft \llbracket P_1 \parallel P_2 \rrbracket$.

Let $D_1 = D[P_1/(P_1 \parallel P_2)]$ and $D_2 = D[P_2/(P_1 \parallel P_2)]$ — essentially, deleting from the event and waiting lists a non-zero number of processes. Then $g(D_1) = \mathbf{R} \triangleleft \llbracket P_1 \rrbracket$ and $g(D_2) = \mathbf{R} \triangleleft \llbracket P_2 \rrbracket$. By induction, $\mathbf{R} \triangleleft \llbracket P_1 \rrbracket \approx \mathbf{R} \triangleleft \llbracket P_{1k} \rrbracket = g(D_{1k})$ and $\mathbf{R} \triangleleft \llbracket P_2 \rrbracket \approx \mathbf{R} \triangleleft \llbracket P_{2k} \rrbracket = g(D_{2k})$.

Now, by Lemma 3.4.1, $\mathbf{R} \triangleleft \llbracket P_1 \parallel P_2 \rrbracket \approx \mathbf{R} \triangleleft \llbracket P_{1k} \parallel P_2 \rrbracket \approx \mathbf{R} \triangleleft \llbracket P_{1k} \parallel P_{2k} \rrbracket = g(D_k)$. Therefore, because \approx is an equivalence relation, $g(D) \approx g(D_k)$ □

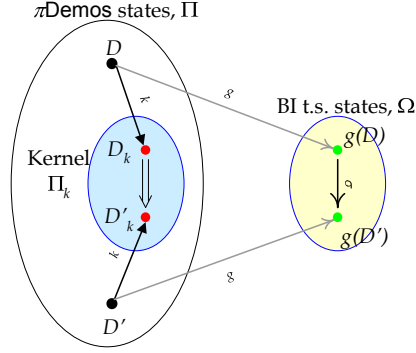


Figure 3.1: The correspondence between states of the transition system and the of operational semantics.

3.6 Correspondence between Π_k and σ Demos states

As we used Σ in the previous section to mean the signature of a state in π Demos, to avoid confusion we shall use the upper-case Greek letter Ω to denote the set of σ Demos (transition system) states. So,

$$g : \Pi \rightarrow \Omega$$

and

$$\bar{f} : \Omega \rightarrow \Pi.$$

The range of g when applied to states in the kernel, Π_k , shall be written Ω_k . That is,

$$\Omega_k \stackrel{\text{def}}{=} \{\mathbf{R} \triangleleft \llbracket P \rrbracket \mid \exists D_k \in \Pi_k. g(D_k) = \mathbf{R} \triangleleft \llbracket P \rrbracket\}$$

It is interesting to consider how our translations between Π and σ Demos states work when we restrict our attention to the kernel sets Π_k and Ω_k . Looking at the definitions of g and \bar{f} , they do not seem to ‘lose’ any information about the state. For some π Demos state in the kernel, $D_k \in \Pi_k$, given its representation in σ Demos, $g(D_k) \in \Omega_k$, we would expect to be able to determine D_k . Similarly, given $\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket)$, we would expect to be able to determine $\mathbf{R} \triangleleft \llbracket P \rrbracket$. We ask, then, the question: does \bar{f} ‘undo’ g , and g ‘undo’ \bar{f} ?

Theorem 3.6.1 For any state $D_k \in \Pi_k$,

$$\bar{f}(g(D_k)) = D_k$$

PROOF OUTLINE *By induction on $|\mathcal{EL}| + |\mathcal{WL}|$ where \mathcal{WL} is the list of processes waiting on resources.*

³ $|\cdot|$ may be read as “list length”. We shall refrain from giving the normal recursive definition.

First consider the resource list, and note that the collection of resource labels is preserved, as is the availability of each resource. Also notice that the each non-parallel (sub)process label is preserved.

The interesting case is where $|\mathcal{EL}| + |\mathcal{WL}| = 1$ and the only process is in the waiting list of some unavailable resource r , \mathcal{WL}_r . Applying g gives a process with $\text{getR}(r, 1)$ at its head, and application of \bar{f} removes this, putting the process back into the waiting list of r .

Notice that we do *not* consider the case where $\text{getR}(r, 1)$ is the command at the head of the event list and r is unavailable, as this is not a kernel state. Therefore, we do not have a problem deciding from $g(D_k)$ whether a process has or has not yet ‘executed’ the $\text{getR}(r, 1)$ command. \square

Theorem 3.6.2 For any state $\mathbf{R} \triangleleft \llbracket P \rrbracket \in \Omega$ — note that we do *not* require that $\mathbf{R} \triangleleft \llbracket P \rrbracket$ be in Ω_k —

$$g(\bar{f}(\mathbf{R} \triangleleft \llbracket P_{[\ell]} \rrbracket)) = \mathbf{R} \triangleleft \llbracket P_{[\ell]} \rrbracket$$

PROOF OUTLINE *By induction on the structure of $\mathbf{R} \triangleleft \llbracket P \rrbracket$.* We first note that the action of g after \bar{f} on an arbitrary state $\mathbf{R} \triangleleft \llbracket P \rrbracket$ results in the same resource collection, \mathbf{R} . We also note that the labels of non-parallel processes are preserved.

Again, there is essentially one interesting case: If $P \equiv \text{getR}(r, 1); \text{body}$ and $r \notin \mathbf{R}$, body will be placed in the waiting list of r in $\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket)$. g will restore the $\text{getR}(r, 1)$ command. Otherwise, for a non-parallel process, the code will be placed in the event list of $\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket)$. g will not add anything to the front, thus restoring $\mathbf{R} \triangleleft \llbracket P \rrbracket$. \square

A *bijection* between two sets A and B is a function that maps every element of A to exactly one element of B , with every element of B being mapped onto exactly once.

Corollary 3.6.3 \bar{f} is the *inverse function* of g (or, equivalently, g is the inverse of \bar{f}) when we consider the sets of kernel states. That is, for $D_k \in \Pi_k$ and $(\mathbf{R} \triangleleft \llbracket P \rrbracket)_k \in \Omega_k$,

$$g(\bar{f}((\mathbf{R} \triangleleft \llbracket P \rrbracket)_k)) = (\mathbf{R} \triangleleft \llbracket P \rrbracket)_k$$

and

$$\bar{f}(g(D_k)) = D_k.$$

Consequently, g is a *bijection* from Ω_k to Π_k and \bar{f} is a bijection from Π_k to Ω_k .

PROOF The proof is immediate from the above theorems and the definitions of inverse functions and bijections. \square




Figure 3.2: Bijection between kernel and transition system

3.7 Correctness of translations

3.7.1 Correctness of g

As mentioned before, in order to demonstrate that Theorem 3.2.1, that the transition system simulates π Demos, is not vacuously true through choosing an artificial definition of g , we must show the following property, which relates to the soundness of the transition system with respect to the operational semantics of π Demos. This is illustrated in Figure 3.1. We make use of the soundness property of the transition system (Theorem 3.3.1) and Theorem 3.6.1.

Theorem 3.7.1 Let \Longrightarrow^* represent the reflexive, transitive closure of the transitions derivable according to the operational semantics. Then $\forall D, D' \in \Pi, \sigma \in Act^*$

$$g(D) \xrightarrow{\sigma} g(D') \text{ implies } D_k \Longrightarrow^* D'_k.$$

PROOF By Theorem 3.3.1,

$$\mathbf{R} \triangleleft [P] \xrightarrow{\sigma} \mathbf{R}' \triangleleft [P'] \text{ implies } \bar{f}(\mathbf{R} \triangleleft [P]) \Longrightarrow^* \bar{f}(\mathbf{R}' \triangleleft [P']).$$

The property applies for every sub-transition to derive $\mathbf{R} \triangleleft [P] \xrightarrow{\sigma} \mathbf{R}' \triangleleft [P']$, so

$$\mathbf{R} \triangleleft [P] \xrightarrow{\sigma} \mathbf{R}' \triangleleft [P'].$$

Now, $\mathbf{R} \triangleleft [P]$ ranges over $g(D_k)$, so

$$g(D_k) \xrightarrow{\sigma} g(D'_k) \text{ implies } \bar{f}(g(D_k)) \Longrightarrow^* \bar{f}(g(D'_k)).$$

By Theorem 3.6.1,

$$g(D_k) \xrightarrow{\sigma} g(D'_k) \text{ implies } D_k \Longrightarrow^* D'_k.$$

Finally, because $g(D) \approx g(D_k)$,

$$g(D) \xrightarrow{\sigma} g(D') \text{ implies } D_k \Longrightarrow^* D'_k,$$

as required. □

3.7.2 Correctness of \bar{f}

Similarly, to demonstrate that we have not chosen an artificial \bar{f} to make Theorem 3.3.1, the soundness of the transition system, vacuously true, we require the following theorem, which relies on the simulation theorem (Theorem 3.2.1) and on Theorem 3.6.2.

