

THE FORMAL SPECIFICATION
OF A CONFERENCE ORGANIZING SYSTEM

T. Clement

Technical Monograph PRG-36
August 1983

Oxford University Computing Laboratory
Programming Research Group
8-11 Keble Road
Oxford OX1 3QD

© 1983. T. Clement

Oxford University Computing Laboratory
Programming Research Group
8-11 Keble Road
Oxford OX1 3QD

Address from September 1983:

Syracuse University
School of Computer and Information Science
313 Link Hall
Syracuse
New York 13210
U.S.A.

1 Introduction

1.1 The problem

This monograph arises from an exercise in information systems design sponsored by IFIP Working Group 8.1 (OleB2). This was intended to provide a means of comparing various design methods, proprietary and otherwise, by asking their exponents to prepare designs from a common specification to a level where the system could be produced by a typical applications programmer.

The problem as posed by IFIP is to produce an information system to support the organization of a conference. They explain that working conferences are normally run by a programme committee, responsible for the technical content, and an organizing committee, who are responsible for finance, publicity, and the hotel and conference facilities required by delegates. This clearly involves cooperation between the committees, and a substantial amount of shared information to be kept up to date.

The description of the existing organization identifies various activities which the committees perform, and with which assistance is sought. The programme committee invites contributions, and receives declarations of intention to submit papers. It also registers the receipt of papers, arranges their refereeing, collates the reports and selects papers, and finally draws up a timetable, with appropriate people chosen to chair each session.

Meanwhile, the organizing committee is responsible for issuing invitations to attend the conference, registering replies, and ensuring that the conference does not grow too large for its accommodation. It is their policy to invite authors of papers, as well as certain categories of IFIP official. They occasionally send several invitations to the same person if they come into more than one of these categories, and they would like the system

to spare them this embarrassment. They also prepare a final list of attendees. Other aspects of the work of the committees are not detailed in the exercise, and were not required as part of the submitted design.

1.2 The goals of formal specification

As producers of computing systems, we find ourselves with the insoluble problem of creating objects whose behaviour, which can be objectively tested, is to correspond to some partial description of a system which we do not fully understand. Our approach is to divide it into two: to find means of implementing precisely a precise specification, and means of agreeing a specification in precise form with our client before constructing a complete system.

Our attack on the first of these problems is through the application of formal notation a mathematical system in which every statement and all the means of manipulation are precisely defined, so that unambiguous specifications can be written, and proposed implementations can be checked for consistency with them, at least in principle.

The second problem remains insoluble, and is one common to all the professions: the lawyer knows the law, but not the client; the doctor knows the body but not the patient; we know about constructing computer systems, but not about the particular application. At the same time, the clients are seeking professional help only because they do not have the knowledge (of the law, of medicine, of computing) to help themselves. They do not understand what information they should give to help their case.

The doctor resolves the situation by posing questions of the patient which are inspired by an internal model of the body. The purpose of a systems design methodology for the designer is to provide a pattern for systems in general within which the nature of a particular system can be determined.

Such a pattern need not even be explicit. If it is, it may be expressed formally or informally. The more detailed it is, the more specific the questions may be, but the fewer the systems it will describe. We should therefore be prepared to deploy a variety of patterns, in the light of our knowledge of their areas of applicability. So far, we have developed them as the problems have arisen, so each is a prototype to be prepared and tested in the course of a specification.

The completed specification serves as a summary of the discussions which led to its creation. The clients should be confident that it represents what they want, and the designer that a system with the given properties exists, even under extra constraints such as a particular choice of machine, budget limitations, speed of response and so on. It can therefore be entered as part of the contract for the supply of the system.

1.3 The notation and the theory

Our starting point for the formal content of this specification is set theory: specifically the axiomatic set theory of [Abr1a18]. Its attractions for our purposes are the ease with which the terms can be given informal interpretations (naive set theory) and its expressive power, which is adequate: consistency of much of mathematics has been shown by exhibiting set theoretic models.

The whole of Abrial's set theory is expressed as a syntax for a language in which its terms and predicates are written (a few lines of BNF or similar notation), and eighteen axioms and a proof rule which together allow us to categorize some of the sentences of the language as theorems. The specification of an informally described system is the inverse process to that of informally interpreting a formal system (for example, giving the naive intuitive explanation of a term in axiomatic set theory). We write sentences which incorporate the words used in the informal description as identifiers.

and which express properties that these identifiers possess. So far as the theory is concerned, these properties are not (necessarily) deducible, but axiomatic the result is a new theory, containing the new identifiers as constants, in which new theorems can be deduced. It must be checked at two levels. The first check is for internal consistency: it is possible to introduce mutually inconsistent axioms and the result is a theory where everything is a theorem. The second is to relate the theory to the informal specification through the interpretation of identifiers. If the interpretation of a theorem corresponds to an untruth in the interpretation, then the theory is not consistent with reality. Conversely, if some property of the real world cannot be shown as a theorem then the theory is incomplete, and admits models which do not have the property. (A simple repair is to add the property as an axiom!)

This completely describes the *principle* of our approach to specification. The *practical* problem in carrying out this programme is that the formally stated properties must be amenable to expression and manipulation, and so the strings of symbols which express them must be short, and the proofs which demonstrate their properties must also be short. Within the framework of the language, the method of attack is through abstraction, the nesting of the final theory inside other theories of greater generality, which introduce new constants to be used in constructing sentences, and axioms concerning these constants which can be used in reasoning about them. Some of this reasoning will be done immediately, to establish theorems of general utility. This method is evident in [Abrial81], in the successive introduction of functions, finite sets, natural numbers, and sequences, each building on the previous ones.

We also make use of some simple notations for generating mathematical text. We define an identifier which expands into an arbitrary piece of mathematics (in many cases, of a form which could be associated with a quantifier) by

$\langle \text{name} \rangle \equiv \langle \text{text} \rangle$

or

<name> _____
|
| <text>
|
|_____

The text may be modified by systematic substitution of identifiers by general terms, so that given, for example

$$\text{XRANGE} \equiv t \in Y \rightarrow X$$

the predicate $\text{XRANGE}[Z \times Z/X]$ becomes $f \in Y \rightarrow Z \times Z$.

We shall save some effort by relaxing our rigour slightly. For example, we shall tend to move freely between the possible isomorphic bracketings of extended cross products. Informal notational conventions can also be convenient in particular, where auxiliary constants are introduced in definitions and theorems, we frequently find predicates of the form $\langle \text{identifier} \rangle \in \langle \text{set} \rangle$. Where the set is denoted by a single capital letter, for instance S , we shall assume that identifiers derived from the lower case letter, such as s, s', \bar{s}, s_1 and so on are members of it, unless otherwise indicated

2 A theoretical background

In the introduction, we identified two desires for a pattern for Information systems to structure our dialogue with our client, and for a theory of information systems to lighten our notation for the specification. We bring the two together here: our pattern (the running informal text) will be the interpretation of the theory (the indented formal text). We shall proceed in three steps, from a consideration of the nature of information systems, through a notation for the commands that act upon them, to an approach to designing a practical system

2.1 Extending the theory for defining systems behaviour

The activities of many organizations depend on knowledge about the state of the world for example, a manufacturing company may decide to order more of a particular item when the number in the warehouse drops below some level. The need for an information system (computer-based or not) arises when the direct observation of the state of the world is too difficult—for example, when the order clerk is too far from the warehouse to go and look on the appropriate shelf. The system being merely a simulacrum of part of the world, the obvious place to begin is in the specification of the part of the world which interests us, and its behaviour. This will determine the boundaries of the simulation, and we must be sure to make them explicit and be prepared to justify them: it seems unnecessary to record the colours of the warehouse shelves until we want to know if we have enough paint to repaint them.

The required observations are only one part of the specification. The state of the world is not static, but changing. In many cases, this is not a continuous change, but a discrete one: the stock level changes because an order is despatched. (On the other hand, the volume of brandy maturing

in a cask decreases with time through evaporation). In the discrete case, we can maintain the simulation of the state by simulating the event in the information system. There will normally be some time difference between the event and its simulation. Its magnitude will not be specified formally here, but, given that the simulation reflects only the state resulting from the events in the system, we shall record the pious hope that it shall be 'as small as reasonably possible'. Making the choice of events also sets boundaries to the system which we must justify. It may be reasonable to ignore the possibility of a warehouse fire, despite its effect on stock levels, but less reasonable to ignore deliveries.

Given the sets of events and observations, the specification must show how the observations are affected by the events, and the order in which events can occur. This ordering has an implicit effect on the possible observations: for example, a person has a single social security number because the 'issue a number' event happens only once for each person. Describing the effects of events may lead to the identifying of further events and observations to be included, and the final form that will be presented here is the result of many iterations of identifying events and specifying their effects. Doing this iteration at the specification stage is cheaper than waiting until implementation or test.

We need notations for presenting the temporal relationships between events and their effect on observations. We begin with the first, which we regard as establishing the patterns of behaviour in the world, and we shall restrict ourselves to the possible orderings in time. There are cases where this is not adequate: for example to record the constraint that goods paid for by cheque are not despatched for ten days, but this loss of expressive power is compensated for by a corresponding simplification of the theory.

The simplest behaviour is to do nothing before stopping. We give this special behaviour a name

SKIP \in BEHAVIOUR

Another fundamental behaviour of mainly technical interest does not even stop:

NULL \in BEHAVIOUR

The next most simple behaviour is to do one event once. We can identify the set of events, which we shall call E , with these behaviours.

$E \subset$ BEHAVIOUR

From these beginnings, we may then construct more complex behaviours. Given two behaviours, we may act like one *then* the other, one *or* the other, or one *and* the other in parallel. We introduce the three corresponding operators

$;$ \in BEHAVIOUR \times BEHAVIOUR \rightarrow BEHAVIOUR
 \square \in BEHAVIOUR \times BEHAVIOUR \rightarrow BEHAVIOUR
 \parallel \in BEHAVIOUR \times BEHAVIOUR \rightarrow BEHAVIOUR

Conversely, if we take a particular behaviour, then we expect it to have been created from the raw materials of E and SKIP or NULL by use of these operators, rather than existing *a priori*. In an equational specification, these would be the generators of the sort, and in Larch ([Gutt83]) special syntax is used for introducing a rule of computational induction over them. In our case, we can introduce this within the language

BEHAVIOUR \Leftarrow {SKIP, NULL} \cup E \cup
(; \cup \square \cup \parallel) (BEHAVIOUR \times BEHAVIOUR)

We can then define functions yielding the events which comprise any behaviour and the orders in which they can occur by recursion on the structure of that behaviour.

