

CAVIAR: A Case Study in Specification

by

Bill Flinn
Standard Telecommunication Laboratories
Harlow, England

and

Ib Holm Sørensen
Programming Research Group
Oxford University Computing Laboratory

Oxford University Computing Laboratory
Wolfson Building
Parks Road
Oxford OX1 3QD

Technical Monograph PRG-48

June 1985

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

Copyright © 1985 Bill Flinn and Ib Holm Sørensen

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

CAVIAR: A Case Study in Specification

Bill Flinn and Ib Holm Sørensen

Abstract

This paper describes the specification, written in the specification language known as Z, of a reasonably complex software system. Important features of the Z approach which are highlighted in this paper include the interleaving of mathematical text with informal prose, the creation of parametrised specifications, and use of the Z schema calculus to construct descriptions of large systems from simpler components.

Contents

- 0. Introduction**
- 1. The Case Study**
- 2. Identification of the Basic Sets**
- 3. The Subsystems of CAVIAR**
- 4. A General Resource-User System**
- 5. Specialisation of the General R-U System**
 - 5.1 An R-U system where resources cannot be shared**
 - 5.2 An R-U system where each user may occupy at most one resource**
 - 5.3 An R-U system where a user occupies
at most one nonsharable resource**
 - 5.4 The specification library**
- 6. Classification and Instantiation**
 - 6.1 Some laws for CAVIAR**
 - 6.2 Matching system with models**
 - 6.3 The hotel reservation subsystem - HR-V**
 - 6.4 The transport reservation subsystem - TR-V**
- 7. The Meeting Attendance Subsystem**
 - 7.1 A pool system**
 - 7.2 The meeting - visitor subsystem**

- 8. The Meeting Resource Subsystems**
 - 8.1 A diary system**
 - 8.2 The conference room booking subsystem**
 - 8.3 The dining room booking subsystem**
 - 8.4 The visitor pool - V-P**
 - 8.5 The construction process**

- 9. The Complete CAVIAR System**
 - 9.1 Combining subsystems to form the system state**
 - 9.2 Operations on CAVIAR**
 - 9.2.1 Operations which involve meetings only**
 - 9.2.2 Operations which involve visitors only**
 - 9.2.3 A general visitor removal operation**

- 10. Conclusion**

- 11. Acknowledgements**

- 12. References and Related Work**

- Appendix: Mathematical Notation**

- Appendix: Schema Notation**

0. Introduction

This paper presents a case study in system specification. The notation used to record the system's properties is known as Z [1, 2, 3]. Z is based on set theory, and its use as a specification language has been developed at the Programming Research Group at Oxford University. Some important aspects of the Z approach are illustrated in this paper.

As is well known, software development can be divided into several phases; requirements analysis, specification, design and implementation. Z can be applied in both the specification and design phases; however, in this paper we will address the specification phase only.

We view a specification as having a two-fold purpose: firstly, to give a formal (mathematical) system description which provides a basis from which to construct a design. Such a mathematical description is essential if we are to prove formally that a design meets its specification. Secondly, to give an informal statement of the system's properties, in order that the specification can be tested (validated) against the (usually informal) statement of requirements. Thus the Z approach is to construct a specification document which consists of a judicious mix of informal prose with precise mathematical statements. The two parts of the document are complementary in that the informal text can be viewed as commentary for the formal text. It can be consulted to find out what aspects of the real world are being described and how it relates to the informally stated requirements. The formal text on the other hand provides the precise definition of the system and hence can be used to resolve any ambiguities present in the informal text. A beneficial side effect for practitioners writing such documents is that their understanding of the system in question is greatly helped by the process of constructing both the formal and the informal descriptions.

It is often the case that the process of abstraction used to construct a specification results in structures which are more general than those actually required for the system being considered. It is part of the Z approach to identify and describe such general structures. These descriptions can be placed in a specification library. Particular cases of these general components can then be used later, either as part of the current system or in subsequent projects.

This specification case study develops a number of general systems which are subsequently constrained and combined to form the complete system description.