Theorem 3.7.2 Let $\xrightarrow{\sigma^*}$ represent the reflexive, transitive closure of the transition relation $\xrightarrow{\sigma}$. Then for all $\mathbf{R} \triangleleft \llbracket P \rrbracket \in \Omega$,

$$\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket) \implies \bar{f}(\mathbf{R}' \triangleleft \llbracket P' \rrbracket) \text{ implies } \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma^*} \mathbf{R}' \triangleleft \llbracket P' \rrbracket$$

PROOF By Theorem 3.2.1, for all $D, D' \in \Pi$,

$$D \implies D' \text{ implies } g(D) \xrightarrow{\sigma^*} g(D').$$

Now, $\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket), \bar{f}(\mathbf{R}' \triangleleft \llbracket P' \rrbracket) \in \Pi$, so

$$\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket) \implies \bar{f}(\mathbf{R}' \triangleleft \llbracket P' \rrbracket) \text{ implies } g(\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket)) \xrightarrow{\sigma^*} g(\bar{f}(\mathbf{R}' \triangleleft \llbracket P' \rrbracket)).$$

Theorem 3.6.2 gives us

$$\bar{f}(\mathbf{R} \triangleleft \llbracket P \rrbracket) \implies \bar{f}(\mathbf{R}' \triangleleft \llbracket P' \rrbracket) \text{ implies } \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma^*} \mathbf{R}' \triangleleft \llbracket P' \rrbracket,$$

as required. □

Chapter 4

Demos and PBI

We have now defined, in abstract terms, the logic that we shall use, **PBI**, and a transition system for Demos programs. It remains to consider how we integrate the two and what we are able to express in the resulting system.

4.1 PBI and σ Demos

Recall that the definition of the semantics of **PBI** given in §1.4 was by a monoid,

$$\mathcal{M} = \langle \mathit{Res} \times \mathit{Prog}, e, \cdot, \sqsubseteq \rangle.$$

The monoid that we shall use to define the logic for Demos models shall be denoted M_D ,

$$\mathcal{M}_D \stackrel{\text{def}}{=} \langle \mathcal{R} \times \mathcal{P}_{\text{ROC}}, \bigcup \triangleleft \llbracket \varepsilon \rrbracket, \circ, =_\alpha \rangle.$$

The composition relation \circ combines states in $\Omega = \mathcal{R} \times \mathcal{P}_{\text{ROC}}$. It simply forms the multiset union of the resource component and the parallel composition of the process component.

$$(\mathbf{R} \triangleleft \llbracket P \rrbracket) \circ (\mathbf{S} \triangleleft \llbracket Q \rrbracket) \stackrel{\text{def}}{=} \mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P \parallel Q \rrbracket$$

The comparison pre-order $=_\alpha$ is a simple equivalence,

$$\mathbf{R} \triangleleft \llbracket P \rrbracket =_\alpha \mathbf{S} \triangleleft \llbracket Q \rrbracket \stackrel{\text{def}}{\iff} \mathbf{R} \equiv \mathbf{S} \wedge P \equiv Q.$$

$=_\alpha$ considers pairs to be equal iff they have the same resource multisets and their process terms are the same, up to the commutativity and associativity of parallelisation, with the inclusion of an arbitrary number of $\llbracket \varepsilon \rrbracket$ processes. Thus $\bigcup \triangleleft \llbracket \varepsilon \rrbracket$ is the identity element under \circ .

The semantics of **PBI** following from these definitions is presented in Table 4.1.

Finally, we must define the relations \uparrow and \downarrow for the multiplicative modalities. These define, respectively, the resources acquired and released by particular

$\mathbf{R} \triangleleft \llbracket P \rrbracket \models r$	iff	$\mathbf{R} \triangleleft \llbracket P \rrbracket =_{\alpha} \wr r \wr \triangleleft \llbracket \varepsilon \rrbracket$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models Q$	iff	$\mathbf{R} \triangleleft \llbracket P \rrbracket =_{\alpha} \wr \wr \triangleleft \llbracket Q \rrbracket$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi \wedge \psi$	iff	$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \models \psi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi \vee \psi$	iff	$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi$ or $\mathbf{R} \triangleleft \llbracket P \rrbracket \models \psi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi \rightarrow \psi$	iff	$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi$ implies $\mathbf{R} \triangleleft \llbracket P \rrbracket \models \psi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi * \psi$	iff	for some $\mathbf{S} \triangleleft \llbracket Q \rrbracket, \mathbf{S}' \triangleleft \llbracket Q' \rrbracket \in \Omega$ $[\mathbf{R} \triangleleft \llbracket P \rrbracket = \mathbf{S} \uplus \mathbf{S}' \triangleleft \llbracket Q \parallel Q' \rrbracket$ and $\mathbf{S} \triangleleft \llbracket Q \rrbracket \models \varphi$ and $\mathbf{S}' \triangleleft \llbracket Q' \rrbracket \models \psi \]$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi \multimap \psi$	iff	for all $\mathbf{S} \triangleleft \llbracket Q \rrbracket \in \Omega$ $\mathbf{S} \triangleleft \llbracket Q \rrbracket \models \varphi$ implies $\mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P \parallel Q \rrbracket \models \psi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle A \rangle \varphi$	iff	for some $\alpha \in A$ $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\alpha} \mathbf{R} \triangleleft \llbracket P' \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P' \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models [A] \varphi$	iff	for all $\alpha \in A$ $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\alpha} \mathbf{R} \triangleleft \llbracket P' \rrbracket$ implies $\mathbf{R} \triangleleft \llbracket P' \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle A \rangle_{\text{new}}^+ \varphi$	iff	$\exists a \in A.$ $a \upharpoonright \mathbf{S} \triangleleft \llbracket Q \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} \mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P' \parallel Q \rrbracket$ and $\mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P' \parallel Q \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models [A]_{\text{new}}^+ \varphi$	iff	$\forall a \in A.$ $a \upharpoonright \mathbf{S} \triangleleft \llbracket Q \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{a} \mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P' \parallel Q \rrbracket$ implies $\mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P' \parallel Q \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle A \rangle_{\text{new}}^- \varphi$	iff	$\exists a \in A. a \upharpoonright \mathbf{S} \triangleleft \llbracket Q_1 \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket = \mathbf{S} \uplus \mathbf{T} \triangleleft \llbracket Q_1 \parallel Q_2 \rrbracket$ and $\mathbf{S} \uplus \mathbf{T} \triangleleft \llbracket Q_1 \parallel Q_2 \rrbracket \xrightarrow{a} \mathbf{T} \triangleleft \llbracket Q'_2 \rrbracket$ and $\mathbf{T} \triangleleft \llbracket Q'_2 \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models [A]_{\text{new}}^- \varphi$	iff	$\forall a \in A. a \upharpoonright \mathbf{S} \triangleleft \llbracket Q_1 \rrbracket$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket = \mathbf{S} \uplus \mathbf{T} \triangleleft \llbracket Q_1 \parallel Q_2 \rrbracket$ and $\mathbf{S} \uplus \mathbf{T} \triangleleft \llbracket Q_1 \parallel Q_2 \rrbracket \xrightarrow{a} \mathbf{T} \triangleleft \llbracket Q'_2 \rrbracket$ implies $\mathbf{T} \triangleleft \llbracket Q'_2 \rrbracket \models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \neg \varphi$	iff	$\mathbf{R} \triangleleft \llbracket P \rrbracket \not\models \varphi$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \top$	always	
$\mathbf{R} \triangleleft \llbracket P \rrbracket \not\models \perp$	for any	$\mathbf{R} \triangleleft \llbracket P \rrbracket$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models I_R$	iff	$\exists P \in \mathcal{P}_{\text{ROC}}. \mathbf{R} \triangleleft \llbracket P \rrbracket =_{\alpha} \wr \wr \triangleleft \llbracket P \rrbracket$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models I_P$	iff	$\exists \mathbf{S} \in \mathcal{R}. \mathbf{R} \triangleleft \llbracket P \rrbracket =_{\alpha} \mathbf{S} \triangleleft \llbracket \varepsilon \rrbracket$
$\mathbf{R} \triangleleft \llbracket P \rrbracket \models I$	iff	$\mathbf{R} \triangleleft \llbracket P \rrbracket =_{\alpha} \wr \wr \triangleleft \llbracket \varepsilon \rrbracket$

Table 4.1: Demos semantics of **PBI**

actions.

$$\begin{aligned} \text{getR}_r^n & \uparrow \{r^n\} \triangleleft \llbracket \varepsilon \rrbracket \\ \text{getB}_b^n & \uparrow \{b^n\} \triangleleft \llbracket \varepsilon \rrbracket \\ \text{getS}_x^n & \uparrow \{x^n\} \triangleleft \llbracket \varepsilon \rrbracket \end{aligned}$$

$$\begin{aligned} \text{putR}_r^n & \uparrow \{r^n\} \triangleleft \llbracket \varepsilon \rrbracket \\ \text{putB}_b^n & \uparrow \{b^n\} \triangleleft \llbracket \varepsilon \rrbracket \\ \text{putS}_x^n & \uparrow \{x^n\} \triangleleft \llbracket \varepsilon \rrbracket \end{aligned}$$

4.1.1 Intuitive account

- Resource propositions r hold in states with no process component where r is the only resource available.
- Process propositions P hold in states with null resource component with P being the only process.
- The additive unit of \wedge is \top , which holds in any state.
- The additive unit of \vee is \perp , which holds in no state.
- The multiplicative resource unit I_R holds in states with no resource component, but allows an arbitrary (possibly parallel) process.
- The multiplicative resource unit I_P holds in states with only a null process component, but allows an arbitrary resource component.
- The (extraneous) multiplicative unit I holds in states with only a null process component and no available resources.
- The additive logical operators \wedge, \vee and \rightarrow are all as in classical (boolean) logic. \wedge allows us to assert that two propositions hold in a state, \vee allows us to assert that at least one of two propositions holds in a state, and \rightarrow allows us to assert that if one proposition holds in a state, so does another.
- The multiplicative conjunction of two propositions will hold in a state, $\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi * \psi$, if the state can be partitioned into two parts, the partition being of resources and processes, with one part making φ hold and the other making ψ hold.
- Multiplicative implication relates to the addition of resources or processes: $\mathbf{R} \triangleleft \llbracket P \rrbracket \models \varphi \multimap \psi$ holds if all states in which φ holds when composed with the state $\mathbf{R} \triangleleft \llbracket P \rrbracket$ give states in which ψ holds.
- $\langle A \rangle \varphi$ will hold if a state can evolve to a new state by a single action $\alpha \in A$ in which φ holds given no change in the resource component of the state. Similarly, $[A] \varphi$ holds in a state if all transitions labelled $\alpha \in A$ lead to states in which φ holds, again with no change in the resource component.