We may define the alphabet of a behaviour as the set of events which it involves, or formally as

$$\alpha \in \text{BEHAVIOUR} \rightarrow \mathcal{P}(E)$$

$$\alpha \text{ SKIP} = \alpha \text{ NULL} = \Phi_E \quad (\text{A1})$$

$$e \in E \Rightarrow \alpha e = \{e\} \quad (\text{A2})$$

$$\alpha(b_1; b_2) = \alpha(b_1 \square b_2) = \alpha(b_1 \parallel b_2) = \alpha b_1 \cup \alpha b_2 \quad (\text{A3})$$

It will be helpful to formally define the pairs of behaviours which have no events in common

$$\text{Independent} \in \text{BEHAVIOUR} \leftrightarrow \text{BEHAVIOUR}$$

$$\text{Independent}(b_1, b_2) \Leftrightarrow \text{Disjoint}(\alpha b_1, \alpha b_2)$$

We can express the orderings of events as the possible sequences in which they can occur. Those sequences after which no more events may occur are defined by

$$\text{finals} \in \text{BEHAVIOUR} \rightarrow \mathcal{P}(\text{seq}(E))$$

$$\text{finals SKIP} = \{\langle \rangle\}$$

$$\text{finals NULL} = \Phi_{\text{seq}(E)}$$

$$\text{finals } e = \{\langle e \rangle\}$$

$$\text{finals}(b_1; b_2) = \{s*t \mid s: \text{finals } b_1; t: \text{finals } b_2\}$$

$$\text{finals}(b_1 \square b_2) = \text{finals } b_1 \cup \text{finals } b_2$$

$$\begin{aligned} \text{finals}(b_1 \parallel b_2) = \{ & t: \text{seq}(E) \mid \alpha t \subset \alpha(b_1 \parallel b_2) \ \& \\ & t \not\vdash \alpha b_1 \in \text{finals } b_1 \ \& \\ & t \not\vdash \alpha b_2 \in \text{finals } b_2 \} \end{aligned}$$

$\text{traces} \in \text{BEHAVIOUR} \rightarrow \mathcal{P}(\text{seq}(E))$

$\text{traces SKIP} = \text{traces NULL} = \{\langle \rangle\}$

$\text{traces } e = \{\langle \rangle, \langle e \rangle\}$

$\text{traces}(b_1; b_2) =$

$\text{traces } b_1 \cup \{s^*t \mid s: \text{finals } b_1; t: \text{traces } b_2\}$

$\text{traces}(b_1 \parallel b_2) = \text{traces}(b_1) \cup \text{traces}(b_2)$

$\text{traces}(b_1 \parallel b_2) = \{t: \text{seq}(E) \mid \alpha t \subset \alpha(b_1 \parallel b_2) \ \&$
 $t \downarrow \alpha b_1 \in \text{traces } b_1 \ \&$
 $t \downarrow \alpha b_2 \in \text{traces } b_2 \}$

and from this we can derive the relationship between sequences of events and the events which can follow them in a behaviour, and the relationship between behaviours and the events which can start them.

$\underline{\text{accepts}} \in \text{BEHAVIOUR} \rightarrow (\text{seq}(E) \leftrightarrow E)$

$t(\underline{\text{accepts}} \ b) \ e \Leftrightarrow t^*(e) \in \text{traces } b$

$\underline{\text{starts}} \in E \leftrightarrow \text{BEHAVIOUR}$

$e \underline{\text{starts}} \ b \Leftrightarrow \langle \rangle (\underline{\text{accepts}} \ b) \ e$

Are these definitions correct? They are clearly internally consistent, since we can find models for BEHAVIOUR, but do they have the properties we expect when we talk of these types of combination of behaviours? They have been carefully constructed to be reasonable, but we should like to have more confidence, which we can gain by establishing a few required properties as theorems

$t \in \text{traces } b \vdash \alpha t \subset \alpha b$

$\vdash \text{traces } b = \text{PRFX}(\text{traces } b)$

$\vdash \text{finals } b \subset \text{traces } b$

$\vdash t(\underline{\text{accepts}} \ b) \ e \Rightarrow t \in \text{traces } b$

$t \in \text{traces } b; t \neq \langle \rangle \vdash \text{hd } t \underline{\text{starts}} \ b$

Proofs of these theorems depend on the computational induction principle enunciated earlier. In general, since we can only construct behaviours which stop after a finite number of events, we might also expect that the traces could be derived as the prefixes of the finals. The behaviour $(e_1; e_2) \parallel (e_2; e_1)$ is an obvious counterexample: in some cases of parallel combination, behaviours may disagree with each other. We shall be more interested in behaviours which do terminate at expected points, which we shall call consistent

Consistent \subset BEHAVIOUR

Consistent $b \Leftrightarrow \text{traces } b = \text{PRFX}(\text{finals } b)$

Consistency is preserved by ; and \square , and inconsistency by ; and \parallel . NULL is not consistent, so we do not expect to use it much in defining other behaviours

In our set theory, all sets, including BEHAVIOUR, have an equality defined amongst their elements. We can be sure that there are many behaviours in the set, by looking at the functions α , finals, and traces defined over it. For example, $e_1 \neq e_1; e_2$, since their alphabets are different. On the other hand, although our intuition tells us that b and $b \square b$ should be the same behaviour, our theory is not sufficiently developed to prove this. As usual we could repair the defect by introducing axioms to state this and other basic algebraic properties as equalities. We should then have to show that they did not equate behaviours detectably different in their alphabet or traces. The alternative is to define equivalence under these functions, and show that the desired properties are theorems for this relationship. This can be made more attractive technically by borrowing the idea of reduction from Larch, which we express by making BEHAVIOUR isomorphic to its equivalence classes under the functions.

$\forall b_1, b_2: \text{BEHAVIOUR.}$

$\alpha b_1 = \alpha b_2 \ \& \ \text{finals } b_1 = \text{finals } b_2 \ \& \ \text{traces } b_1 = \text{traces } b_2 \Rightarrow b_1 = b_2$

We can now use standard equality theory rather than define a new relation, and take as a natural (semantic) model of BEHAVIOUR a subset of $\mathcal{P}(E) \times \mathcal{P}(\text{seq}(E)) \times \mathcal{P}(\text{seq}(E))$, defined in part by the interrelationships of the theorems above. The theorems below establish the standard algebraic properties of the operators that we expect to hold under interpretation. Such properties are a useful source of conjectures whose establishment gives grounds for confidence in the definitions made. They are also sufficient on their own to perform useful proofs.

$$\vdash (b_1; b_2); b_3 = b_1; (b_2; b_3) \quad (\text{C1})$$

$$b; \text{SKIP} = \text{SKIP}; b = b \quad (\text{C2})$$

$$b \square b = b \quad (\text{C3})$$

$$b_1 \square b_2 = b_2 \square b_1 \quad (\text{C4})$$

$$(b_1 \square b_2) \square b_3 = b_1 \square (b_2 \square b_3) \quad (\text{C5})$$

$$b \square \text{NULL} = b \quad (\text{C6})$$

$$b \parallel b = b \quad (\text{C7})$$

$$b_1 \parallel b_2 = b_2 \parallel b_1 \quad (\text{C8})$$

$$(b_1 \parallel b_2) \parallel b_3 = b_1 \parallel (b_2 \parallel b_3) \quad (\text{C9})$$

$$b \parallel \text{SKIP} = \text{SKIP} \parallel b = b \quad (\text{C10})$$

$$b; \text{NULL} = b \parallel \text{NULL} \quad (\text{C11})$$

$$(b_1 \square b_2); b_3 = (b_1; b_3) \square (b_2; b_3) \quad (\text{C12})$$

$$b_1; (b_2 \square b_3) = (b_1; b_2) \square (b_1; b_3) \quad (\text{C13})$$

Since \square and \parallel are both commutative and associative, and have units, we may define

$$\begin{aligned} \square &\in \mathcal{F}(\text{BEHAVIOUR}) \rightarrow \text{BEHAVIOUR} \\ \square &= \text{continuation}(\square, \text{NULL}) \end{aligned}$$

$$\begin{aligned} \parallel &\in \mathcal{F}(\text{BEHAVIOUR}) \rightarrow \text{BEHAVIOUR} \\ \parallel &= \text{continuation}(\parallel, \text{SKIP}) \end{aligned}$$

where continuation, as defined in appendix A, places the dyadic operator between elements of the set. Another derived operator is $?$, defined by

$$\begin{aligned} ? &\in \text{BEHAVIOUR} \rightarrow \text{BEHAVIOUR} \\ ?b &= b \sqcup \text{SKIP} \end{aligned}$$

The behaviour $?b$ represents a choice in the world between behaving like b and doing nothing.

It is convenient when describing and reasoning about behaviours to be able to deal with partial definitions, or constraints, as well as complete definitions. If one behaviour constrains another, then we expect it to control the possible orderings of the events which it comprises, whilst leaving any others unconstrained.

$$\begin{aligned} \underline{\text{sat}} &\in \text{BEHAVIOUR} \leftrightarrow \text{BEHAVIOUR} \\ b_1 \underline{\text{sat}} b_2 &\Leftrightarrow \forall t: \text{traces } b_1, t \downarrow \alpha b_2 \in \text{traces } b_2 \ \& \\ &\quad \forall t: \text{finals } b_1, t \downarrow \alpha b_2 \in \text{finals } b_2 \end{aligned}$$

This clearly has some resemblance to the definition of \parallel . Intuitively, the behaviour is to be completed by describing the rest of it to run in parallel. Formally,

$$\begin{aligned} \vdash b_1 \underline{\text{sat}} b_2 \ \&\ \alpha b_2 \subset \alpha b_1 &\Leftrightarrow b_1 = b_1 \parallel b_2 \\ &\Leftrightarrow \exists b: \text{BEHAVIOUR}, b_1 = b_2 \parallel b \end{aligned}$$

and it follows from the algebraic properties of \parallel that this restriction of $\underline{\text{sat}}$ defines a semilattice on BEHAVIOUR , with \parallel as the join operation.

$$\underline{\text{sat}}_S = \underline{\text{sat}} \cap \{(b, b') \mid b, b': \text{BEHAVIOUR} \ \&\ \alpha b' \subset \alpha b\}$$

$$\vdash b \text{ sat}_S b \quad (S1)$$

$$b_1 \text{ sat}_S b_2 \ \& \ b_2 \text{ sat}_S b_1 \Rightarrow b_1 = b_2 \quad (S2)$$

$$b_1 \text{ sat}_S b_2 \ \& \ b_2 \text{ sat}_S b_3 \Rightarrow b_1 \text{ sat}_S b_3 \quad (S3)$$

$$b \text{ sat}_S \text{SKIP} \quad (S4)$$

Theorems S1 and S4 clearly apply also to sat. The fundamental theorem for reasoning about behaviours follows directly from the definition.

$$\vdash s \subset \alpha b_2 \ \& \ b_1 \text{ sat} b_2 \Rightarrow \{t \downarrow S \mid t: \text{traces } b_1\} \subset \{t \downarrow S \mid t: \text{traces } b_2\}$$

The following theorems allow us to establish satisfaction of one behaviour by another in each case, if the alphabet restrictions are met. sat_S can be substituted for sat and the transitive property used to construct chains of satisfaction.