1. The Case Study

This specification of a Computer Aided Visitor Information And Retrieval system resulted from the analysis of a manual system concerned with recording and retrieval of data about arrangements for visitors and meetings at a large industrial site. Standard Telecommunications Laboratories (U.K.) sponsored the study in order to investigate the feasibility of converting to a computer based solution. Of particular concern were the interrelation of the stored information, the quality of the user interface and the volume of data which was required to be processed. The customer provided as input to the study an informal requirements document. We attempt to provide in this paper an outline of the steps involved in development of the eventual formal specification. It is important to stress at the outset that we view the task of constructing such a specification to be an iterative process, involving several attempts at construction of a model for the system interspersed with frequent dialogues with the customer to clarify details which are ambiguous or undefined in the initial requirements document, and frequent redrafting to clarify the structure of the document.

At an early stage in the analysis it became clear that the CAVIAR system consisted of several highly independent subsystems. Each subsystem records important relationships within the complete system and these separate subsystems are themselves related according to some simple rules. Most of the operations to be provided in the user interface can be explained as functions which transform one particular subsystem only, leaving the others invariant. These observations led to the decision to first define the subsystems in isolation and then to describe the complete system by combining the definitions of the subsystems. Once this decision had been taken, it also became clear that each of the individual subsystems, when viewed at an appropriate level of abstraction, was a particular instance of a general structure. From this vantage point it was natural to specify each of the subsystems by "refining" a specification which describes the underlying general system.

The process of analysis as presented here begins with an identification of the sets which appear to be important from the customer's point of view. Next the relationships between these sets are investigated and a preliminary classification of the subsystems follows. The third phase consists of developing an appropriate general mathematical structure in which to place these subsystems. Various ways of specialising (restricting) the general structure are then investigated and particular subsystems are modelled by instantiation. Finally the subsystem models are combined.

2. Identification of the Basic Sets

We now present a brief account of the existing system, emphasizing the important concepts in boldface. **Visitors** come to the site to attend **meetings** and/or consult Company employees. A visitor may require a **hotel reservation** and/or **transport reservation**. Each meeting is also required to take place in a designated **conference room**, at a certain **time**. A meeting may require the use of a **dining room** for lunch, on a particular date. Booking a dining room requires **lunch information** including the number of places needed. Each conference room booking requires **session information** about resources required for use in the meeting, e.g., viewgraphs, projectors. The main operations required at the user interface can briefly be described as facilities for booking, changing and cancelling the use of resources. We list below the sets together with the names that we shall adopt for referring to them.

Set	Name
Meetings	M
Visitors	V
Conference Rooms	CR
Dining Rooms	DR
Lunch Information	LI
Session Information	SI
Hotel reservation	HR
Transport reservation	TR

The informal interpretation of these sets is straight forward and for the purpose of this specification no further detail is necessary. Note that the question of modelling time remains to be resolved; at this point we simply observe that hotel reservations are made for particular dates, transport reservations are made for certain *times* on particular dates, and conference room bookings are made for *sessions* on particular dates. We shall not specify the term *session* further apart from noting that a date is always associated with a session; it could, for example, denote complete mornings or afternoons, or hourly or half-hourly intervals, depending on the way conference rooms are allocated.

The notion of time and the relationship between the different units of time used within the system can be formalised by asserting the existence of three sets as follows:

Date
Session
Time

together with two total functions

date-of-session : Session \rightarrow Date
date-of-time : Time \rightarrow Date.

3. The Subsystems of CAVIAR

The first approach to a mathematical model stems from the realisation that several of the sets listed above can be viewed as resources and other sets viewed as users of those resources. We can identify the following subsystems of CAVIAR in this framework (i. e., Resource-User systems). Observe that in different subsystems the same set may appear in differing roles.

System	Resources	Users
CR-M	Conference rooms	Meetings
DR-M	Dining rooms	Meetings
M-V	Meetings	Visitors
HR-V	Hotel reservations	Visitors
TR-V	Transport reservations	Visitors

Once we have made this mathematical abstraction it seems worthwhile to develop a general theory of such resource-user systems for the following reasons:


1. A specification of such a general system would be more useful as part of a "specification library" than a specific instance of such a system. Re-usability is much more likely to be achieved by having generic specifications available which can be *instantiated* to provide particular systems.
2. Particular subsystems of the general system can be constructed as *special cases* of the general specification in various ways. This will amply repay care and time spent on the general case. Furthermore, such instantiation may well result in a more compact implementation.