- $\mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle A \rangle^+ \varphi$ will hold if by some action $\sigma \in A$, there is a transition to a new state with a *larger* resource collection, $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P' \rrbracket$, with $\sigma \upharpoonright \mathbf{S}$, and $\mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P' \rrbracket \models \varphi$. Similarly, $[A]^+$ relates to all transitions with labels in A leading to states in which a proposition holds.
- Dually, the negative multiplicative modalities relate to transitions leading to states with *smaller* resource collections. They relate to transitions of the form $\mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R} \triangleleft \llbracket P' \rrbracket$, with the difference in resource specified by $\sigma \upharpoonright \mathbf{S}$.
- Finally, because our pre-order is an equality, our logic is classical so we have negation, \neg . This allows us to assert that a proposition does not hold in a given state. The inclusion of negation makes \rightarrow redundant: we could use the equivalence $\varphi \rightarrow \psi \Leftrightarrow \neg\varphi \vee \psi$.

If we want to reason about an arbitrary type of action occurring, we would form the disjunction of diamond modalities or the conjunction of the box modalities. We shall introduce the following abbreviations:

$$\begin{aligned} \mathbf{R} \triangleleft \llbracket P \rrbracket \models [A]_{\mathbf{any}} \varphi & \text{ iff } \mathbf{R} \triangleleft \llbracket P \rrbracket \models [A]_{\mathbf{new}}^- \varphi \wedge [A]_{\mathbf{new}}^+ \varphi \wedge [A] \varphi \\ \mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle A \rangle_{\mathbf{any}} \varphi & \text{ iff } \mathbf{R} \triangleleft \llbracket P \rrbracket \models \langle A \rangle_{\mathbf{new}}^+ \varphi \vee \langle A \rangle_{\mathbf{new}}^- \varphi \vee \langle A \rangle \varphi. \end{aligned}$$

Recalling the definition of I_R , the formula $r * I_R$ holds in states where r is the only available resource with an arbitrary process term. We shall soon see a formula representing r being *one* of the free resources. Recalling the definition of I_P , the formula $P * I_P$ holds in states where P is the only process in the state with an arbitrary collection of resources.

4.2 Deadlock in PBI

As presented in §1.3.2, deadlock freedom in Hennessy-Milner logic is defined as:

$$\mathfrak{D}_f^{hml} = \bigwedge_{i \in \mathbb{N} \cup \{\omega\}} [A]^i \langle A \rangle \top$$

We must make a few adaptations to this to make it suitable for us. Firstly, we must select an appropriate modality: we shall choose the **any** modalities. Secondly, we wish to relax the requirement that the process never terminates to requiring that if it is in a non-terminated state, it can take an action. We do this by requiring I_P , the proposition indicating clean termination, or non-termination in every state. Consequently, the **PBI** formula for deadlock freedom in Demos is:

$$\mathfrak{D}_f^D = \bigwedge_{i \in \mathbb{N} \cup \{\omega\}} [-]_{\mathbf{any}}^i (\langle A \rangle_{\mathbf{any}} \top \vee I_P)$$

4.3 Examples

We shall start by considering the modalities. The additive modalities, as we have said, relate to transitions in which the resource component of state does not change. We have, then,

$$\mathbf{R} \triangleleft \llbracket \mathbf{hold}(5); \mathbf{putR}(r, 1) \rrbracket \models \langle - \rangle \top,$$

which expresses that the process can act without changing the available resource collection. We also have

$$\lambda r \int \triangleleft \llbracket \mathbf{hold}(5); \mathbf{putR}(r, 1) \rrbracket \parallel \llbracket \mathbf{getR}(r, 1); \mathbf{putR}(r, 1) \rrbracket \models [-](r * I_R).$$

which expresses that every action that does not change the available resource collection gives a state in which r is the only available resource. Notice that we do not have $\models \langle - \rangle \neg(r * I_R)$: the \mathbf{getR} transition, the only transition leading to a state in which r is unavailable, does not match the semantics of the additive modality.

The additive modality does not refer to actions that will change the available resource collection, so

$$\mathbf{R} \triangleleft \llbracket \mathbf{putR}(r, 1) \rrbracket \models [-] \perp,$$

which expresses that the process can make no action that does not change the available resource collection.

The multiplicative release modality is used for $\mathbf{put}X$ transitions. So, for example,

$$\lambda r, s \int \triangleleft \llbracket \mathbf{putB}(b, 1) \rrbracket \models \langle - \rangle_{\mathbf{new}}^+ (b * r * s)$$

by the (only) transition

$$\lambda s, r \int \triangleleft \llbracket \mathbf{putB}(b, 1) \rrbracket \xrightarrow{\mathbf{putB}_b^1} \lambda b, r, s \int \triangleleft \llbracket \varepsilon \rrbracket$$

and that $\mathbf{putB}_b^1 \upharpoonright \lambda b \int \triangleleft \llbracket \varepsilon \rrbracket$.

The multiplicative acquisition modality is used for $\mathbf{get}X$ transitions:

$$\lambda b, r, s \int \triangleleft \llbracket \mathbf{putB}(b, 1) \rrbracket \parallel \llbracket \mathbf{getB}(b, 1) \rrbracket \models [-]_{\mathbf{new}}^- (r * s * I_R)$$

by the only \mathbf{get} transition

$$\lambda b, r, s \int \triangleleft \llbracket \mathbf{putB}(b, 1) \rrbracket \parallel \llbracket \mathbf{getB}(b, 1) \rrbracket \xrightarrow{\mathbf{getB}_b^1} \lambda r, s \int \triangleleft \llbracket \mathbf{putB}(b, 1) \rrbracket \parallel \llbracket \varepsilon \rrbracket$$

In general, we will have the following:

$$\mathbf{R} \triangleleft \llbracket P \rrbracket \models \begin{array}{l} [-]_{\mathbf{new}}^+ \varphi_+ \\ \wedge [-]_{\mathbf{new}}^- \varphi_- \\ \wedge [-] \varphi \end{array}$$

where φ_+ is some formula true in all states following a transition by any process with \mathbf{put} at its head, φ_- following a transition by any process with \mathbf{get} at its head, and φ following a transition by any other process.

* relates to *partitions* of the state, both of processes and resources. Returning to the intelligence agency example, recalling that \top holds for *any* state, we see that the following holds:

$$\{C, F, 6, 5\} \triangleleft \llbracket P \rrbracket \models F * \top$$

where C denotes the CIA line, 6 denotes the MI6 line *etc.* by $\{F\} \triangleleft \llbracket \varepsilon \rrbracket \models F$ and $\{C, 6, 5\} \triangleleft \llbracket P \rrbracket \models \top$. So in general, just as we can in propositional **BI**, we may indicate that a resource r is available by the formula $r * \top$.

Returning to the jobshop example, we may wish to express that h and m are the only available resources. Clearly, $\{h, m\} \triangleleft \llbracket P \rrbracket \models h * m * \top$, but also $\{h, m, x\} \triangleleft \llbracket P \rrbracket \models h * m * \top$. Rather than using \top , we should have used I_R : $\{h, m\} \triangleleft \llbracket P \rrbracket \models h * m * I_R$. This expresses that there are three partitions in the state, one having only the resource h , another having the resource m and the final one being an arbitrary process with no resource.

Suppose that we have a proposition ε_C which denotes termination of the CIA process. As has been illustrated, it is not the case that in every possible execution the CIA process terminates.¹

$$\{C, F, 6, 5\} \triangleleft \llbracket P \rrbracket \not\models \bigwedge_{i \in \mathbb{N}_\omega} [-]_{\mathbf{any}}^i (\langle - \rangle_{\mathbf{any}} \top \vee \varepsilon_C)$$

This formula expresses that there is a state from which the system can make no transition and the CIA process has not terminated cleanly. We might like to consider the effect of removing MI6's line: would this allow the CIA to complete its task in all conditions? We may express this in the logic as:

$$\{C, F, 6, 5\} \triangleleft \llbracket P \rrbracket \stackrel{?}{=} 6 * \bigwedge_{i \in \mathbb{N}_\omega} [-]_{\mathbf{any}}^i (\langle - \rangle_{\mathbf{any}} \top \vee \varepsilon_C)$$

Now, $\{6\} \triangleleft \llbracket \varepsilon \rrbracket \models 6$. For the above formula to be true, then, all we have to show is that $\{C, F, 5\} \triangleleft \llbracket P \rrbracket \models \bigwedge_{i \in \mathbb{N}_\omega} [-]_{\mathbf{any}}^i (\langle - \rangle_{\mathbf{any}} \top \vee \varepsilon_C)$. (Notice that the choice of partition here is entirely deterministic.) Because MI6's line is removed, the MI6 process will not be able to proceed to take MI5's line or the FBI's. Consequently, the CIA process will terminate properly, so we do have²:

$$\{C, F, 6, 5\} \triangleleft \llbracket P \rrbracket \models 6 * \bigwedge_{i \in \mathbb{N}_\omega} [-]_{\mathbf{any}}^i (\langle - \rangle_{\mathbf{any}} \top \vee \varepsilon_C).$$

This illustrates how $*$ may be used to reason about *subtraction* of resources from a system.