$$\vdash b_1 \parallel b_2 \text{ sat}_S b_2 \quad (S5)$$

$$\text{Independent}(b_1, b_2) \Rightarrow b_1 \text{ sat} b_2 \ \& \ b_2 \text{ sat} b_1 \quad (S6)$$

$$\text{Independent}(b_1, b) \ \& \ b_1 \text{ sat} b_2 \Rightarrow b_1 \text{ sat} b_2 \sqcap b \quad (S7)$$

$$b_1 \text{ sat} b \ \& \ b_2 \text{ sat} b \Rightarrow b_1 \sqcap b_2 \text{ sat} b \quad (S8)$$

$$\text{Independent}(b_1, b_2) \Rightarrow b_1; b_2 \text{ sat}_S b_1 \parallel b_2 \quad (S9)$$

$$\text{Independent}(b_1, b_4) \ \& \ \text{Independent}(b_2, b_3) \ \&$$

$$b_1 \text{ sat} b_3 \ \& \ b_2 \text{ sat} b_4 \Rightarrow b_1; b_2 \text{ sat} b_3; b_4 \quad (S10)$$

Let us now turn to the effect of the events on the observations. We may imagine stationing an observer where he can see the events as they occur in the part of the world whose behaviour we have described. (Such an observer makes a suitable operator for the information system). He may choose to record various levels of information: he may count the events that occur, he may record which events have occurred, he may record the order in which they occur; or he may even record the times at which they happen. From these basic observations of the world, we may deduce others.

and the richer the information recorded, the more we shall be able to deduce subsequently. When we developed the theory of behaviours, we decided to restrict ourselves to specifying only relative orderings in time, with the understanding that we could not express some constraints in that framework. We shall take an analogous decision here, and deal only with observations defined by the order of events, which we can formalize as functions in $\text{seq}(E) \rightarrow X$ for some X . This is adequate for the IFIP specification presented here, but could not be used, for example, for deriving the interest payable on a savings account. A specification of the IFIP problem based on a full treatment of time is given in [Gustafsson81] and makes an interesting contrast with our approach.

Since any trace of a behaviour may occur, the domain of the observation should include them all. Conversely, as no other sequences will be observed, the domain need be no larger than that. We shall define our observations over precisely the observable sequences of events.

2.2 Defining the alphabet of events

So far we have treated the events in the system as an arbitrary, unstructured set. If we examine a typical information system application, however, we find that we have the 'same' event occurring to different members of some set of entities, more correctly, that the events form families whose members are distinguished by the value of one or more parameters. For example, in the conference system we have the family of events corresponding to the submission of papers and parameterized by the paper and its author, the family of events for refereeing parameterized by paper and referee, and so on.

We should like to use the natural notation $f P$, where f is the family and P the parameter value, to denote an event. We also need to capture the

intuition that an event is identified uniquely by f and P . We can formalize this directly by giving axiom schemas for this type of term.

$$f \in \text{FAMILY} \Rightarrow f P \in E$$

$$f_1, f_2 \in \text{FAMILY} \ \& \ f_1 P_1 = f_2 P_2 \Rightarrow f_1 = f_2 \ \& \ P_1 = P_2$$

In general, we expect the parameter terms to be tuples, so that we see events of the form $\text{Referee}(p, r)$. Terms of the form f or P are to be taken as parameters of the schema.

This extension of the axiom schemas of our set theory to enable us to reason with a greater variety of terms can be compared with the 'abstract syntax' approach of [Sufrin83]. In that treatment, each f in FAMILY would be associated with an injection from its parameter set into E , and the disjointness of the ranges of the injections asserted. To lighten the load of stating these properties, a new syntactic form is introduced, and the associated proof rules given informally. The proof rules for events we have introduced above are then theorems of the construction. We can complete the 'isomorphism' between the two approaches by defining injections from the parameter sets for each family

$$\text{rep} \in \text{FAMILY} \rightarrow (X \rightarrow E)$$

$$(\text{rep } f) x = f x$$

This can be extended to give the set of events corresponding to a set of parameters in a family

$$\text{from} \in \text{FAMILY} \times \mathcal{P}(X) \rightarrow \mathcal{P}(E)$$

$$f \text{ from } S = \{f s \mid s: S\}$$

so we can talk of $\text{Submit } \text{from} (P \times A)$ and so on.

2.3 Backtracking

Once we have defined the behaviour of our world, we know the possible sequences of events which may be observed to happen in it. A corresponding sequence of events should be entered into the system which imitates the world. We can make the system reject sequences of events which could not possibly happen, but we cannot prevent the entry of sequences which might have happened but didn't. We can make provision for recovery from the simplest form of error - accidentally entering the wrong event - by providing a new event, OOPS, with the intention (that we shall make formal later) that entering OOPS 'undoes' the effect of the previous event. This new event is distinct from the events in the world.

$$\text{OOPS} \in E$$

$$E_{\text{OOPS}} = E \cup \{\text{OOPS}\}$$

A system with capacity for undoing erroneously entered events must accept sequences of events from the set E_{OOPS} . Given such a sequence, we can define the corresponding sequence of desired events, where the erroneous entries have been cancelled.

$$\text{edit} \in \text{seq}(E_{\text{OOPS}}) \rightarrow \text{seq}(E)$$

$$\text{edit}(\langle \rangle) = \langle \rangle$$

$$\text{edit}(\langle \text{OOPS} \rangle) = \langle \rangle$$

$$\forall t: \text{seq}(E_{\text{OOPS}}).$$

$$e \in E \Rightarrow \text{edit}(t * \langle e \rangle) = \text{edit}(t) * \langle e \rangle \ \&$$

$$e \in E_{\text{OOPS}} \Rightarrow \text{edit}(t * \langle e, \text{OOPS} \rangle) = \text{edit}(t)$$

The rather unconventional right-recursive definition is much more convenient in definition and proof than the left-recursive form. Examination of the definition will show that we have defined the behaviour in cases left

unmentioned in the informal description two consecutive OOPSes cancel, leaving the event before 'done', and an initial OOPS will be ignored. The former is a reasonably natural choice (and certainly more efficiently implementable), while the latter is made for the technical convenience of having edit total. As one would expect, edit leaves a sequence of events not containing OOPS unchanged. (The proof is by induction, trivially).

$$\vdash \forall t: \text{seq}(E). \text{edit}(t) = t$$

We should like to characterize the set of sequences of events with corrections which should be accepted by a system with given behaviour. Any such sequence, when edited, must give a legitimate trace of the behaviour. The same property must also hold for every prefix so that impossible events are rejected immediately.

$$\begin{aligned} \text{backtracks} &\in \text{BEHAVIOUR} \rightarrow \mathcal{P}(\text{seq}(E_{\text{OOPS}})) \\ \text{backtracks}(b)(t) &\Leftrightarrow \\ &\forall t': \text{PRFX}(t). \text{edit}(t') \in \text{traces}(b) \end{aligned}$$

There are certain properties that we expect of backtracks. First, where an operator makes no mistakes, the original sequence of events should be valid: that is, every member of traces(b) is in backtracks(b). Second, OOPS should always be allowed as an event. Finally, as events are entered one at a time, we expect the set to be prefix closed. It might be possible to use these properties to define backtracks. Instead, we must establish them as theorems and thus justify our more constructive approach. The first is trivial, given the theorem on edit and the prefix closure of traces; the second follows from case analysis on the end of the trace, and the third is proved by induction.

$$\begin{aligned} b &\in \text{BEHAVIOUR} \\ &\vdash \text{traces}(b) \subseteq \text{backtracks}(b) \\ &\vdash \forall t: \text{backtracks}(b). t * \langle \text{OOPS} \rangle \in \text{backtracks}(b) \\ &\vdash \forall t: \text{backtracks}(b). \\ &\quad \forall t': \text{PRFX}(t). t' \in \text{backtracks}(b) \end{aligned}$$

The observations on the system must now be defined over $\text{backtracks}(b)$, rather than $\text{traces}(b)$. It is possible to define functions whose results depend on the presence of OOPS events: we could, for instance, just count the number of errors made by the operator of the system. On the other hand, we expect the observations for which the system was designed to depend for their results only on the events correctly entered. We can characterize such functions as *forgiving*.

$$\begin{aligned} \text{Forgiving}(S) = & \\ & \{r: \text{seq}(E_{\text{OOPS}}) \leftrightarrow S \mid \\ & \forall t_1, t_2: \text{seq}(E_{\text{OOPS}}); s: S. \\ & \text{edit}(t_1) = \text{edit}(t_2) \Rightarrow r(t_1, s) \Leftrightarrow r(t_2, s)\} \end{aligned}$$

Given a sequence of events within the backtracks of a behaviour, we expect the possible next events to be those acceptable to the edited trace within $\text{traces}(b)$. If we define a new relation to include OOPS

$$\begin{aligned} \underline{\text{accepts}}_B \in \text{BEHAVIOUR} \rightarrow (\text{seq}(E_{\text{OOPS}}) \leftrightarrow E_{\text{OOPS}}) \\ t (\underline{\text{accepts}}_B b) e \Leftrightarrow t * \langle e \rangle \in \text{backtracks}(b) \end{aligned}$$

we can state this as

$$b \in \text{BEHAVIOUR} \vdash \text{Forgiving}(E_{\text{OOPS}})(\underline{\text{accepts}}_B b)$$

The proof follows from the definitions of the terms and case analysis on the accepted event.

The function edit defines equivalence classes over $\text{seq}(E_{\text{OOPS}})$, each of which contains a unique member of $\text{seq}(E)$. We can systematically define relations on the enhanced traces from those on the traces of events in the world by making use of this.

$$\begin{aligned} \text{enhancement} \in (\text{seq}(E) \leftrightarrow X) \rightarrow (\text{seq}(E_{\text{OOPS}}) \leftrightarrow X) \\ \text{enhancement}(r) = \text{edit}; r \end{aligned}$$

Clearly, by construction, since forgivingness requires only that the equivalence classes of edit be respected by a relation.

$$r \in (\text{seq}(E) \leftrightarrow S) \vdash \text{Forgiving}(S)(\text{enhancement}(r))$$

We can therefore define observations for the system in the natural way over the sequences of real events which comprise the traces, and characterize the functions over the backtracks that we actually require by enhancement

3 The specification

We have now completed the theoretical background, and have some idea of the shape the system will have. To make this pattern more specific, we shall adopt the image of an 'intelligent filing cabinet', which allows the recording of particular events as they occur, and the retrieval of information for guiding the activities of the committees, but does not attempt to automate any of their activities.

3.1 The behaviour of the system

We should begin by giving the behaviour a name. This provides a new constant in our theory whose properties we may then state

CRIS ϵ BEHAVIOUR

We have developed mechanisms for constructing a behaviour from its subcomponents. A natural division in the informal description is that between the programme committee, who are responsible for the papers submitted, and the organizing committee, who must take care of the delegates. We shall begin with the events concerning the papers once they have arrived

A paper comes into the system by being submitted by its author. It may then be sent for consideration by one or more referees. Normally, this will result in a verdict being returned, but in exceptional circumstances we shall allow the refereeing process to be truncated at the request of one or other party. When all the referees' reports are received, a decision can be made and the author informed. An accepted paper must be assigned to a conference session. We also provide for the possible acknowledgement of the receipt of a paper.