4. A General Resource-User System

We consider a system parametrised over three sets;

$$[T, R, U]$$

Informally, T is to be thought of as a set of *time slots*, R is a set of *resources* and U is a set of *users*. We describe a general resource-user system as a function from T to the set of relations between R and U . Thus we have a rather general framework: for each time slot $t \in T$, some users are occupying or using some resources. The set T will later be instantiated with different sets in the various applications. Notice that considering *relations* between R and U allows us the possibility of a user occupying several different resources simultaneously, as is shown informally in the following diagram:



Formally, the structure we are describing is captured by a function of type

$$T \rightarrow (R \leftrightarrow U)$$

We shall now incorporate this into a *schema definition*. This schema is parametrised over the sets T , R and U , and contains some useful ancillary concepts in addition to the function τ above which will be useful in later analysis. In Z specifications it is common to introduce such derived components: as specifiers of software we are neither in the position of a pure mathematician looking for a particularly sparse set of concepts and axioms with which to define a mathematical structure, nor are we in the position of an implementor trying to minimise storage. The component *in-use*, which gives the set of resources in use at any point of time, will be useful in contexts where

we are not concerned with the user component of the system state. The function *users*, which gives the users occupying resources at any point of time, will be used in situations where we do not require the information about resources. We also note that there may be occasions when we wish to consider the set of *inverse* relations generated by *ru*; we call this function *ur*.

$$\begin{array}{l}
 \text{R-U} \\
 \text{ru} \quad : T \rightarrow (R \leftrightarrow U) \\
 \text{in-use} : T \rightarrow \mathcal{P} R \\
 \text{users} \quad : T \rightarrow \mathcal{P} U \\
 \text{ur} \quad : T \rightarrow (U \leftrightarrow R) \\
 \\
 \forall t: T \cdot \\
 \quad \text{in-use}(t) = \text{dom}(\text{ru}(t)) \wedge \\
 \quad \text{users}(t) = \text{rng}(\text{ru}(t)) \wedge \\
 \quad \text{ur}(t) = (\text{ru}(t))^{-1}
 \end{array}$$

The *initial state* of this system is defined by making $\text{ru}(t)$ the empty relation for each t .

$$\text{Init-R-U} \triangleq [\text{R-U} \mid \text{rng}(\text{ru}) = \{ \{ \} \}]$$

Our first theorem proves that such an initial state is reasonable and assures us of the consistency of the definition of R-U.

Theorem 1.

$$\vdash \exists \text{R-U} \cdot \text{Init-R-U}$$

In the interests of readability we have not given proofs of theorems stated in this paper.

We continue by defining the appropriate operations for this structure. The first step is to identify *commonalities*. For our purposes, the operations that we wish to consider on this structure are concerned with making a new booking, i. e., adding a new pair (r, u) to an existing relation at some time t , cancelling an existing booking, i. e., removing such an (r, u) pair, or modifying in some other way the relation that exists at some particular time. In fact we shall be a little more general and define a class of operations on R-U which allows the image of a set of time values to be altered. This is because we anticipate such operations as booking a conference room

for a meeting which lasts for several time slots. Of course a booking which involves only a single time slot is a particular case.

Thus we may summarise the common part of all the operations as follows. Their description involves: a state before, R-U which introduces ru, in-use, users and ur; a state after, R-U' which introduces ru', in-use', users and ur; a set of time values, t? which denotes an input. The operations always leave the function ru unchanged except for times in t?. Formally this is captured by

$$\begin{array}{l}
 \Delta R-U \\
 R-U \\
 R-U' \\
 t? : P \ T \\
 \hline
 t? \triangleleft ru' = t? \triangleleft ru
 \end{array}$$

We now have a successful booking operation defined as follows

$$\begin{array}{l}
 R-U-Book \\
 \Delta R-U \\
 r? : R \\
 u? : U \\
 \hline
 \forall t : t? \cdot \\
 (r?, u?) \notin ru(t) \wedge \\
 ru'(t) = ru(t) \cup \{ (r?, u?) \}
 \end{array}$$

Thus R-U-Book inherits all the properties of $\Delta R-U$. Furthermore, it takes two additional (input) parameters $r?:R$ and $u?:U$, and is constrained by a predicate which imposes a requirement on the input parameters and also further relates the before and after states.