It is a simple matter to ask if a partition exists that is non-deadlocking:

$$\mathbf{R} \triangleleft \llbracket P \rrbracket \stackrel{?}{=} \mathcal{D}_f^D * \mathcal{D}_f^D$$

This will hold only if there are deadlock-free partitions $\mathbf{S}_1 \triangleleft \llbracket P_1 \rrbracket$ and $\mathbf{S}_2 \triangleleft \llbracket P_2 \rrbracket$ such that $\mathbf{R} \triangleleft \llbracket P \rrbracket = \mathbf{S}_1 \uplus \mathbf{S}_2 \triangleleft \llbracket P_1 \parallel P_2 \rrbracket$. If we require a non-trivial partition (*i.e.*, without one partition being null-process), we would ask:

$$\mathbf{R} \triangleleft \llbracket P \rrbracket \stackrel{?}{=} (\mathcal{D}_f^D \wedge \neg I_P) * (\mathcal{D}_f^D \wedge \neg I_P)$$

¹The following formula comes from the definition of deadlock — see §4.2

²But not overall deadlock freedom

Now let us consider $\neg*$. Suppose that we want to consider the effect of installing a new line from the hub to the FBI. Would this result in deadlock freedom? The question would be posed in the logic as:

$$\wr C, F, 6, 5 \wr \triangleleft \llbracket P \rrbracket \stackrel{?}{=} F \neg* \mathfrak{D}_f^D$$

The only state (model) satisfying $\models F$ is $\wr F \wr \triangleleft \llbracket \varepsilon \rrbracket$. Thus to answer the above question, we must resolve satisfaction of $\wr C, F, F, 6, 5 \wr \triangleleft \llbracket P \rrbracket \stackrel{?}{=} \mathfrak{D}_f^D$. It turns out that this is true: either the CIA process or the MI6 process may end up blocked, but the other will eventually release the resource that it is waiting on.

The multiplicative modalities are also quite interesting: we could now use the positive (given the possible inclusion of processes, a better term might be *composition*) modalities $\langle A \rangle_{\text{new}}^+$ and $[A]_{\text{new}}^+$ to indicate a new process being generated. This occurs with an *Entity* command, and would require \wr to be extended by

$$\tau_{\text{Entity}}^{\text{class}} \wr \wr \triangleleft \llbracket \text{class} \rrbracket.$$

Similarly, we could use the positive (a better term now might be *decomposition*) modalities to indicate processes being removed. This could be through the termination case of parallel composition or through an explicit *kill* command (which does not exist in Demos). With such rich multiplicative modalities, it may be desirable to exclude the case where the number of processes changes in this manner from the additive modality. Further consideration of this, though, is beyond our current scope. We shall, therefore, use the additive modalities for Entity declarations *etc.*, and leave \wr and \wr defined only where necessary: on the resource actions.

It should be noted, though, that the total composition relation described here makes proof-search very non-deterministic. Take, for example, the formula $\langle - \rangle_{\text{new}}^+ \top \neg* \varphi$. Not allowed if we define composition *partially*, $\mathbf{R} \triangleleft \llbracket P \rrbracket \circ \mathbf{S} \triangleleft \llbracket \varepsilon \rrbracket = \mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P \rrbracket$, to evaluate this we may have to consider adding every σ Demos process that can release a resource!

4.3.1 Sub-deadlock

Consider again the jobshop example of §2.4. The worker process (the input stream of jobs to do) can, as it is, never deadlock. Consequently, if we consider the whole process we see that it is always possible to derive a transition using the worker process³, so the whole system may never deadlock. Now, this is not of very much use to us: what we really want to know is whether or not the jobbers deadlock.

Deadlock of a process, whether it be a sub-process or not, is defined to occur where no transition is derivable at any reachable future state. For a sub-process $P_{k[\ell]}$, then, a state is deadlock-free with respect to $P_{k[\ell]}$ iff there is no reachable state from which $P_{k[\ell]}$ does not act in any transition from any future state. We can write this in **PBI** as:

$$\mathfrak{D}_f^{P,k} \stackrel{\text{def}}{\iff} \neg \bigvee \langle - \rangle_{\text{any}} \left(\neg \varepsilon_k \wedge \bigwedge [-]_{\text{any}}^j [K]_{\text{any}} \perp \right)$$

³This is not the case if we consider a *fair* transition system, §6.1

Note that we have used a set of action labels indicating action by $P_{k[\ell]}$, K . This set can be defined by pattern-matching,

$$K = \{a \mid a \in \text{Labels and } a = \ell : _ \}$$

Chapter 5

Synchronous parallel rule

5.1 Motivation

Consider a kitchen in which there is a knife, a pair of scissors, a mother and her child. The mother uses the knife and then the scissors before returning them. Adopting our Demos representation for processes, and letting k represent the knife and s the pair of scissors, her activity is modelled by the following sequence of actions:

$$P_m = \text{getR}(k, 1); \text{getR}(s, 1); \text{hold}(5); \text{putR}(s, 1); \text{putR}(k, 1).$$

The child only wants to use the scissors for a short period of time, and they are not allowed to use the knife. Their Demos representation is:

$$P_c = \text{getR}(s, 1); \text{hold}(1); \text{putR}(s, 1),$$


We want to ask if having the mother and the child in the kitchen at the same time will cause deadlock. Letting putative consequence (asking whether a formula is satisfied in a particular model) be written $?=$, the question is represented formally by:

$$\{s, k\} \triangleleft \ll [P_m \parallel P_c] \text{ ?} = \mathfrak{D}_f^D.$$

In the logics presented above, HML, CTL and modal- μ , the complexity of model-checking such questions increases as the state space generated by the transition system increases. Figure 5.1 shows the program fragment of the states that may be generated using the transition system with the parallel composition rule presented above.

Branching occurs where more than one transition is derivable from a particular state — there is *non-deterministic choice* of the next state to enter into. For example, at state $(1, 1)$, either the child may take the scissors or the mother may take the knife. For the transition system presented above, the only time we may have such non-deterministic choice is where we have the parallel composition of processes where more than one is active (*i.e.*, not waiting on a resource).

The astute reader may have noticed that, when considering deadlock, it was not essential to distinguish whether the first action was or was not the mother



The axes are labelled by numbers representing the line of code being executed:

	P_m	P_c
1	getR(k , 1)	getR(s , 1)
2	getR(s , 1)	hold(1)
3	hold(5)	putR(s , 1)
4	putR(s , 1)	ε
5	putR(k , 1)	—
6	ε	—

So, for example, $(4, 3) = \llbracket \text{putR}(s, 1); \text{putR}(k, 1) \parallel \text{putR}(s, 1) \rrbracket$.

Arrowed lines indicate which states may be derived from each state.

Figure 5.1: Reachable state space for kitchen example derived by asynchronous transition system

picking up the knife, and then considering the real interaction. Why, then, did we need to generate a whole collection of states to distinguish this possibility? In terms of process interaction, it would have been more efficient to force the first action to be the mother taking the knife. Of course, the reader could argue that it is the rôle of the modeller to weed-out such irrelevances. Unfortunately, though, the modeller may not know in advance which processes we wish to model interaction between, so this cannot be expected.

5.2 Soundness: try and resource splitting

Let us consider how we may formulate a rule that allows synchronous execution of processes. If we include the simple, intuitive rule

$$(\text{SPAR}) : \frac{\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket \quad \mathbf{S} \triangleleft \llbracket Q \rrbracket \xrightarrow{\sigma} \mathbf{S}' \llbracket Q' \rrbracket}{\mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P \parallel Q \rrbracket \xrightarrow{\sigma} \mathbf{R}' \uplus \mathbf{S}' \triangleleft \llbracket P' \parallel Q' \rrbracket},$$

the semantics of the transition system become unsound with respect to the operational semantics of Demos. The unsoundness comes, in part, from the fact that we have considered only a simple split of resources: in deducing a transition, we may allocate one processes resources it does not need, thereby depriving the other process of resources. This error may manifest itself in a `try` command choosing the wrong branch.

To give a concrete example of this, consider the jobshop example (§2.4) with one extra hammer. Suppose that one worker has taken a medium-difficulty job to do, and the other has taken a hard job to do. The state in Ω , then, is:

$$\{m, h, h\} \triangleleft \llbracket MID; J \parallel HARD; J \rrbracket$$

where *MID* is the procedure for a medium-difficulty process

```
try [getR(m, 1)] then {hold(20); trace('Mid'); putR(m, 1)}
etry [getR(h, 1)] then {hold(10); trace('Mid'); putR(h, 1)} ,
```

and *HARD* denotes

```
getR(h, 1); hold(20); trace('Hard'); putR(h, 1)
```

Define

```
MIDh  ≡ hold(10); trace('Mid'); putR(h, 1)
HARDh ≡ hold(20); trace('Hard'); putR(h, 1),
```

representing the code to be executed if we have a middle- or hard-difficulty job, respectively, and have picked up a hammer.

The derivation from this state, using the unsound transition rule, is presented in Figure 5.2.

Clearly this is wrong — we did not use the mallet on the easy job as we should have done. The error came from the fact that it was possible for us to ‘give’ the mallet to the process executing *HARD*, which did not need it, whilst *MID* did not know that the mallet existed.

For this reason, a logical resource weakening rule for transitions,

$$\text{(WEAK)} : \frac{\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket}{\mathbf{R} \uplus \mathbf{S} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma} \mathbf{R}' \uplus \mathbf{S} \triangleleft \llbracket P' \rrbracket},$$

would also be unsound.