To formalize this, let us represent the papers by the set P, referees by the set R, their verdicts by V, authors by A, and times for presentation by T. We shall assume these (and the rest of the sets introduced) to be finite, so that all our combinators on behaviours can be used freely. As usual, lower case single letter identifiers and their variants belong to the upper case sets. We shall allow for two possible decisions, acceptance and refusal:

{Submit, Referee, Report, Cancel, Decide,
Inform, Schedule, Acknowledge} \subseteq FAMILY

Dec = {ACCEPTED, REJECTED}

From the viewpoint of the referee, some of the papers will arrive, and he will either referee them or return them. If he does neither, then he may be told not to bother. He cannot referee the same paper twice.

```

REF _____
|
|  Referee(r, p);
|  ( [] {Report(r, p, v) | v: V} [] Cancel(r, p)
|  _____

```

$\forall r: R; p: P. \text{CRIS } \underline{\text{sat}}_g \text{ ?REF}$

If this is the only restriction on refereeing, then an arbitrary number of papers may be sent to any one referee. This is in accord with our declared policy of leaving decisions to the users: we shall expect to make the workload of the referees available to the programme committee to help them make a selection (They should also be aware of other constraints beyond the scope of this system, such as, for example, a referee being a spouse of the author)

Note that there is a difference between undoing the selection of a referee by using OOPS and cancelling the appointment with Cancel. In the latter

case. the referee can not subsequently be reappointed to the paper (as one might expect) Cancellation can also take place at any time after the appointment. whereas the use of OOPS applies only to the immediately preceding event. whatever that may be

The author. having submitted a paper. may get an acknowledgement. and will be notified of the programme committee's decision.

AUTHORP _____

```
?Acknowledge(p, a);  
[] {Inform(p, a, d) | d: Dec}
```

$\forall p: P; a: A. \text{CRIS } \underline{\text{sat}}_s \text{ ?AUTHORP}$

We can then describe how the programme committee must handle a paper
They may choose to acknowledge the receipt of a paper at least until a decision is reached. Subsequently, they must inform the author and schedule the paper if necessary, where scheduling requires the choice of a time slot for the paper to be presented

DEC _____

```
Decide(p, ACCEPTED);  
(Inform(p, a, ACCEPTED) ||  
 [] {Schedule(p, t) | t: T})  
[] Decide(p, REJECTED); Inform(p, a, REJECTED)
```

PAPER

$$\square ((\text{Submit}(p, a);$$
$$\quad ?\text{Acknowledge}(p, a) \parallel (||\{?\text{REF} \mid r: R\});$$
$$\quad \text{DEC})$$
$$\quad \mid a: A)$$

$\forall p: P. \text{CRIS } \underline{\text{sat}}_s \text{ ?PAPER}$

We can establish that

$$\vdash \forall p: P; r: R. \text{PAPER } \underline{\text{sat}}_s \text{ ?REF}$$
$$\vdash \forall a: A; p: P. \text{PAPER } \underline{\text{sat}}_s \text{ AUTHORP}$$

by using the satisfaction theorems of the previous chapter as lemmas. Hence, given the transitivity of $\underline{\text{sat}}_s$, we can be sure that any system which satisfies ?PAPER will meet the rest of the specification so far.

In scheduling the papers, we have one more, rather unexpected, 'actor' in our behaviour: the conference time slots. These must be allocated to at most one paper.

$\forall t: T. \text{CRIS } \underline{\text{sat}}_s \text{ ?} \square (\text{Schedule}(p, t) \mid p: P)$

The programme committee may also take suggestions for potential authors, issue a call for papers, and then register letters of intent. We shall assume that these are not mandatory, and that they will not happen once a paper has actually been submitted

{Suggest_A, Call, Intend} c FAMILY

PRESUBMIT _____

?(Suggest_A(a); Call(a)); ?Intention(a)

Vp: P; a: A. CRIS sat_s ?(PRESUBMIT; AUTHORP)

The organizing committee must coordinate the invitations to delegates. These delegates may be suggested, in which case they will certainly be sent an invitation. They may then accept the invitation, and will receive a reply from the organizers. We shall also deal with the self-invited, returning the answer 'no'. Formally, we use D for the set of delegates.

{Suggest_D, Invite, Accept, Reply} c FAMILY

DEL _____

Invite(d);
?(Accept(d); \square {Reply(d, d') | d': Decision})

DELEGATE _____

(Suggest_D(d); DEL)
 \square DEL
 \square (Accept(d); Reply(d, REJECTED))

Vd: D. CRIS sat_s ?DELEGATE

Authors, as we have established, submit papers and expect replies from the programme committee. They are also entitled to invitations as a result of submitting a paper. Their complete behaviour as far as the system is concerned is thus

```
AUTHOR _____
|
| PRESUBMIT; Submit(p, a); (DEL[a/d] || AUTHORP)
|_____
```

$\forall p: P; a: A. \text{CRIS } \underline{\text{sat}}_s \text{ ?AUTHOR}$

We can readily show that

$\vdash \text{AUTHOR } \underline{\text{sat}}_s \text{ PRESUBMIT; AUTHORP}$

and hence that satisfaction of the requirements of AUTHOR guarantees satisfaction of those of AUTHORP.

We have now defined the behaviour of the world as seen from various viewpoints. We want all these activities to go on, as far as possible, in parallel we can also say at this stage that there are no other events. We can therefore define CRIS as the parallel combination of all the behaviours it must satisfy

$$\text{CRIS} = \{ \{ \{ \text{?PAPER } \mid p: P \} \parallel \{ \{ \text{?DELEGATE } \mid d: D \} \parallel \{ \{ \text{?AUTHOR } \mid p: P; a: A \} \parallel \{ \{ \text{? } \square \{ \text{Schedule}(p, t) \mid p: P \} \mid t: T \} \} \} \} \}$$

3.2 The observations

Having defined the behaviour in which we are interested, we can now discuss some observations that we should like to make. It is convenient to name the results of applying some of the functions on behaviours to CRIS.

```
TRACES      = traces(CRIS)
BTRACKS     = backtracks(CRIS)
accepts    = acceptsg(CRIS)
```

We intend to define some functions from BTRACKS to provide the substitute observations which the user needs, but we must first decide what observations are required. Our ultimate authority is the client, but as usual we should like some pattern for shaping our dialogue with him.

The observations we can define are bounded above by the information which BTRACKS conveys. The restrictions implied by the theory itself were discussed as it was developed: by choosing specific events associated with particular information we have limited ourselves further. It is no use trying to ask about consequences of events which we have decided not to record, and if we find that we need to do so, we must revise our ideas of the events in the world. Completeness is confirmed by exhibiting a definition of the observation.

Let us consider the observations that the user might want to make of the state of the refereeing. These might include the referees still to produce reports, the papers with reports outstanding, papers which need referees appointed, the referees' reports on a paper, and so on, and indeed, the actual observations chosen may be constrained by aspects of the design not yet settled. For example, if a referee's verdict is expressed on a scale out of ten, then we may conveniently show all the verdicts for a particular paper in one screenful, whereas if they are more involved they may need to be taken one at a time. At this level of abstraction of specification,

therefore, we should aim to find a small set of observations from which we can expect to derive others, and delay the determination of the actual set to be presented to the user until later

There are two properties of CRIS that help us in this. In general, the observations define the effect, or "meaning", of a sequence of events. Our intuition in this case, however, is that individual (real) events have a meaning independent of the rest of the sequence in which they appear (although these sequences are constrained by the behaviour to be meaningful collectively). For example, $\text{Submit}(p, a)$ means that paper p has been received from author a , and it occurs only in sequences where no other author has submitted the paper. It therefore seems that knowing the set of events which has occurred might be enough to deduce any further interesting observations.

Further support for this idea comes from the second property of CRIS, that events happen once if at all: to know that they have occurred is to know how often they have occurred. Further, if the behaviour tells us that it is one of a sequence of events, then we know the order in which those events occurred, and if it is part of a set of independent alternatives, that the events of the other alternatives will not occur. Only where events may occur in parallel is there any doubt as to the actual order of occurrence, but parallel combination in CRIS is used only with independent threads of behaviour: we are unlikely to want to know if a given paper was received before some person was invited to the conference. Thus, given an observation from which we can determine if a given event has occurred, we expect to be able to define almost any (forgiving) function definable from the traces themselves, the remainder being 'uninteresting'.

For sequences of real events, this observation can be just the alphabet of the trace. For the more general case of traces with corrections, we have

$$\begin{aligned} \text{obs} &\in \text{BTRACKS} \rightarrow \mathcal{P}(E) \\ \text{obs} &= \text{enhancement}(\alpha) \end{aligned}$$

$\vdash \text{ran obs} = \{\alpha t \mid t: \text{TRACES}\}$

which, applying the theorem of alphabets of behaviours, gives

$\vdash \text{ran obs} \subset \mathcal{P}(\alpha \text{ CRIS})$

This observation is rather inconveniently monolithic, and we can consider breaking it up into parts that, taken together, provide the same information: that is, if we have f_1 in $\mathcal{P}(E) \rightarrow X_1$, f_2 in $\mathcal{P}(E) \rightarrow X_2$, and so on, then the combined function f_1, f_2, \dots defines an isomorphism between ran obs and the range of the combination.

Our programmer's intuition says that this can 'obviously' and 'naturally' be done by providing one observation for each family, which for Suggest_A would yield the set of potential authors suggested to the committee, for Referee the relationship between referees and papers, and so on. With the suspicion of words in inverted commas born of experience, we should like a little more justification for this proposal, particularly as different programmers may choose different interpretations of the observations: for example, by returning the referees who are still to report on a paper rather than all the referees who have ever been asked to report on it.

The justification is a construction of the isomorphism between the unstructured and the structured observations. We have already established by axiom the isomorphism between the parameter values of a family and the set of events they define. Projection of the set of events which have occurred onto a single family is clearly not injective: if, however, we combine projections so that every event is represented somewhere in the image, then the resulting operation will be. This shows that one observation per family is necessary, and that it is sufficient to consider only the events in the alphabet of the behaviour, as defined by the parameters of the families.