Notice that we are making the predicate

$$\forall t : t? \bullet (r?, u?) \in ru(t)$$

a pre-condition for a successful booking. In fact, we can show that this condition is sufficient for performing a successful booking, i. e., if we are in a valid system state with the required input parameters of the correct type available and furthermore the above condition holds, then there exists a resulting valid system state which is related to the starting state according to the R-U-Book schema. Formally, this is the content of the following result:

Theorem 2

$$\begin{array}{l} R-U \wedge [t?: \mathbb{P} T; r?: R; u?: U \mid \forall t : t? \bullet (r?, u?) \in ru(t)] \\ \vdash \\ \exists R-U' \bullet R-U\text{-Book} \end{array}$$

A successful cancellation operation may be defined via

$\begin{array}{l} R-U\text{-Cancel} \\ \Delta R-U \\ r? : R \\ u? : U \end{array}$
$\begin{array}{l} \forall t : t? \bullet \\ (r?, u?) \in ru(t) \wedge \\ ru'(t) = ru(t) - \{ (r?, u?) \} \end{array}$

The pre-condition for successful cancellation is that the pair $(r?, u?)$ is related by $ru(t)$ for all time values t in $t?$; i. e., the following theorem holds.

Theorem 3

$$\begin{array}{l} R-U \wedge [t? : \mathbb{P} T; r? : R; u? : U \mid \forall t : t? \bullet (r?, u?) \in ru(t)] \\ \vdash \\ \exists R-U' \bullet R-U\text{-Cancel} \end{array}$$

So far we have only specified successful operations; thus these descriptions are incomplete. We could at this stage define robust operations by introducing appropriate error recovery machinery. In the interests of simplicity we shall not give a general treatment of errors; however we shall indicate in a later section how the descriptions of the operations at the user interface can be completed.

We shall define two further operations on this structure. The first involves deleting a resource and all use of that resource. This is an operation to be treated with caution: see Theorem 7 below.

$$\begin{array}{l}
 \text{R-U-De1-Res} \\
 \hline
 \Delta R-U \\
 r? : R \\
 \hline
 \forall t : t? \cdot \\
 \quad r? \in \text{dom } ru(t) \wedge \\
 \quad ru'(t) = \{ r? \} \triangleleft ru(t)
 \end{array}$$

Informally, this operation may be described as follows. Consider each element t in $t?$ and the corresponding relation $ru(t)$ in turn. All elements $(r?, u)$ are to be removed from $ru(t)$.

Theorem 4

$$\begin{array}{l}
 R-U \wedge [t? : P \ T; r? : R \mid \forall t : t? \cdot r? \in \text{dom } ru(t)] \\
 \vdash \\
 \exists R-U' \cdot R-U\text{-De1-Res}
 \end{array}$$

Corresponding to deleting a resource there is an operation which, given a user value $u?$, deletes all pairs $(r, u?)$ from the relations associated with time values in $t?$. This is defined as follows:

$$\begin{array}{l}
 \text{R-U-De1-User} \\
 \hline
 \Delta R-U \\
 u? : U \\
 \hline
 \forall t : t? \cdot \\
 \quad u? \in \text{rng } ru(t) \wedge \\
 \quad ru'(t) = ru(t) \triangleright \{ u? \}
 \end{array}$$

Theorem 5

$$\begin{array}{l}
R-U \wedge | t? : P T; u? : U | \forall t : t? \cdot u? \in \text{rng } ru(t) \] \\
\vdash \\
\exists R-U' \cdot R-U\text{-Del-User}
\end{array}$$

So far we have listed theorems that a specifier is obliged to prove; viz the result that the initial state satisfies the required definition (and therefore that the specification is consistent) and the theorems that explicitly give the pre-conditions for each operation.