Two solutions exist to overcome this particular problem. One is to adopt the rule:

$$\text{(SPAR)} : \frac{\mathbf{R} \uplus \mathbf{H} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma_1} \mathbf{H} \triangleleft \llbracket P' \rrbracket \quad \mathbf{R}' \uplus \mathbf{H} \triangleleft \llbracket Q \rrbracket \xrightarrow{\sigma_2} \mathbf{H} \triangleleft \llbracket Q' \rrbracket}{\mathbf{R} \uplus \mathbf{R}' \uplus \mathbf{H} \triangleleft \llbracket P \parallel Q \rrbracket \xrightarrow{\sigma_1 \uparrow \sigma_2} \mathbf{H} \triangleleft \llbracket P' \parallel Q' \rrbracket}.$$

This allows both derivations to ‘see’ all the unused resources in the system.

The second solution is to disallow the command at the head of any process being ‘try’ when considering the synchronous parallel rule. We shall adopt the first solution.

5.3 Soundness: Synchrony

The presentation so far has failed to take into one thing into account: we do not necessarily want to consider processes executing synchronously. Take, for example, the following processes:

$$\begin{aligned} P_1 &\stackrel{\text{def}}{=} \text{getR}(r, 1); \text{getR}(s, 1); \text{putR}(r, 1); \text{putR}(s, 1) \\ P_2 &\stackrel{\text{def}}{=} \text{do}(2)\{\text{putB}(b, 1); \} \text{getR}(s, 1); \text{getR}(r, 1); \dots \end{aligned}$$

The only path derivable from here using the synchronous rule is:

$$\begin{aligned} &\{r, s\} \triangleleft \llbracket P_1 \parallel P_2 \rrbracket \\ \xrightarrow{\text{getR} \upharpoonright \text{putB}} &\{s, b\} \triangleleft \left[\begin{array}{l} \text{getR}(s, 1); \text{putR}(r, 1); \text{putR}(s, 1) \\ \parallel \\ \text{do}(1)\{\text{putB}(b, 1); \} \text{getR}(s, 1); \text{getR}(r, 1); \dots \end{array} \right] \\ \xrightarrow{\text{getR} \upharpoonright \text{putB}} &\{b, b\} \triangleleft \left[\begin{array}{l} \text{putR}(r, 1); \text{putR}(s, 1) \\ \parallel \\ \text{getR}(s, 1); \text{getR}(r, 1); \dots \end{array} \right] \\ \Rightarrow &\dots \end{aligned}$$

By executing synchronously, we missed the deadlock that would have occurred if the two **putB**s had executed first. In effect, the transition system that we have described so far in this section could be simulated in the operational semantics by removing all **hold** commands and inserting **hold**(*k*), for some constant *k*, between all the other commands.

How, then, can we reconcile our desire for a state space reduction with the necessity of considering all interactions between processes? The answer lies in

identifying which commands could affect the execution of other processes — we shall say that they *interfere* with other processes. We can choose to branch only on these commands, whilst executing non-interfering commands synchronously. So, for example, if a process has a series of commands on a resource that is not referenced by any other process, we will not introduce any new paths to be examined when model checking.

We shall, therefore, require a condition, (*), that dictates when to use the synchronous parallel rule and when to use the asynchronous rule(s). To capture interference some auxilliary functions are required. Firstly, the multiset of resources accessed by a non-parallel process:

$$\begin{aligned}
res(\mathbf{getR}(r, n)) &\stackrel{\text{def}}{=} \wr r^n \wr \\
res(\mathbf{putR}(r, n)) &\stackrel{\text{def}}{=} \wr r \wr \\
res(\mathbf{getB}(b, n)) &\stackrel{\text{def}}{=} \wr b^n \wr \\
res(\mathbf{putB}(b, n)) &\stackrel{\text{def}}{=} \wr b \wr \\
res(\mathbf{sync}(x)) &\stackrel{\text{def}}{=} \wr \mathbf{sync}_x \wr \\
res(\mathbf{getS}(x, n)) &\stackrel{\text{def}}{=} \wr \mathbf{sync}_x^n \wr \\
res(c; P) &\stackrel{\text{def}}{=} res(c) \uplus res(P) \\
res(c, C) &\stackrel{\text{def}}{=} res(c) \uplus res(C) \\
res\left(\begin{array}{l} \mathbf{try} [e_1] \mathbf{then} P_1 \\ \mathbf{etry} [e_2] \mathbf{then} P_2 \end{array}\right) &\stackrel{\text{def}}{=} res(e_1) \uplus res(e_2) \uplus res(P_1) \uplus res(P_2) \\
res(\text{anything else}) &\stackrel{\text{def}}{=} \wr \wr
\end{aligned}$$

where \mathbf{sync}_x is a token indicating synchronisation on a pool x .

The command that may next be executed by a non-parallel process is its head, defined as follows:

$$head(c) \stackrel{\text{def}}{=} c \quad head(c; P) \stackrel{\text{def}}{=} c$$

The function \overrightarrow{res} shall return the multiset of resources accessed by parallel processes:

$$\begin{aligned}
\overrightarrow{res}(\llbracket \varepsilon \rrbracket) &\stackrel{\text{def}}{=} \wr \wr \\
\overrightarrow{res}(\llbracket P \parallel Q \rrbracket) &\stackrel{\text{def}}{=} res(P) \uplus \overrightarrow{res}(Q), \quad P \text{ not parallel}
\end{aligned}$$

Take the set of indices of a parallel process P to be In_P ; that is, if $P = \llbracket P_1 \parallel \dots \parallel P_n \rrbracket$ then $In_P = \{1, \dots, n\}$. The set of indices of non-interfering processes (processes with non-interfering commands at their head) shall be written C'_P , defined as:

$$C'_P \stackrel{\text{def}}{=} \left\{ i \mid res(head(P_i)) \uplus \overrightarrow{res}\left(\left\| \parallel_{k \in In_P \setminus \{i\}} P_k \right\| \right) = \emptyset \right\}$$

To ensure that the synchronous parallel rule does not prevent us entering the true deadlocked state, we shall also require that it only be applied to processes that are able to act. This set of processes is:

$$A_P^{\mathbf{R}} \stackrel{\text{def}}{=} \{i \mid i \in In_P \wedge (head(P_i) = \mathbf{getR}(r, n) \implies res(head(P_i)) \subseteq \mathbf{R})\}$$

We shall execute all non-interfering processes synchronously until there are none left, at which point we shall use the non-deterministic asynchronous parallel rules. Thus the condition is:

$$C'_P \cap A_P^{\mathbf{R}} \neq \emptyset \quad (*)$$

The synchronous parallel execution rule, then, is:


$$(\text{SPAR}) : \frac{\mathbf{R}_i \uplus \mathbf{H} \triangleleft \llbracket P_i \rrbracket \xrightarrow{\sigma_i} \mathbf{H} \triangleleft \llbracket P'_i \rrbracket}{\mathbf{H}(\biguplus_{i \in I'_P} \mathbf{R}_i) \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma_i |_{i \in I}} \mathbf{H} \triangleleft \llbracket \parallel_{k \notin I'_P} P_k \parallel_{i \in I'_P} P'_i \rrbracket}} \quad (*)$$

The asynchronous parallel execution rules are unchanged, other than having the condition $\neg(*)$.

In Figure 5.3, we present an example series of transitions derivable with this rule. Notice that there is no choice of which transition to derive — this particular path is derived completely deterministically from the given state. Following a τ -transition to remove the null process, we would now use the asynchronous transition rules to consider the interfering executions. Figure 5.4 shows the reduced state space of the kitchen example. Notice that it is much smaller than that derived by the asynchronous system, with only one branching point.

$$\begin{array}{l}
\{r, s, t\} \triangleleft \left[\begin{array}{c} \text{getB}(r, 1) \\ \text{putR}(u, 1) \\ \text{hold}(3) \end{array} \parallel \left[\begin{array}{c} \text{getR}(u, 1) \\ \text{hold}(4) \\ \text{putR}(u, 1) \end{array} \parallel \left[\begin{array}{c} \text{getR}(t, 1) \\ \text{getR}(s, 1) \\ \text{hold}(4) \\ \text{putR}(t, 1) \\ \text{putR}(s, 1) \end{array} \right] \right] \right] \quad C'_P = \{1, 3\} \\
\begin{array}{c} \text{getB} \\ \xrightarrow{\text{getR}} \end{array} \{s\} \triangleleft \left[\begin{array}{c} \text{putR}(u, 1) \\ \text{hold}(3) \end{array} \parallel \left[\begin{array}{c} \text{getR}(u, 1) \\ \text{hold}(4) \\ \text{putR}(u, 1) \end{array} \parallel \left[\begin{array}{c} \text{getR}(s, 1) \\ \text{hold}(4) \\ \text{putR}(t, 1) \\ \text{putR}(s, 1) \end{array} \right] \right] \right] \quad C'_P = \{3\} \\
\begin{array}{c} \text{getR} \\ \xrightarrow{\quad} \end{array} \{s\} \triangleleft \left[\begin{array}{c} \text{putR}(u, 1) \\ \text{hold}(3) \end{array} \parallel \left[\begin{array}{c} \text{getR}(u, 1) \\ \text{hold}(4) \\ \text{putR}(u, 1) \end{array} \parallel \left[\begin{array}{c} \text{hold}(4) \\ \text{putR}(t, 1) \\ \text{putR}(s, 1) \end{array} \right] \right] \right] \quad C'_P = \{3\} \\
\begin{array}{c} \tau \\ \xrightarrow{\quad} \end{array} \{s\} \triangleleft \left[\begin{array}{c} \text{putR}(u, 1) \\ \text{hold}(3) \end{array} \parallel \left[\begin{array}{c} \text{getR}(u, 1) \\ \text{hold}(4) \\ \text{putR}(u, 1) \end{array} \parallel \left[\begin{array}{c} \text{putR}(t, 1) \\ \text{putR}(s, 1) \end{array} \right] \right] \right] \quad C'_P = \{3\} \\
\begin{array}{c} \text{putR} \\ \xrightarrow{\quad} \end{array} \{t\} \triangleleft \left[\begin{array}{c} \text{putR}(u, 1) \\ \text{hold}(3) \end{array} \parallel \left[\begin{array}{c} \text{getR}(u, 1) \\ \text{hold}(4) \\ \text{putR}(u, 1) \end{array} \parallel \left[\begin{array}{c} \text{putR}(s, 1) \end{array} \right] \right] \right] \quad C'_P = \{3\} \\
\begin{array}{c} \text{putR} \\ \xrightarrow{\quad} \end{array} \{s, t\} \triangleleft \left[\begin{array}{c} \text{putR}(u, 1) \\ \text{hold}(3) \end{array} \parallel \left[\begin{array}{c} \text{getR}(u, 1) \\ \text{hold}(4) \\ \text{putR}(u, 1) \end{array} \parallel \left[\begin{array}{c} \varepsilon \end{array} \right] \right] \right] \quad I'_P = \{3\}
\end{array}$$