This simple process of projection and label stripping uniformly gives us the historical rather than the current observation of processes like refereeing: justifying the current view is more involved, and we shall return to it when we consider implementations in the next chapter. In the meantime, let us define a pattern for the observations for a given family and parameter set

$$\text{Observation} \in \text{FAMILY} \times \mathcal{P}(X) \rightarrow (\text{TRACES} \rightarrow \mathcal{P}(X))$$

$$x \in \text{Observation}(f, S)(t) \Leftrightarrow x \in S \ \& \ f x \in \alpha t$$

Submitted = Observation(Submit, P × A)
 Refereeing = Observation(Referee, P × R)
 Cancelled = Observation(Cancel, P × R)
 Reports = Observation(Report, (P × R) × V)
 Acknowledged = Observation(Acknowledge, P × A)
 Decided = Observation(Decide, P × Dec)
 Informed = Observation(Inform, P × (A × Dec))
 Scheduled = Observation(Schedule, T × P)
 Suggested_A = Observation(Suggest_A, A)
 Called = Observation(Call, A)
 Intending = Observation(Intend, A)
 Suggested_D = Observation(Suggest_D, D)
 Invited = Observation(Invite, D)
 Accepted = Observation(Accept, D)
 Replied = Observation(Reply, D × Dec)

4 Implementing the system

The functional specification of the system was completed in the previous chapter: by the end, we have a description of what the system is to do, but with no indication of how, even to the extent of indicating hardware or user interface. At this stage, the project should be reviewed. The client can decide how much a system with the specified capabilities is worth, and this limits the budget of the implementation phase. The designer can then decide if the system can be provided within the budget, possibly by exploring the outlines of some designs in private. If both parties are then satisfied, the contracts can be signed and the implementation begun.

4.1 The properties of implementations

In this section, we shall establish a framework for talking about a particular implementation by developing the properties that all implementations must have.

We can regard the implemented system as a 'black box', which is provided with means for indicating the occurrence of individual events (including OOPS), and for determining the current values of the observations at any time. The box presumably has internal state, since the observations depend on the sequence of events input, in a way which mimics the observations defined for the sequences of events in the specification. An attempt to enter an event which is not acceptable as a continuation to the previous sequence causes no change in the observations. We can infer that the state must be rich enough to support the computation of the results of the observations and the effects of the operations which follow the indication of an event.

Formally, let us assume that the internal states are in some set DB and denote members of it by identifiers like d. To allow the results of the observations to be computed, the implementation must provide a function

$$\text{obs}_M \in \text{DB} \rightarrow \text{ran obs}$$

where the partiality of the function allows for 'unused' states. This reminds that the implementation must provide at least as many states as there are distinct observations.

The effect of entering an event in a given state may be formalized as a function

$$\text{effect} \in E_{\text{OOPS}} \rightarrow (\text{DB} \rightarrow \text{DB})$$

and this can be extended naturally to sequences of events

$$\begin{aligned} \text{effect} &\in \text{seq}(E_{\text{OOPS}}) \rightarrow (\text{DB} \rightarrow \text{DB}) \\ \text{effect} &= \text{extension}(\text{effect}) \end{aligned}$$

where extension is defined in appendix A. If the implementation is to mimic the specification, then, for some InitDB in DB, we must have

$$\forall t: \text{BTRACKS}. \text{obs } t \approx \text{obs}_M(\text{effect } t \text{ InitDB})$$

If, however, we want to interpret t in this axiom as "the sequence of events entered by the user", then this requirement is not strong enough: the result of any sequence must be correct. Our goal for establishing the correctness of an implementation then becomes

$$G \equiv \forall t: \text{seq}(E_{\text{OOPS}}). \text{obs } t = \text{obs}_M(\text{effect } t \text{ InitDB})$$

The function `obs` is so far defined only over `BTRACKS`, and we should formalize what we mean by observing arbitrary sequences.

$$\forall t: \overline{\text{BTRACKS}}. \text{obs } t = \text{obs } (\text{purify } t)$$

where

$$\begin{aligned} \text{purify } e &\in \text{seq}(E_{\text{OOPS}}) \rightarrow \text{BTRACKS} \\ \text{purify}(\langle \rangle) &= \langle \rangle \\ \text{purify } s \text{ accepts } e &\Rightarrow \\ &\text{purify}(s * \langle e \rangle) = \text{purify}(s) * \langle e \rangle \\ \neg \text{purify } s \text{ accepts } e &\Rightarrow \text{purify}(s * \langle e \rangle) = \text{purify}(s) \end{aligned}$$

The interpretation of `purify` clearly satisfies our informally stated requirement of no observable change caused by illegal events. Purification is clearly idempotent, and hence an identity on `BTRACKS`: hence

$$\vdash \forall t: \text{seq}(E_{\text{OOPS}}). \text{obs } t = \text{obs}(\text{purify } t)$$

The obvious strategy for proving that `G` holds for an implementation is by induction. In expanding out the basis step, we arrive at the goal

$$G1 \equiv \text{obs } \langle \rangle = \text{obs}_M \text{ InitDB}$$

and for the inductive step

$$\begin{aligned} G2 \equiv \text{obs } t = \text{obs}_M (\text{effect } t \text{ InitDB}) &\vdash \\ \text{obs}(t * \langle e \rangle) = \text{obs}_M (\text{effect } (t * \langle e \rangle) \text{ InitDB}) & \end{aligned}$$

This latter goal can be subdivided according to the validity of the event introduced to give

$$G3 \equiv \text{obs } t = \text{obs}_M(\text{effect } t \text{ InitDB}); \text{purify } t \underline{\text{accepts}} e \\ \vdash \text{obs}(t^*\langle e \rangle) = \text{obs}_M(\text{effect}(t^*\langle e \rangle) \text{ InitDB})$$

(a particular form of G2), and

$$G4 \equiv \text{obs } t = \text{obs}_M(\text{effect } t \text{ InitDB}); \neg \text{purify } t \underline{\text{accepts}} e \\ \vdash \text{obs}_M(\text{effect } t \text{ InitDB}) = \\ \text{obs}_M(\text{effect}(t^*\langle e \rangle) \text{ InitDB})$$

which together subsume G2 (the latter restating what we have always known, that illegal events have no effect). If we assume the forgiveness of obs, then we have

$$\text{Forgiving}(\text{obs}) \vdash \forall t: \text{BTRACKS}. \text{obs } t = \text{obs}(\text{edit } t)$$

and hence

$$\text{Forgiving}(\text{obs}) \vdash \\ \forall t: \text{seq}(E_{\text{OOPS}}). \text{obs}(\text{purify } t) = \text{obs}(\text{edit}(\text{purify } t)) \\ = \text{obs } t$$

which yields new subgoals from G3 by analysis of the last event

$$G5 \equiv \text{obs } t = \text{obs}_M(\text{effect } t \text{ InitDB}) \\ ; \text{purify } t \underline{\text{accepts}} e; e \in E \vdash \\ \text{obs}(t^*\langle e \rangle) = \text{obs}_M(\text{effect}(t^*\langle e \rangle) \text{ InitDB})$$

$$G6 \equiv \text{obs}_M(\text{effect } \text{OOPS} \text{ InitDB}) = \text{obs}_M \text{ InitDB}$$

$$\begin{aligned}
 G7 \equiv \text{obs } t &= \text{obs}_M(\text{effect } t \text{ InitDB}); \text{purify } t \text{ accepts } e \\
 &\vdash \text{obs}_M(\text{effect OOPS } (\text{effect } (t^*(e)) \text{ InitDB})) = \\
 &\quad \text{obs}_M(\text{effect } t \text{ InitDB})
 \end{aligned}$$

Of the goals so far, only G1 and G5 depend on the form of obs. By adopting the definition of the previous section, we may rewrite these as

$$G8 \equiv \text{obs}_M \text{ InitDB} = \Phi_E$$

$$\begin{aligned}
 G9 \equiv \text{obs } t &= \text{obs}_M(\text{effect } t \text{ InitDB}); \text{purify } t \text{ accepts } e \\
 &\quad ; e \in E \vdash \text{obs}_M(\text{effect } e \text{ (effect } t \text{ InitDB)}) = \\
 &\quad \quad \text{obs}_M(\text{effect } t \text{ InitDB}) \cup \{e\}
 \end{aligned}$$

respectively. The goal set to be established after this analysis is {G4, G6, G7, G8, G9}.

4.2 A particular implementation

inspection of the goals shows that, depending on the conditions satisfied by the previous history and the event, the observable effect to be produced by the event is no change (G4), the result before the previous event (G7), or is derived from the previous observation by adding an event (G9), a process which is guaranteed in the case of CRIS to produce a change in the observation. We might therefore take as the state a pair of sets of events

$$DB = \mathcal{P}(E) \times \mathcal{P}(E)$$

where one represents the current and the other the previous observation. The observation function is trivial

$$\text{obs}_M(c, p) = c$$

Formally, we can take G3 and G4 and our new found (although not unexpected) knowledge that events always have an effect, and deduce that

$$\begin{aligned}
 G3; G4 \vdash \forall t_1, t_2: \text{seq}(E_{\text{OOPS}}); e: E. \\
 \text{effect } t_1 \text{ InitDB} = \text{effect } t_2 \text{ InitDB} \Rightarrow \\
 \text{purify } t_1 \underline{\text{accepts}} e \Leftrightarrow \text{purify } t_2 \underline{\text{accepts}} e
 \end{aligned}$$

Informally, we can interpret this as the discovery that, if the system is to perform as specified, then the states of the implementation must contain enough information to decide when events are allowed. OOPS, of course, is always possible. An implementation must have this property to have any chance of correctness.

If we take the stronger condition of equality of the observation implying equivalent acceptance, we may rewrite the precondition on correctness as

$$\begin{aligned}
 G11 \equiv \forall t_1, t_2: \text{TRACES}; e: E \mid \alpha t_1 = \alpha t_2. \\
 t_1 \underline{\text{accepts}} e \Leftrightarrow t_2 \underline{\text{accepts}} e
 \end{aligned}$$

We should like to establish this for CRIS without having to enumerate all its traces. We can generalize the property to all behaviours

$$\begin{aligned}
 \text{DBLike} \subset \text{BEHAVIOUR} \\
 \text{DBLike}(b) \Leftrightarrow \forall t_1, t_2: \text{traces}(b); e: E \mid \alpha t_1 = \alpha t_2. \\
 t_1 (\underline{\text{accepts}} b) e \Leftrightarrow t_2 (\underline{\text{accepts}} b) e \ \& \\
 \text{finals } b \ t_1 \Leftrightarrow \text{finals } b \ t_2
 \end{aligned}$$

The clause involving `finals` strengthens the condition of G11 to include the 'event' of termination. It is required of the first behaviour in a sequential combination to preserve the weaker condition, and is preserved by the constructors in a similar way. It is therefore technically convenient to fold the two conditions together when stating the theorems below.

Definitions of properties of behaviours in terms of their sets of traces are not very helpful, since these sets can be large. It is possible to deduce some sufficient (but not necessary) conditions on the structure and alphabet of behaviours which are more readily determined.