For the specifications that we shall develop from now on these theorems have been omitted in the interests of brevity.

In addition to these obligatory results, there are other "optional" theorems that are a consequence of the specification, and which often give insight into the structure being developed.

Two such results for our system are as follows:

Theorem 6

$$R-U\text{-Book} ; R-U\text{-Cancel} \quad \vdash \quad ru' = ru.$$

Informally, this theorem states that if we make a booking and follow it immediately by a cancellation using the same input parameters, then the state of the system does not change.

Theorem 7

$$\begin{array}{l}
R-U\text{-Del-Res} \\
\vdash \\
\text{in-use}' = \text{in-use} \circ (\lambda t:t? \cdot \text{in-use}(t) - \{r?\}) \wedge \\
\text{users}' = \text{users} \circ \\
\quad (\lambda t:t? \cdot \text{users}(t) - \{u : U \mid ur(t)\{\{u\}\} = \{r?\} \})
\end{array}$$

This theorem makes precise the informal comment made earlier about the need for caution with the R-U-Del-Res operation. This theorem shows that resources are removed from the system structures, which we do expect, but furthermore the operation can also remove existing users.

There is a similar result concerning the R-U-Del-User operation.

5. Specialisation of the General R-U System

We shall now specialise the general R-U system into particular classes of the system. These specialisations are motivated by the observation that for some of the instances listed earlier, at any given time a resource may be related to only one user, or a user may occupy only one resource, or both.

5.1 An R-U system where resources cannot be shared

The first case we define is the class where each resource may be utilised by at most one user, but each user may occupy several resources. We denote this system by $R \succ U$ (where " \succ " is just a character in the name) and define it formally by

$$R \succ U \triangleq [R-U \mid \text{rng}(ru) \subseteq R \rightarrow U]$$

The initial state of this system is given by the same condition as for Init-R-U ; thus we have

$$\text{Init-}R \succ U \triangleq [R \succ U \mid \text{rng}(ru) = \{ \}]$$

All operations are described in terms of

$$\Delta R \succ U \triangleq R \succ U \wedge R \succ U'$$

The operations on this system may be defined as special cases of the general operations for R-U. We first consider the booking operation.

$$R \succ U\text{-Book} \triangleq \Delta R \succ U \wedge [R-U\text{-Book} \mid \forall t: t? \cdot r? \notin \text{dom } ru(t)]$$

The qualifying predicate is included in indicate that there is a further pre-condition for booking a resource in a $R \succ U$ system.

We now have two parts to the pre-condition for this operation; firstly this qualifying predicate, and secondly the pre-condition arising from R-U-Book. In fact the former implies the latter, as is easily checked.

The cancellation operation is defined as follows:

$$R \succ U\text{-Cancel} \triangleq R-U\text{-Cancel} \wedge \Delta R \succ U$$

On considering the two deletion operations defined for R-U, we observe that R-U-Del-Res is equivalent to a cancellation in our present context, because the resource is associated with only one user. We therefore need only the operation which deletes a user.

$$R\triangleright U\text{-Del-User} \triangleq R\text{-U-Del-User} \wedge \Delta R\triangleright U$$

5.2 An R-U system where each user may occupy at most one resource

The second case we define is the class where each user may occupy at most one resource but resources may be shared amongst users. We denote this system by R<U and define it formally by

$$R<U \triangleq [R\text{-U} \mid \text{rng}(ur) \subseteq U \rightarrow R]$$

The initial state of this system is also given by the predicate for Init-R-U. We have

$$\text{Init-R}<U \triangleq [R<U \mid \text{rng}(ru) = \{ \}]$$

The operations are described in terms of

$$\Delta R<U \triangleq R<U \wedge R<U$$

We now define the booking operation for the system.

$$R<U\text{-Book} \triangleq \Delta R<U \wedge [R\text{-U-Book} \mid \forall t : t? \cdot u? \notin \text{rng } ru(t)]$$

As before, a qualifying predicate is needed and again as before the constraint given here implies the earlier pre-condition for the general R-U-Book operation.