Figure 5.3: Path derived, deterministically, using synchronous parallel rule



Filled circles represent states in which $(*)$ holds. Double lines indicate a transition derived using the synchronous parallel rule.

Figure 5.4: State space for kitchen example derived by the synchronous transition system

As mentioned before, though, the reduced state space comes at a cost. The transition system is not complete with respect to the operational semantics of Demos programs because the operational semantics defines asynchronous action. In the sequel, \models_a shall represent satisfaction in **PBI** with the normal, asynchronous semantics, and \models_s shall represent satisfaction with the synchronous parallel rule. We see, then, that \models_s is not complete with respect to \models_a .

Proposition 5.3.1 \models_s is not complete with respect to \models_a . That is, there is a formula φ and σ Demos state $\mathbf{R} \triangleleft \llbracket P \rrbracket$ with $\mathbf{R} \triangleleft \llbracket P \rrbracket \models_a \varphi$ and $\mathbf{R} \triangleleft \llbracket P \rrbracket \not\models_s \varphi$.

PROOF Take $\mathbf{R} \triangleleft \llbracket P \rrbracket$ to be the first state in the sequence above. Then $\mathbf{R} \triangleleft \llbracket P \rrbracket \models_a \langle - \rangle_{\text{any}} s * t * I_R$ but $\mathbf{R} \triangleleft \llbracket P \rrbracket \not\models_s \langle - \rangle_{\text{any}} s * t * I_R$. \square

Moreover, \models_s is not even sound with respect to \models_a (and, therefore, the operational semantics of Demos):

Proposition 5.3.2 \models_s is not sound with respect to \models_a . That is, there is a formula φ and σ Demos state $\mathbf{R} \triangleleft \llbracket P \rrbracket$ with $\mathbf{R} \triangleleft \llbracket P \rrbracket \models_s \varphi$ but $\mathbf{R} \triangleleft \llbracket P \rrbracket \not\models_a \varphi$.

PROOF With $\mathbf{R} \triangleleft \llbracket P \rrbracket$ as before, $\mathbf{R} \triangleleft \llbracket P \rrbracket \not\models_a \neg(\langle - \rangle_{\text{any}} s * t * I_R)$ but $\mathbf{R} \triangleleft \llbracket P \rrbracket \models_s \neg(\langle - \rangle_{\text{any}} s * t * I_R)$. \square

This holds even in a system without negation, *e.g.*, consider evaluating $\langle - \rangle(s * I_R)$ on the first state of the above sequence.

We postulate, though, that there is a syntactically defined class of formulæ, Φ_s , for which \models_s is sound and complete. We anticipate that deadlock freedom, \mathcal{D}_f^D , is in this class. Though the definition of this class is beyond us at the moment, we shall proceed to show that deadlock freedom is, indeed, sound and complete in the synchronous system.

5.4 Deadlock freedom equivalence in asynchronous and synchronous systems

Recall that the formula for deadlock freedom for Demos in our HML-type logic is:

$$\mathfrak{D}_f^D = \bigwedge_{i \in \mathbb{N} \cup \{\omega\}} [Act]_{\mathbf{any}}^i (\langle Act \rangle_{\mathbf{any}}^\top \vee I_P).$$

This formula will be equivalent in both transition systems if

- any state reachable in the synchronous system that is deadlocked is also reachable in the asynchronous system, and
- any deadlocked state reachable in the asynchronous system is also reachable in the synchronous system.

The first condition relates to soundness of deadlock detection in the synchronous system (dually, completeness of deadlock freedom); the second relates to completeness of deadlock detection in the synchronous system (dually, soundness of deadlock freedom).

More formally, for arbitrary $\kappa, \lambda \in \mathbb{N}_\omega$ there exist κ', λ' such that:

$$\begin{aligned} \text{If } \mathbf{R} \triangleleft \llbracket P \rrbracket \models_s \langle - \rangle_{\mathbf{any}}^\kappa (\neg(\langle - \rangle_{\mathbf{any}}^\top \vee I_P)) \text{ then} \\ \mathbf{R} \triangleleft \llbracket P \rrbracket \models_a \langle - \rangle_{\mathbf{any}}^{\kappa'} (\neg(\langle - \rangle_{\mathbf{any}}^\top \vee I_P)) \end{aligned} \quad (5.1)$$

$$\begin{aligned} \text{If } \mathbf{R} \triangleleft \llbracket P \rrbracket \models_a \langle - \rangle_{\mathbf{any}}^\lambda (\neg(\langle - \rangle_{\mathbf{any}}^\top \vee I_P)) \text{ then} \\ \mathbf{R} \triangleleft \llbracket P \rrbracket \models_s \langle - \rangle_{\mathbf{any}}^{\lambda'} (\neg(\langle - \rangle_{\mathbf{any}}^\top \vee I_P)) \end{aligned} \quad (5.2)$$

5.4.1 Soundness of deadlock detection

In order to show the soundness of deadlock detection, we shall firstly show that any state reachable in one transition by the synchronous system is reachable by the asynchronous system.

Lemma 5.4.1 There exists a finite number of transitions, n , such that:

$$\text{If } \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma}_s \mathbf{R}' \triangleleft \llbracket P' \rrbracket \text{ then } \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma}_a^n \mathbf{R}' \triangleleft \llbracket P' \rrbracket$$

PROOF *By induction on the structure of synchronous transition derivations.* The non-parallel cases are trivially true — in neither system do we use the parallel execution rules.

The interesting case is where we have a **try** command in parallel with some other process, as presented in §5.2.

The other cases for (SPAR) follow quite easily from the fact that we can ‘decompose’ the derivation using (SPAR) to form two asynchronous derivations. From the fact that the two processes do not interfere — otherwise, we would not be using (SPAR) — we can deduce that the resources necessary to take both actions will be available in the asynchronous case if they were in the synchronous case.

Returning to the `try` command case, assume that

$$\mathbf{R} \triangleleft \llbracket P \rrbracket \equiv \mathbf{S} \uplus \mathbf{S}' \uplus \mathbf{H} \triangleleft \llbracket \text{try } [e_1] \text{ then } P_1 \text{ etry } [e_2] \text{ then } P_2 \rrbracket \parallel Q \rrbracket.$$

If we are to derive a transition, we must have the sub-derivation $\mathbf{S} \uplus \mathbf{H} \triangleleft \llbracket \text{try } [e_1] \text{ then } P_2 \text{ etry } [e_2] \text{ then } P_2 \rrbracket \xrightarrow{\text{try}} \mathbf{H} \triangleleft \llbracket P' \rrbracket$, where $P' \equiv P_1$ or P_2 . We know that neither e_1 nor e_2 require any resources referenced by process Q by (*). Consequently, any transition from process Q will not use the resources required by e_1 or e_2 , and so they are not contained in \mathbf{S}' . Therefore, any resources that could be used by the `try` command, thereby affecting the path chosen, will not be hidden by inclusion in \mathbf{S}' . Thus adding the resources in \mathbf{S}' to the environment will not affect the transition, so we could equivalently derive:

$$\mathbf{S} \uplus \mathbf{S}' \uplus \mathbf{H} \triangleleft \llbracket \text{try } \dots \rrbracket \xrightarrow{\text{try}} \mathbf{H} \uplus \mathbf{S}' \triangleleft \llbracket P' \rrbracket$$

As this does not use a parallel execution rule — it uses (TRY), this holds for both systems. By induction, there exists an integer m such that

$$\mathbf{S}' \uplus \mathbf{H} \triangleleft \llbracket Q \rrbracket \xrightarrow[\sigma]{m}_a \mathbf{H} \triangleleft \llbracket Q' \rrbracket.$$

If, then, we apply the asynchronous parallel rule to first derive

$$\mathbf{S} \uplus \mathbf{S}' \uplus \mathbf{H} \triangleleft \llbracket \text{try } \dots \rrbracket \parallel Q \rrbracket \xrightarrow{\text{try}} \mathbf{S}' \uplus \mathbf{H} \triangleleft \llbracket P' \rrbracket \parallel Q \rrbracket,$$

we can apply the same sequence of m transitions on Q to derive the required state. \square

It follows easily by transfinite (ordinal) induction that any state reachable by a sequence of transitions in the synchronous system is reachable by a sequence of transitions in the asynchronous system.