$\vdash \text{DBLike}(\text{SKIP})$
 $\vdash \text{DBLike}(e)$
 $b_1, b_2 \in \text{BEHAVIOUR}; \text{DBLike}(b_1); \text{DBLike}(b_2)$
 $\vdash \text{Independent}(b_1, b_2) \Rightarrow \text{DBLike}(b_1; b_2)$
 $\vdash \text{Separable}(b_1, b_2) \Rightarrow \text{DBLike}(b_1 \parallel b_2)$
 $\vdash \text{DBLike}(b_1 \parallel b_2)$

where

$\text{Separable} \in \text{BEHAVIOUR} \leftrightarrow \text{BEHAVIOUR}$
 $\text{Separable}(b_1, b_2) \Leftrightarrow$
 $\forall t_1, t_2: \text{traces}(b_1 \parallel b_2) \mid \alpha t_1 = \alpha t_2.$
 $(t_1 \in \text{traces } b_1 \Leftrightarrow t_2 \in \text{traces } b_1) \ \&$
 $(t_1 \in \text{traces } b_2 \Leftrightarrow t_2 \in \text{traces } b_2)$

Separability is not, of course, expressed as a structural condition, but we shall show two forms of separable behaviour in the next section. In each case above, the preconditions make the proof trivial, and are relatively strong: independence of behaviours joined by $;$ implies that each event (as in CRIS) occurs only once. We can use these theorems to show

$\vdash \text{DBLike}(\text{CRIS})$
 $\vdash G11$

Now we have shown the adequacy of the state, let us complete the construction by building the operations upon it, and determining the initial state. Let us split the effect functions along the lines suggested by the conditions of G4 and G9.

$$\begin{aligned} \forall e: E. \text{ effect } e = \text{Id} \oplus \text{ effect}_{\mathbf{N}} e \ \& \\ \text{dom } \text{effect}_{\mathbf{N}} e = \\ (\text{effect } t \text{ InitDB} \mid t: \text{seq}(E_{\text{OOPS}}) \\ \mid \text{purify } t \text{ accepts } e) \end{aligned}$$

With this construction, clearly,

$$\forall t: \text{seq}(E_{\text{OOPS}}). \text{ effect } t = \text{effect } (\text{purify } t)$$

Then, as a consequence of G11, we have

$$\begin{aligned} \vdash \forall t: \text{seq}(E_{\text{OOPS}}); e: E. \\ \text{effect } t \text{ InitDB} \in \text{dom } \text{effect}_{\mathbf{N}} e \Leftrightarrow \text{purify } t \text{ accepts } e \end{aligned}$$

which immediately establishes G4. If we choose

$$\begin{aligned} \forall (c, p): \text{dom}(\text{effect}_{\mathbf{N}} e). \\ \text{effect}_{\mathbf{N}} e (c, p) \in \{(c \cup \{e\}, S) \mid S: \mathcal{P}(E)\} \end{aligned}$$

then G9 is also established, and we can take

$$\text{InitDB} \in \{(\Phi_E, S) \mid S: \mathcal{P}(E)\}$$

to satisfy G8. Consideration of the requirements of OOPS dictates that these be strengthened to

$$\begin{aligned} \forall (c, p): \text{dom}(\text{effect}_{\mathbf{N}} e). \text{ effect}_{\mathbf{N}} e (c, p) = (c \cup \{e\}, c) \\ \text{InitDB} = (\Phi_E, \Phi_E) \end{aligned}$$

with

$$\text{effect OOPS } (c, p) = (p, c)$$

given which, the remaining goals G6 and G7 are readily established.

4.3 A more efficient implementation

There are three areas in which the efficiency of the proposed implementation is dubious particularly with regard to the space required the representation of the previous state for OOPS, the form in which the domains of the operations are expressed, and the representation of the state as an arbitrary set of events. We shall deal with these in turn.

We know that each event occurs at most once, and hence the effect of a legal, real event is to add precisely that event to the observation. This suggests that we can represent one state by the single event needed to derive it from the other. It is convenient to derive the previous state from the current state.

$$DB_1 = \mathcal{P}(E) \times E_{OOPS}$$

We can write a retrieve function to the previous state as

$$\begin{aligned} \text{ret} &\in DB_1 \gg \{\text{effect } t \text{ InitDB} \mid t: \text{TRACES}\} \\ e \in c &\Rightarrow \text{ret}(c, e) = (c, c - \{e\}) \\ e \in c \ \& \ e \neq \text{OOPS} &\Rightarrow \text{ret}(c, e) = (c, c \cup \{e\}) \\ \text{ret}(c, \text{OOPS}) &= (c, c) \end{aligned}$$

and, since this is a bijection, immediately derive the corresponding operations from the previous implementation by defining

$$\begin{aligned} \text{effect } e \in E_{OOPS} &\rightarrow (DB_1 \rightarrow DB_1) \\ \text{effect}_N e \in E_{OOPS} &\rightarrow (DB_1 \rightarrow DB_1) \\ \text{effect } e &= \text{ret}; \text{effect } e; \text{ret}^{-1} \end{aligned}$$

This gives

$$\begin{aligned} e \in E &\Rightarrow \text{effect}_N e (c, e') = (c \cup \{e\}, e) \\ \text{effect OOPS } (c, \text{OOPS}) &= (c, \text{OOPS}) \end{aligned}$$

$$\begin{aligned}
e \in E; e \notin c &\Rightarrow \text{effect OOPS}(c, e) = (c \cup \{e\}, e) \\
e \in E; e \in c &\Rightarrow \text{effect OOPS}(c, e) = (c - \{e\}, e) \\
\text{InitDB} &= (\Phi_E, \text{OOPS})
\end{aligned}$$

Now let us turn our attention to the domain conditions for the operations. We can reshape the collected domain conditions as a relation analogous to the acceptance relations on sequences.

$$\begin{aligned}
\underline{\text{accepts}} \in \text{DB} &\leftrightarrow E \\
d \underline{\text{accepts}} e &\Leftrightarrow d \in \text{dom}(\text{effect}_{\text{N}} e)
\end{aligned}$$

What we are seeking is a relation $\underline{\text{accepts}}_{\text{M}}$ equivalent to this, but expressed in a form more oriented to computation than a set of pairs of states and events. Equivalence in this context means giving the same result over the subset of DB reached by sequences of events.

$$\begin{aligned}
\underline{\text{accepts}}_{\text{M}} \in \text{DB} &\leftrightarrow E \\
\forall t: \text{seq}(E_{\text{OOPS}}); e: E. &\text{effect } t \text{ InitDB } \underline{\text{accepts}} e \Leftrightarrow \\
&\text{effect } t \text{ InitDB } \underline{\text{accepts}}_{\text{M}} e
\end{aligned}$$

Making use of the results on $\underline{\text{accepts}}$ and purify , this can be reduced to an equivalent goal

$$\begin{aligned}
G12 \equiv \forall t: \text{BTRACKS}; e: E. \\
\text{effect } t \text{ InitDB } \underline{\text{accepts}}_{\text{M}} \Leftrightarrow t \underline{\text{accepts}} e
\end{aligned}$$

and if we agree to construct $\underline{\text{accept}}_{\text{M}}$ as

$$\begin{aligned}
\underline{\text{accepts}}_{\text{A}} \in \mathcal{P}(E) &\leftrightarrow E \\
\underline{\text{accepts}}_{\text{M}} &= \text{obs}_{\text{M}}; \underline{\text{accepts}}_{\text{A}}
\end{aligned}$$

then this finally reduces to showing

$$G13 \equiv \forall t: \text{TRACES}; e: E. \alpha t \underline{\text{accepts}}_{\text{A}} e \Leftrightarrow t \underline{\text{accepts}} e$$

G11 gives us grounds for thinking that such a relation may exist

Let us look at a particular event. $\text{Reports}(r, p, v)$. Looking at the specification of the behaviour, it seems that the view REF gives us sufficient context to determine when the event can occur after $\text{Referee}(r, p)$, but not if any $\text{Report}(r, p, v')$ or $\text{Cancel}(r, p)$ have occurred. We should attempt to formalize these ideas and justify our intuition. The basic test of a state looks at the events which have occurred and those which have not

$$\text{FPAIR} = \mathcal{P}(E) \times \mathcal{P}(E)$$

$$\underline{\text{sel}} \in \text{FPAIR} \leftrightarrow \mathcal{P}(E)$$

$$(o, n) \underline{\text{sel}} S \Leftrightarrow o \subset S \ \& \ \text{Disjoint}(n, S)$$

We can take a set of these, and define the new test as the disjunction of the tests of its elements

$$\text{FILTER} = \mathcal{P}(\text{FPAIR})$$

$$\underline{\text{sel}} \in \text{FILTER} \leftrightarrow \mathcal{P}(E)$$

$$F \underline{\text{sel}} S \Leftrightarrow \exists f: F. f \underline{\text{sel}} S$$

and, given a filter for each event

$$\underline{\text{has filter}} \in \text{BEHAVIOUR} \rightarrow (E \leftrightarrow \text{FILTER})$$

we can construct, given some filter F for e

$$e (\underline{\text{has filter}} \text{ CRIS}) F \Rightarrow S \underline{\text{accepts}}_A e \Leftrightarrow F \underline{\text{sel}} S$$

and hence can define in general

$$\begin{aligned} e (\underline{\text{has filter}} \ b) F &\Leftrightarrow \\ \forall t: \text{traces } b. F \underline{\text{sel}} \alpha t &\Leftrightarrow t (\underline{\text{accepts}} \ b) e \ \& \\ \alpha F &= \alpha b \end{aligned}$$

where

$$\alpha \in \text{FILTER} \rightarrow \mathcal{P}(E)$$

$$\alpha F = \sqcup \{e \cup e' \mid (e, e') : F\}$$

The 'no junk' condition is useful when we come to construct filters for particular shapes of behaviour

Then, from the basic observations

$$e, e' \in E; e' \neq e$$

$$\vdash e \text{ (has filter SKIP) } \Phi_{\text{FPAIR}}$$

$$\vdash e \text{ (has filter } e) \quad \{(\Phi_E, \{e\})\}$$

$$\vdash e' \text{ (has filter } e) \quad \Phi_{\text{FPAIR}}$$

it follows that

$$\text{Reports}(r, p, v) \text{ (has filter Reports}(r, p, v))$$

$$\{(\Phi_E, \{\text{Reports}(r, p, v)\})\},$$

$$\text{Reports}(r, p, v)$$

$$\text{(has filter } \sqcup \{\text{Reports}(r, p, v) \mid v : V\})$$

$$\{(\Phi_E, \{\text{Reports}(r, p, v \mid v : V)\})\}$$

by applying lemma

$$H: \text{Deterministic}(b_1, b_2);$$

$$\vdash e \text{ (has filter } (b_1 \sqcup b_2))$$

$$\{(\{s, s \cup \{e : E \mid e \text{ starts } b_2\}\} \mid (s, s') : F_1\} \cup$$

$$\{(\{s, s \cup \{e : E \mid e \text{ starts } b_1\}\} \mid (s, s') : F_2\},$$

$$\text{Reports}(r, p, v) \text{ (has filter REF)}$$

$$\{(\{\text{Referee}(r, p)\}, \{\text{Reports}(r, p, v \mid v : V) \cup$$

$$\{\text{Cancel}(r, p)\} \})\}$$

by applying the previous lemma and

$$\begin{aligned}
 & H; e \text{ starts } b_2; b_1 \text{ has ends } F_1; \text{Independent}(b_1, b_2) \\
 & \quad \vdash e \text{ (has filter } (b_1; b_2)) \\
 & \quad \quad \{ (s, s' \cup t') \mid (s, s') : F_1; t' : \mathcal{P}(E) \\
 & \quad \quad \quad \mid (\Phi_E, t') \in F_2 \} \cup \\
 & \quad \quad \{ (s, s') \mid (s, s') : F_2 \mid s' \neq \Phi_E \}
 \end{aligned}$$

and that this is also a filter in CRIS by applying

$$\begin{aligned}
 & H; e \in \alpha b_1; \text{Independent}(b_1, b_2) \\
 & \quad \vdash e \text{ (has filter } (b_1; b_2)) F_1
 \end{aligned}$$

$$\begin{aligned}
 & H; e \in \alpha b_2; \neg e \text{ starts } b_2; \text{Independent}(b_1, b_2) \\
 & \quad \vdash e \text{ (has filter } (b_1; b_2)) F_2
 \end{aligned}$$

$$\begin{aligned}
 & H; e \in \alpha b_1; e \notin \alpha b_2; \\
 & \quad \vdash e \text{ (has filter } (b_1 \parallel b_2)) F_1
 \end{aligned}$$

where

$$H \equiv e \text{ (has filter } b_1) F_1; e \text{ (has filter } b_2) F_2$$

Deterministic \in BEHAVIOUR \leftrightarrow BEHAVIOUR

Deterministic(b_1, b_2) \Leftrightarrow

$$\begin{aligned}
 & \forall e: E. e \text{ starts } b_1 \Rightarrow e \notin \alpha b_2 \ \& \\
 & \quad e \text{ starts } b_2 \Rightarrow e \notin \alpha b_1
 \end{aligned}$$

has ends \in BEHAVIOUR \leftrightarrow FILTER

b has ends $F \Leftrightarrow$

$$\forall t: \text{traces}(b). F \text{ sel } \alpha t \Leftrightarrow \text{finals } b \ t$$

Determinism implies separability (see the previous section). Behaviours of the form $e; b_1 \square b_1$ are also separable. We can build a structural analysis for has ends by treating termination as some special event occurring and finding the filter for that event.

This process of going from simpler to more complex behaviours in stages, guessing at each step how to modify the filter to preserve its correctness, and then proving the construction correct, can be carried out for each event. We might expect to arrive reasonably quickly at a sufficient set of lemmas, covering all the special cases which have particularly compact forms of filter.

Finally, let us consider how to represent sets of events. We can gain more insight by looking at the results of the separate functions which collectively yield a result isomorphic to that of the original observation `obs`.

```
obsS = Submitted, Refereeing, Cancelled, Reports,
      Acknowledged, Decided, Informed, Scheduled,
      SuggestedA, Called, Intending, SuggestedD,
      Invited, Accepted, Replied
```

What is the range of this function? If we look at one of the components, `Reports` for example, we can see that

$$\text{Reports} \in \text{BTRACKS} \rightarrow \mathcal{P}((P \times R) \times V)$$

which is conventionally written

$$\text{Reports} \in \text{BTRACKS} \rightarrow (P \times R \leftrightarrow V)$$

Such rewritings into relational form apply to all the observations involving cross products. In this case, we can go further. After some manipulation, we can write the range as

```
ran Reports =
  {(top Report)-1(α t) |
   t: {t | Report from P×R×V | t: TRACES}}
```

Reports \in BTRACKS \rightarrow (P \times R \rightarrow V)

we need to demonstrate. for t in TRACES

G14 = {Report(p, r, v₁), Report(p, r, v₂)} \subset at \Rightarrow v₁ = v₂

We have by enumeration

t \in traces(\exists {Report(r, p, v) | v: V}) \vdash G14

and, since

\vdash CRIS sat \exists {Report(r, p, v) | v: V}

we can use the theorem concerning traces and satisfaction to give the desired result. Following this pattern in each case, we can show that

obs_s \in BTRACKS \rightarrow RANGE

where

RANGE = ((P \rightarrow A) \times (P \leftrightarrow R) \times (P \leftrightarrow R) \times (P \times R \rightarrow V) \times
P(P) \times (P \rightarrow Dec) \times (P \rightarrow A \times Dec) \times (T \rightarrow P) \times
P(A) \times P(A) \times P(A) \times P(D) \times
P(D) \times P(D) \times (D \rightarrow Dec))

and further, that

dom Refereeing \subset dom Submitted
dom Reports \subset Refereeing
Cancelled \subset Refereeing
Disjoint(dom Reports, Cancelled)
Acknowledged \subset dom Submitted
 $\forall p$: dom Informed; a: A; d: Dec | Informed p = (a, d).
Decided p = d & Submitted p = a

We can conclude that the set

$$DB_2 = RANGE \times E_{COOPS}$$

implements DB_1 in an obvious way

As a last step, we shall explore the sets which are isomorphic to the subset of RANGE defined by the theorems interrelating observations, to present the structure in a way more appropriate to a filing system and to attempt to have fewer unreachable states in the obvious implementation

In a manual filing system, we might consider allocating one drawer to files labelled by paper, containing the authorship, refereeing and reporting information relevant to that paper, another drawer to files on people (identified by elements of Q, which is henceforth to include A, D, and R) recording their status as delegates and authors, and a timetable to be filled in as scheduling progresses. Formally, we can define two sub-observations

$$obs_p = Submitted, Refereeing, Cancelled, Reports, \\ Acknowledged, Decided, Informed$$

$$obs_q = Suggested_A, Called, Intending, Suggested_D, \\ Invited, Accepted, Replied$$

such that

$$obs_s = obs_p, Scheduled, obs_q$$

Consider obs_p

$$obs_p \in BTRACKS \rightarrow RANGE_p \\ RANGE_p = ((P \rightarrow A) \times (P \leftrightarrow R) \times (P \leftrightarrow R) \times (P \times R \rightarrow V) \times \\ P(P) \times (P \rightarrow Dec) \times (P \rightarrow A \times Dec))$$

We want to convert elements of $RANGE_p$ to functions in

$$PF = P \rightarrow PAPER$$

for some (structured) set PAPER, which we can do in four reversible stages. The first is to convert each component of $RANGE_p$ to a partial function over P.

$$RP_1 = ((P \rightarrow Q) \times (P \rightarrow P_1(Q)) \times (P \rightarrow P_1(Q)) \times \\ (P \rightarrow Q \rightarrow V) \times (P \rightarrow \{ACKD\}) \times (P \rightarrow Dec) \times \\ (P \rightarrow (A \times Dec)))$$

$$stage_1 \in RANGE_p \gg RP_1$$

$$stage_1 = Id \times \# \times \# \times CP \times const\ ACKD \times Id \times Id$$

where the bijection property comes immediately from that of the components (of which $\#$, CP and $const$ are defined in appendix A).

To turn this into a single function with cross-product target, the use of the combinator \cdot is appropriate. We can show that this is bijective if it is restricted to functions with common domain, so the second stage is to reduce the functions to this form. It follows from the restrictions on $RANGE$ that the domain of Submitted t includes that of any other function, for any t in BTRACKS. Domain extension can be done conveniently and reversibly by 'filling in the gaps' with a new value.

$$\text{extend} \in Y \rightarrow (P(X) \times (X \rightarrow Y) \rightarrow (X \rightarrow Y)) \\ \text{extend } y (S, f) = const(S, y) \# f$$

$$\vdash \text{extend } y \downarrow (S \times (S \rightarrow \overline{\{y\}})) \in S \times (S \rightarrow \overline{\{y\}}) \gg (S \rightarrow Y)$$

$$RP_2 = ((P \rightarrow Q) \times (P \rightarrow \mathcal{P}(Q)) \times (P \rightarrow \mathcal{P}(Q)) \times \\ (P \rightarrow Q \rightarrow V) \times (P \rightarrow \{ACKD\}_\perp) \times (P \rightarrow Dec_\perp) \times \\ (P \rightarrow (A \times Dec)_\perp))$$

$$stage_2 \in RP_1 \gg RP_2$$

$$stage_2(s, r, c, v, a, d, i) = \\ (s, \text{extend } \Phi(\text{dom } s, r), \text{extend } \Phi(\text{dom } s, c), \\ \text{extend } \Phi(\text{dom } s, v), \text{extend } i_{\{ACKD\}}(\text{dom } s, a), \\ \text{extend } i_{Dec}(\text{dom } s, d), \text{extend } i_{A \times Dec}(\text{dom } s, i))$$

where. for all sets S

$$i_S \notin S \\ S_\perp = S \cup \{i_S\}$$

Then

$$RP_3 = P \rightarrow (A \times \mathcal{P}(Q) \times \mathcal{P}(Q) \times (Q \rightarrow V) \times \\ \{ACKD\}_\perp \times Dec_\perp \times (A \times Dec)_\perp)$$

$$stage_3 \in RP_2 \rightarrow RP_3$$

$$stage_3(s, r, c, v, a, d, i) = s, r, c, v, a, d, i$$

Lastly. given

$$V_c = V \cup \{PENDING, CANCELLED\}$$

we can simplify the range of the resulting function

$$PAPER = (A \times (Q \rightarrow V_c) \times \{ACKD\}_\perp \times Dec_\perp \times \{VSENT\}_\perp)$$

$$stage_4 \in RP_3 \rightarrow PAPER$$

$$\text{stage}_4(s, r, c, v, a, d, i_{A \times \text{Dec}}) =$$

$$(s, \text{extend PENDING}(r, \text{extend CANCELLED}(c, v)),$$

$$a, d, \{\text{VSENT}\}_1)$$

$$i \in A \times D \Rightarrow \text{stage}_4(s, r, c, v, a, d, i) =$$

$$(s, \text{extend PENDING}(r, \text{extend CANCELLED}(c, v)),$$

$$a, d, \text{VSENT})$$

and, since the composition of these stages over ran obs_p is bijective, it is a relatively simple matter to determine how the events should update this new representation

Similarly, we can construct

$$QF \in Q \rightarrow \text{PERSON}$$

$$\text{PERSON} = \{\text{SUGGESTED}_A\}_1 \times \{\text{CALLED}\}_1 \times \{\text{INTENDING}\}_1 \times$$

$$\{\text{SUGGESTED}_D\}_1 \times \{\text{INVITED}\}_1 \times \{\text{ACCEPTED}\}_1 \times$$

$$\text{Dec}_1$$

to give an implementation

$$DB_3 = PF \times QF \times (T \rightarrow P) \times E_{\text{OOPS}}$$

5 Summary

In the previous chapters, we have presented a scheme for describing the behaviour of the real world in terms of sequences of events communicated to a central point, and used it in the particular case of describing the process of conference organization. From this, we derived the sequences of events which could be fed into a machine equipped with a simple system for correcting errors in data entry, and described a set of observations which could be made from the machine and which would mimic the situation in the world. We then developed a design for the state of the machine, to a level from which the implementation in terms of arrays, files and other conventional data structures should be trivial.