Lemma 5.4.2 For any $\kappa \in \mathbb{N}_\omega$ there is a $\kappa' \in \mathbb{N}_\omega$ such that:

$$\text{If } \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow[\sigma]{\kappa}_s \mathbf{R}' \triangleleft \llbracket P' \rrbracket \text{ then } \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow[\sigma]{\kappa'}_a \mathbf{R}' \triangleleft \llbracket P' \rrbracket$$

\square

We now prove that if deadlock is detected according to the synchronous transition system then the state really can deadlock.

Theorem 5.4.3 (Soundness of deadlock detection (5.1))

$$\mathbf{R} \triangleleft \llbracket P \rrbracket \models_s \langle - \rangle_{\mathbf{any}}^\kappa (\neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P)) \implies \mathbf{R} \triangleleft \llbracket P \rrbracket \models_a \langle - \rangle_{\mathbf{any}}^\kappa (\neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P))$$

PROOF If we dissect the formula $\langle - \rangle_{\mathbf{any}}^\kappa \neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P)$, it expresses that $\neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P)$ holds in some state reachable from $\mathbf{R} \triangleleft \llbracket P \rrbracket$. Suppose, then, that the synchronous transition system leads to a state $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_s \neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P)$. By Lemma 5.4.2, we know that we can reach this state in the asynchronous system. We must, then, show that $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_a \neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P)$.

Firstly, notice that $\neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P)$ is logically equivalent to $\neg(\langle - \rangle_{\mathbf{any}} \top) \wedge \neg I_P$. Consequently, from the semantics of \wedge , we have $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_s \neg I_P$. As this is not

a modal formula, the differing transition system will not affect its satisfaction, so $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_a \neg I_P$. All that remains to be proven, then, is that $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_s \neg \langle - \rangle_{\mathbf{any}} \top$ implies $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_a \neg \langle - \rangle_{\mathbf{any}} \top$. Assuming $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_s \neg \langle - \rangle_{\mathbf{any}} \top$, there are two cases to consider: (*) holds and (*) does not hold.

If (*) does not hold, we are unable to derive a transition using the asynchronous parallel execution rules, and the synchronous parallel rule may not be used. Immediately, therefore, we are unable to derive any transition in the purely asynchronous system, so $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_a \neg \langle - \rangle_{\mathbf{any}} \top$.

If (*) holds, we know that there is at least one process with a non-interfering command at its head. Furthermore, we know that if this command is a `getX` (where X could be `R`, `B` or `S`) command, the resource is available. There is no reason, therefore, why we cannot make a transition from this state for every process indexed by $C'_P \cap A_P^{\mathbf{R}}$. We may deduce that this case does not arise: $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_a \neg \langle - \rangle_{\mathbf{any}} \top$ is never true if (*) holds, so we may vacuously conclude $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_s \neg \langle - \rangle_{\mathbf{any}} \top$. \square

5.4.2 Completeness of deadlock detection

In order to show that the synchronous system is complete with respect to the asynchronous system — that is, if we detect deadlock in the asynchronous system, we will in the synchronous system — we shall first show that if it is not possible to make a transition in the asynchronous system from an arbitrary state, it is not possible to make a transition in the synchronous system.

Lemma 5.4.4 For all states $\mathbf{R} \triangleleft \llbracket P \rrbracket$,

$$\mathbf{R} \triangleleft \llbracket P \rrbracket \models_a \neg \langle - \rangle_{\mathbf{any}} \top \implies \mathbf{R} \triangleleft \llbracket P \rrbracket \models_s \neg \langle - \rangle_{\mathbf{any}} \top$$

PROOF By induction on $\mathbf{R} \triangleleft \llbracket P \rrbracket$

Assume that $\mathbf{R} \triangleleft \llbracket P \rrbracket \models_a \neg \langle - \rangle_{\mathbf{any}} \top$. As usual, we shall consider only the case where P is the parallel composition of processes and (*) holds. Otherwise, we would use exactly the same transition system and the statement would be easily provable. Then there is no derivation concluding, for any state $\mathbf{R}' \triangleleft \llbracket P' \rrbracket$, $\mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma}_a \mathbf{R}' \triangleleft \llbracket P' \rrbracket$. Consequently, as P is the parallel composition of processes, there is no transition derivable for any of these. In the synchronous system, (*) holds, so we would require at least one derivation of a transition for a sub-process to derive a transition for P . By induction, we can make no transition for a sub-process in the synchronous system, so we can derive no transition for P in the synchronous system. ¹ \square

To prove the completeness of deadlock detection is quite complicated. We shall, therefore, present only an outline of the proof. Moreover, we shall consider only acquisition through `getR`; the cases for `getS` and `getB` are similar, as is acquisition through `try`.

¹An alternative proof would note that this case is vacuously true: if (*) holds, we must be able to make a transition, so we must be able to make a transition in the asynchronous system. We have given the full proof to illustrate how we could prove the property with a more relaxed condition (*).

Theorem 5.4.5 (Completeness of deadlock detection (5.2)) If there is a path to a deadlocked state in the asynchronous system, there is a path to a state in the synchronous system. That is, for arbitrary $\mathbf{R} \triangleleft \llbracket P \rrbracket$,

$$\mathbf{R} \triangleleft \llbracket P \rrbracket \models_a \langle - \rangle_{\mathbf{any}}^{\kappa} \neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P) \implies \mathbf{R} \triangleleft \llbracket P \rrbracket \models_s \langle - \rangle_{\mathbf{any}}^{\kappa} \neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P)$$

PROOF OUTLINE We start by assuming an arbitrary state $\mathbf{R}' \triangleleft \llbracket P' \rrbracket$ such that $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_a \neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P)$. This is equivalent to $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_a \neg(\langle - \rangle_{\mathbf{any}} \top \wedge \neg I_P)$. By the semantics of \wedge , $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_a \neg(\langle - \rangle_{\mathbf{any}} \top)$ and $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_a \neg I_P$. By Lemma 5.4.4, we see that $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_s \neg(\langle - \rangle_{\mathbf{any}} \top)$. Whether I_P holds in a state is independent of the transition system, so $\mathbf{R}' \triangleleft \llbracket P' \rrbracket \models_s I_P$.

Now, notice that the only rule we may fail to derive a transition by is (GETR); it is the only rule for which the side-condition² does not form a tautology with the other rules. As the state is deadlocked, the command at the head of every process, therefore, must be $\mathbf{getR}(r_i, 1)$. Moreover, as the side condition is not met (because we are not able to make a transition from this state), we do not have any resource r_i that is requested.

We may only use the rule (SPAR) to execute certain commands. It may not be used to pass any $\mathbf{getR}(r_i, 1)$ command where either r_i is not available or r_i is referenced by some other process. Beyond here, the asynchronous system will be used. Notice that (*) necessarily does not hold where we fail to derive a transition in the asynchronous system. We see, then, that

$$\begin{aligned} & \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma^k}_a \mathbf{R}' \triangleleft \llbracket \mathbf{getR}(r_1, 1) \parallel \dots \parallel \mathbf{getR}(r_n, 1) \rrbracket \\ \text{implies } & \mathbf{R} \triangleleft \llbracket P \rrbracket \xrightarrow{\sigma^\ell}_s \mathbf{R}' \triangleleft \llbracket \mathbf{getR}(r_1, 1) \parallel \dots \parallel \mathbf{getR}(r_n, 1) \rrbracket \end{aligned}$$

As $\neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P)$ is equivalent in the two systems, then,

$$\mathbf{R} \triangleleft \llbracket P \rrbracket \models_a \langle - \rangle_{\mathbf{any}}^{\kappa} \neg(\langle - \rangle_{\mathbf{any}} \top \vee I_P).$$

□

It is interesting that in the above proofs we have not used the fact that we do not use the (SPAR) rule to pass by $\mathbf{putR}(r, 1)$ commands where r is referenced by any other process. We could, therefore, remove this restriction and preserve the correctness of deadlock detection. Such a relaxation, though, would decrease the size of the class Φ_s of formulæ equivalent in the synchronous and asynchronous systems. The synchronous system would, for example, be incomplete in detecting whether any sub-process could ever have to wait on a resource.

It is also interesting to consider the behaviour of the synchronous parallel rule on the following program:

$$\begin{array}{c} \mathbf{getR}(r, 1) \\ \mathbf{getR}(s, 1) \end{array} \parallel \begin{array}{c} \mathbf{getR}(s, 1) \\ \mathbf{getR}(r, 1) \end{array} \parallel \begin{array}{c} \mathbf{while}(\top) \{ \\ \quad \mathbf{putB}(t, 1) \\ \quad \mathbf{getB}(t, 1) \\ \} \end{array}$$

The synchronous execution rule, assuming a resource state $\{r, s, t\}$, will continually ‘execute’ the third process. Though this is fine with respect to detecting

²The side-condition is implicit in the structure of the initial resource component

overall system deadlock, we will not be able to detect that the first two processes may become stuck: they will never pass their first commands, so they may both always make a transition.

Chapter 6

Conclusions

We have presented some of a theoretical framework capable of conducting analyses of discrete-event processes that interact through shared, discrete resources. This has involved taking the Logic of Bunched Implications and extending its semantics to be able to consider both processes and resources. In order to describe such processes, we have taken a simulation language, Demos, and defined its non-temporal action. Demos is an ideal language for this task: as well as being designed to model concurrently executing processes interacting through shared resources, its operational semantics are defined. This has allowed us to show that the transition system used by the logic that defines the action of processes is correct with respect to the non-temporal action of processes in Demos. We defined a kernel of Demos states capable of representing processes non-temporally.