The philosophy of our approach, of taking part of the world and describing its behaviour before embarking on a functional description of a system, is also that of the Jackson Systems Design method ([Jackson83]). This too has been strongly influenced by CSP, drawing on it as a structured approach to programming concurrent systems. A JSD specification is a collection of sequential programs with state, and the method produces programs which emphasise human comprehension rather than efficiency, although they are at least 'walkable' (to use Burstall's terminology). These specifications are then transformed systematically into programs more efficiently executable using current computer systems. In contrast, our approach draws on the more mathematical end of the CSP literature, and deals less systematically with implementation issues.

We can also compare the approach here with the constructive specification style of [Jones80]. The specification in that case would have started with a proposal for the state, perhaps that of DB_1 , and operations corresponding to the events and defined on that state: that is to say, at a much later stage than here. Subsequent refinements would require the construction of a retrieve function to show which abstract state corresponded to a given representation. This then forms the basis for proofs of adequacy of the representation and correctness of the operations constructed on it.

The operations of the specification would include preconditions similar to those derived in the previous chapter to ensure their activation only at the correct time

The advantage of the method presented here is to make explicit information which is only a consequence of the operational definition. The domain conditions of the operations are there to satisfy the requirement that operations should occur only in a given order. This requirement is stated in this specification and the consequences for the implementation derived later. Certain incidental advantages follow. OOPS was introduced by stating general properties, rather than as a separate operation on a specific state. We had at each level of implementation a characterization of the precise subset of the states which were actually reachable, which facilitated the refinement of the state by means of bijections, and hence guaranteed adequacy and made the definition of refined operations a more or less mechanical task. We were also able to separate out the rather complex state oriented characterizations of domain from the rather simple actions of the operations.

The other point of comparison is CSP ([Hoare80, Hoare81]). The original incentive to develop the theory of the second chapter came from this work, although it was developed with the idea of describing a limited class of terminating behaviours, rather than potentially infinite processes. There are thus differences in the interpretation of the terms: what we are doing is 'tapping' the line of communication between the committees and the rest of the world, and attempting to describe the possible sequences. The attempt may take the form of postulating various process-like entities, such as authors and referees, which may even make 'silent transitions' as a result of activity we cannot see.

There are some differences between the model suggested here and that of [Hoare80]. The use of ' \surd ' rather than explicitly separating out the final traces is, for finite behaviours, purely one of taste. Our model is more explicit about alphabets, but allows the distinguishing of two behaviours by alphabet alone, which is hard to justify on the basis of experiments on

communications. Finally, the behaviours, unlike the processes of [Hoare80] as restricted by technical note (2), do not necessarily terminate deterministically. This allows us to use τ , which captures a common situation in a world where we are not privy to all the information which influences others.

The implementation is by no means complete, since we have deliberately abstracted from issues of the user interface. This was explicit in the choice of observations, and it is now time to honour our promise to consult the client. We have also made an implicit abstraction by saying that P and Q identify papers and people, without explaining how the operator is to communicate these identifiers to the system. Work remains to be done on formalizing these aspects of communication between user and system.

Appendix A
Some extensions to the theory of [Abrial81]

There are a few simple notions which we have used in the body of the specification which are not defined in [Abrial81]. They are presented here for completeness.

The first is the disjointness relationship between sets. The abstraction is rather more attractive in use than its definition.

$$\begin{aligned} \text{Disjoint} &\in \mathcal{P}(X) \leftrightarrow \mathcal{P}(X) \\ \text{Disjoint}(S_1, S_2) &\Leftrightarrow S_1 \cap S_2 = \emptyset_X \end{aligned}$$

The same applies to the second, which provides a definition of the alphabet of a sequence.

$$\begin{aligned} \alpha &\in \text{seq}(X) \rightarrow \mathcal{P}(X) \\ \alpha s &= \text{ran } s \end{aligned}$$

Given an associative and commutative binary operator, and an element of its source set (usually a unit), there is a natural promotion to finite sets of operands.

$$\begin{aligned} \text{Associative} &\subset X \times X \rightarrow X \\ \text{Associative}(\cdot) &\Leftrightarrow \forall x_1, x_2, x_3: X. x_1 \cdot (x_2 \cdot x_3) = (x_1 \cdot x_2) \cdot x_3 \\ \text{Commutative} &\subset X \times X \rightarrow Y \\ \text{Commutative}(\cdot) &\Leftrightarrow \forall x_1, x_2: X. x_1 \cdot x_2 = x_2 \cdot x_1 \end{aligned}$$

$\text{continuation} \in (X \times X \rightarrow X) \times X \rightarrow (\mathcal{P}(X) \rightarrow X)$
 Associative(f) & Commutative(f) & $x' \in S \Rightarrow$
 $\text{continuation}(f, x)(\Phi_x) = x$ &
 $\text{continuation}(f, x) S =$
 $f(x', \text{continuation}(f, x)(S - \{x'\}))$

Given an indexed set of homogeneous relations, we can convert a sequence of indices into a relation by composition.

$\text{extension} \in (I \rightarrow (X \leftrightarrow X)) \rightarrow (\text{seq}(I) \rightarrow (X \leftrightarrow X))$
 $\text{extension } F \langle \rangle = \text{Id}$
 $\text{extension } F(1^{\wedge} s) = F(i); (\text{extension } F s)$

The remaining combinators are simple rearrangements of structures to put them into a more convenient form.

$\text{const} \in \mathcal{P}(X) \times Y \gg (X \rightarrow Y)$
 $\text{const}(S, y) = \lambda x: X \mid x \in S. y$

$\text{curry} \in (X \times Y \rightarrow Z) \gg (X \rightarrow (Y \rightarrow Z))$
 $\forall(x, y): \text{dom } f. \text{curry } f \ x \ y = f(x, y)$

$\text{CP} \in (X \times Y \rightarrow Z) \gg (X \rightarrow (Y \rightarrow Z))$
 $f^{\text{CP}} = (\text{curry } f) \downarrow \{\Phi\}$

$\# \in (X \leftrightarrow Y) \gg (X \rightarrow \mathcal{P}_1(Y))$
 $\text{dom } \underline{x}^{\#} = \text{dom } \underline{x}$
 $\forall x: \text{dom } \underline{x}. y \in \underline{x}^{\#}(x) \Leftrightarrow x \underline{x} y$

References

- [Abrial81] *A course on system specification*
J. R. Abrial (PRG internal document)
- [Gustafsson81] *A declarative approach to conceptual information modelling*
M. R. Gustafsson, T. Karlsson, J. Bubenko
(SYSLAB Report 8)
- [Gutttag83] *An introduction to the Larch shared language*
J. V. Gutttag (IFIP'83 Proceedings (to appear))
- [Hoare80] *A model for communicating sequential processes*
C. A. R. Hoare (Technical monograph PRG-22)
- [Hoare81] *A theory of communicating sequential processes*
C. A. R. Hoare, S. D. Brookes, A. W. Roscoe
(Technical monograph PRG-16)
- [Jackson83] *System Development*
M. Jackson (Prentice-Hall International 1983)
- [Jones80] *Software Development - A Rigorous Approach*
C. B. Jones (Prentice-Hall International 1980)
- [Ollie82] *Information systems design methodologies*
T. W. Ollie, H. G. Sol, A. A. Verrijn-Stuart (eds)
(North Holland 1982)
- [Suftrin83] *Disjoint Unions, Recursive Types, Enumerated Types*
B. A. Suftrin (PRG internal document)

OXFORD UNIVERSITY COMPUTING LABORATORY
PROGRAMMING RESEARCH GROUP TECHNICAL MONOGRAPHS

SEPTEMBER 1983

This is a series of technical monographs on topics in the field of computation. Copies may be obtained from the Programming Research Group, (Technical Monographs, 8-11) Keble Road, Oxford, OX1 3QD, England.

- PRG-2 Dana Scott
Outline of a Mathematical Theory of Computation
- PRG-3 Dana Scott
The Lattice of Flow Diagrams
- PRG-5 Dana Scott
Data Types as Lattices
- PRG-6 Dana Scott and Christopher Strachey
Toward a Mathematical Semantics for Computer Languages
- PRG-7 Dana Scott
Continuous Lattices
- PRG-8 Joseph Stoy and Christopher Strachey
*OS6 - an Experimental Operating System
for a Small Computer*
- PRG-9 Christopher Strachey and Joseph Stoy
The Text of OSPub
- PRG-10 Christopher Strachey
The Varieties of Programming Language
- PRG-11 Christopher Strachey and Christopher P. Wadsworth
*Continuations. A Mathematical Semantics
for Handling Full Jumps*
- PRG-12 Peter Mosses
The Mathematical Semantics of Algol 60
- PRG-13 Robert Milne
*The Formal Semantics of Computer Languages
and their Implementations*
- PRG-14 Shan S. Kuo, Michael H. Linck and Sohrab Saadai
A Guide to Communicating Sequential Processes
- PRG-15 Joseph Stoy
The Congruence of Two Programming Language Definitions
- PRG-16 C. A. R. Hoare, S. D. Brookes and A. W. Roscoe
A Theory of Communicating Sequential Processes
- PRG-17 Andrew P. Black
Report on the Programming Notation 3R
- PRG 18 Elizabeth Fielding
*The Specification of Abstract Mappings
and their implementation as B^T-trees*

- PRG-19 Dana Scott
Lectures on a Mathematical Theory of Computation
- PRG-20 Zhou Chao Chen and C A R Hoare
Partial Correctness of Communicating Processes and Protocols
- PRG-21 Bernard Sulrin
Formal Specification of a Display Editor
- PRG-22 C A R. hoare
A Model for Communicating Sequential Processes
- PRG-23 C A R Hoare
A Calculus of Total Correctness for Communicating Processes
- PRG-24 Bernard Sulrin
Reading Formal Specifications
- PRG-25 C B Jones
Development Methods for Computer Programs including a Notion of Interference
- PRG-26 Zhou Chao Chen
The Consistency of the Calculus of Total Correctness for Communicating Processes
- PRG-27 C A R. Hoare
Programming is an Engineering Profession
- PRG-28 John Hughes
Graph Reduction with Super-Combinators
- PRG-29 C A R. Hoare
Specifications, Programs and Implementations
- PRG-30 Alejandro Teruel
Case Studies in Specification: Four Games
- PRG-31 Ian D Cottam
The Rigorous Development of a System Version Control Database
- PRG-32 Peter Henderson Geraint A Jones and Simon B Jones
The Lispkit Manual
- PRG-33 C A R. Hoare
Notes on Communicating Sequential Processes
- PRG-34 Simon B Jones
Abstract Machine Support for Purely Functional Operating Systems
- PRG-35 S D Brookes
A Non-deterministic Model for Communicating Sequential Processes
- PRG-36 T. Clement
The Formal Specification of a Conference Organizing System