6.1 Further work

A number of aspects remain to be considered.

Shared resource ¹ Demos programs do not allow us to represent shared resource: we cannot represent, for example, that processes P_1 and P_2 have access to a resource whilst P_3 does not.² Our process logic, similarly, is not able to make such a distinction, and it is not clear to see how we may extend it to do so without significantly modifying the logic.

Fairness and justice We have allowed processes to execute in an arbitrary manner (*e.g.*, in §5.4.2 we allowed the third process to be executed all the time, ignoring the (blocked) first two processes). Such biased execution may be considered to model improperly the actual characteristics of a system. Following from the discussion of fairness in [GPSS80], Manna and Pnueli in [MP83] suggest two characteristics of proper executions: *fairness* and *justice*. We adapt their definition to our notion of concurrently executing processes.

¹As suggested by D. J. Pym

²This could represent the resource having been taken by P_1 or P_2 which are “close to” each other, or the resource being generally unavailable to P_3 .

- An execution sequence is said to be *fair* if it terminates finitely, it is infinite and all finite processes terminate or otherwise cannot continue to execute, or it is infinite and all non-halting processes are executed infinitely often.
- An execution sequence is said to be *just* if whenever a process is always able to take a transition, that process will act infinitely many times.

[CES86] describes how fairness may be integrated into a model-checker.

Adopting such a fairness condition in the example of §5.4.2 above would prevent the third process executing through every transition, and we could deduce that the system can deadlock. It appears particularly important to require such a notion of fairness when we consider a parallel execution rule that may favour particular processes, such as the one we defined.

Non-syntactic resource ownership Our approach would be more general had we not relied upon the fact that at a given state it is possible to deduce what resources are held by a process from the syntax of the process code. For example, if the following were allowed in Demos: `while [putB(r,1)]` (or even `try[putB(r,1)]`), we would not be able to deduce from the syntax the resources held by the process. It would, therefore, be necessary to keep a record of the resources held by each sub-process in the transition system. Just as processes explicitly hold other processes that are synchronised on them, we would have

$$\mathbf{R} \triangleleft \left[\left[\frac{P}{\{r, s, t\}} \right] \right]$$

representing a process P holding resources r, s and t . We chose not to do this in order to reduce the burden of notation.

Variables Most programmers will recognise the usefulness of variables, and they are, indeed, allowed in Demos 2000.³ It would be a relatively simple matter to extend the transition system to take this into account. We would introduce a *variable state* component to take this into account. Transitions, then, would be of the form $\mathbf{R} \triangleleft \llbracket P \rrbracket s \xrightarrow{\sigma} \mathbf{R}' \triangleleft \llbracket P' \rrbracket s'$. Variable state would be simply a map from variable labels to variable values, $s : Var \rightarrow Float$. We chose not to introduce this as our account is theoretical, and whilst it is not difficult to introduce variables into an implementation, proofs can become less succinct with their inclusion.

The semantics of **PBI** presented in this document have been defined with proof-search in mind. If we were to write a model-checker, though, it is anticipated that it will not be able to check all formulæ; for example, to check the formula

$$\mathbf{R} \triangleleft \llbracket P \rrbracket \models I_R \text{ -* } \varphi$$

could involve us considering adding all syntactically correct processes being placed in composition with P . Clearly, as the set of all processes is infinite,

³That is, numeric variables *e.g.*, $x := 4; y := 5$, rather than resource-variables.

such a simple implementation would lead to a non-complete model-checker. It would be acceptable, however, to provide a model checker complete for only a fragment of the logic. It may also be necessary to consider a logic other than the general Hennessy-Milner-style logic used here, similar to CTL or modal- μ . Finally, we have not given a deduction system for our logic. Such a system should be proven correct (sound and complete) with respect to the semantics defined here.

To conclude, we have seen how we can extend the Logic of Bunched Implications to consider processes in a simple manner to achieve an expressive and clear way of formally reasoning about processes that interact on shared resources.

Bibliography

- [AB96] Andrea Asperti and Nadia Busi. Mobile petri nets. Technical Report UBLCS-96-10, University of Bologna, Department of Computer Science, May 1996.
- [BS01] Julian C. Bradfield and Colin Stirling. Modal logics and mu-calculi: an introduction. In J. Bergstra, A. Ponse, and S. Smolka, editors, *Handbook of Process Algebra*, pages 239–330. Elsevier, <http://www.dcs.ed.ac.uk/~jcb/>, 2001.
- [BTa] Graham Birtwistle and Chris Tofts. Getting demos models right. Part I: Practice.
- [BTb] Graham Birtwistle and Chris Tofts. Getting demos models right Part II: ... and theory.
- [BT93] Graham Birtwistle and Chris Tofts. Operational semantics for process-based simulation languages. Part 1: π demos. *Transactions of the Society for Computer Simulation*, 10(4):299 – 334, 1993.
- [BT01] Graham Birtwistle and Chris Tofts. *DEMOS 2000 — A Semantically Justified Simulation Language*, August 2001. <http://www.demos2k.org>.
- [CES86] E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 8(2):244–263, 1986.
- [FOL] The Free On-line Dictionary of Computing. <http://foldoc.doc.ic.ac.uk>.
- [Gir87] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- [GPSS80] Dov M. Gabbay, Amir Pnueli, Saharon Shelah, and Jonathan Stavi. On the temporal analysis of fairness. In *Proceedings of the 7th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 163–173. ACM Press, 1980.
- [HM80] Matthew C. B. Hennessy and Robin Milner. On observing nondeterminism and concurrency. In J. W. de Bakker and Jan van Leeuwen,

editors, *Automata, Languages and Programming, 7th Colloquium*, volume 85 of *Lecture Notes in Computer Science*, pages 299–309, Noordwijkerhout, The Netherlands, 14–18 July 1980. Springer-Verlag, New York, Berlin, Heidelberg.

- [HM85] Matthew C. B. Hennessy and Robin Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the ACM (JACM)*, 32(1):137–161, 1985.
- [Hoa85] C.A.R. Hoare. *Communicating Sequential Processes*. International Series in Computer Science. Prentice-Hall, 1985. Series editor: C. A. R. Hoare.
- [MBC⁺95] M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis. *Modelling with Generalized Stochastic Petri Nets*. Series in Parallel Computing. Wiley, 1995.
- [Mil89] Robin Milner. *Communication and Concurrency*. International Series in Computer Science. Prentice-Hall, 1989. Series editor: C. A. R. Hoare.
- [MP83] Zohar Manna and Amir Pnueli. How to cook a temporal proof system for your pet language. In *Proceedings of the 10th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, pages 141–154. ACM Press, 1983.
- [MS] Faron Moller and Perdita Stevens. Edinburgh Concurrency Workbench user manual (version 7.1). Available from <http://www.dcs.ed.ac.uk/home/cwb/>.
- [Mur89] Tadao Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, April 1989.
- [OED73] OED. *The Shorter Oxford English Dictionary*. Oxford University Press, third edition, 1973.
- [OP99] Peter W. O’Hearn and David J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–243, June 1999.
- [Pącz89] Paweł Pączkowski. Proving total correctness of concurrent programs without using auxiliary variables. Technical Report ECS-LFCS-89-100, University of Edinburgh, Laboratory for Foundations of Computer Science, November 1989.
- [Per58] Alan J. Perlis. Announcement. *Communications of the ACM*, 1(1):1, 1958. Letter from Saul Gorn.
- [Pit02] Andrew M. Pitts. Operational semantics and program equivalence. In G. Barthe, P. Dybjer, and J. Saraiva, editors, *Applied Semantics, Advanced Lectures*, volume 2395 of *Lecture Notes in Computer Science, Tutorial*, pages 378–412. Springer-Verlag, 2002. International Summer School, APPSEM 2000, Caminha, Portugal, September 9–15, 2000.

- [Plo81] Gordon D. Plotkin. A structural approach to operational semantics. Technical Report DAIMI FN-19, University of Aarhus, Denmark, Computer Science Department, 1981. <http://www.dcs.ed.ac.uk/home/gdp/>.
- [Plo03] Gordon D. Plotkin. The origins of structural operational semantics. *Journal of Functional and Logic Programming*, 2003.
- [POY02] David J. Pym, Peter W. O’Hearn, and Hongseok Yang. Possible worlds and resources: The semantics of **BI**. *Journal of Theoretical Computer Science (to appear)*, 2002.
- [Pym02] David J. Pym. *The Semantics and Proof Theory of the Logic of Bunched Implications*, volume 26 of *Applied Logic Series*. Kluwer Academic Publishers, Dordrecht, July 2002.
- [TB94] Chris Tofts and Graham Birtwistle. Denotational semantics for process-based simulation languages. Part I: π demos. Technical Report School of Computer Studies, University of Manchester, 1994.
- [TB95] Chris Tofts and Graham Birtwistle. Denotational semantics for process-based simulation languages. Part I: μ demos. Technical Report School of Computer Studies, University of Manchester, 1995.
- [TB01] Chris Tofts and Graham Birtwistle. Operational semantics of DEMOS 2000. Technical Report HPL-2001-263, Hewlett-Packard Laboratories, Graham Birtwistle, School of Computer Studies, University of Leeds, 2001.
- [Tof01] Chris Tofts. Compiling DEMOS 2000 into Petri nets. Technical Report HPL-2001-274, Hewlett-Packard Laboratories, 2001.