The cancellation operation is defined as follows:

$$R<U\text{-Cancel} \triangleq R\text{-U-Cancel} \wedge \Delta R<U$$

On considering the two deletion operations defined for R-U, we observe that this time R-U-Del-Res is equivalent to a cancellation in our present context, because a user may be associated with only one resource. We therefore need only the operation which deletes a resource.

$$R<U\text{-Del-Res} \triangleq R\text{-U-Del-Res} \wedge \Delta R<U$$

5.3 An R-U system where a user occupies at most one nonsharable resource

The third and last specialisation we define shares all the properties of the systems defined in the preceding two sections. It is therefore defined as the conjunction of the two schemas above. In this system each user may occupy at most one resource and each resource may be occupied by at most one user. Formally we have

$$R \sqsubseteq U \hat{=} R \triangleright U \wedge R \triangleleft U$$

The initial state of this system is clearly defined by

$$\text{Init-}R \sqsubseteq U \hat{=} [R \sqsubseteq U \mid \text{rng}(ru) = \{ \}]$$


The operations on this system are given by the conjunction of the operations defined for each of the two earlier systems. For this system we require only the booking and cancellation operations. Thus we have

$$R \sqsubseteq U\text{-Book} \hat{=} R \triangleright U\text{-Book} \wedge R \triangleleft U\text{-Book}$$

$$R \sqsubseteq U\text{-Cancel} \hat{=} R \triangleright U\text{-Cancel} \wedge R \triangleleft U\text{-Cancel}$$

5.4 The specification library

We have now constructed four specifications which might be considered to form the nucleus of a specification library for resource-user systems. We may summarise the relationships between the four classes of system schematically as follows:



6. Classification and Instantiation

6.1 Some laws for CAVIAR

In this section, in order to illustrate the clarification process which took place during requirements analysis, we list some observations about the CAVIAR system which emerged during dialogue with the customer. We formalise the important constraints as laws which need to be taken account into account in the development which follows.

1. At any time a conference room is associated with only one meeting.
2. At any time a meeting may be associated with more than one conference room.

Law 1 is reasonably obvious: it would be difficult to hold more than one meeting in a given room. Law 2 is not obvious: it was unclear from the informal description whether or not a meeting could occupy more than one room. In fact the customer believed initially that a meeting could only take up one room, but a counter-example was found amongst the supporting documentation.

3. At any time a meeting is associated with only one dining room.
4. At any time participants from several meetings can occupy the same dining room.

These laws followed from the informal information provided that all visitors in a particular meeting would go to lunch in the same dining room. It was further established that all seats in a dining room were treated as indistinguishable, so further meetings could be accommodated if enough seats were available. Further clarification was necessary regarding lunch times: it transpired that there were "early" and "late" lunches; however this was handled by "doubling up" each dining room. For example, a booking would be made for "DR 1, early" and this was a different dining room from "DR 1, late."

5. At any time a visitor is associated with only one meeting.
6. At any time a meeting may involve several visitors.

Law 5 had to be checked out with the customer.

7. At any time a hotel room is associated with only one visitor and vice versa.
8. At any time a transport reservation is associated with only one visitor and vice versa.

Law 7 was natural, but law 8 was less so. It was established that even if the transport department decided to use a minibus, a separate transport reservation would be issued to each visitor.

6.2 Matching system with models

In this section we first consider each CAVIAR subsystem in turn and match it to the appropriate model. In fact we have enough structure available to define two subsystems directly and we do this in the remainder of this section.

- (1) We first consider the conference room - meeting system CR-M.

From laws 1 and 2 we see that CR-M is an instance of the $R \times U$ subsystem.

- (2) The dining room - meeting subsystem DR-M.

Applying laws 3 and 4 we find that DR-M is an instance of $R \times U$.

However this system does not contain any information about numbers of seats or the lunch details, so we will need to extend this system later.

- (3) The meeting - visitor subsystem M-V.

From laws 5 and 6 M-V is an instance of $R \times U$.

However we have not documented the fact that meetings have to be created before visitors can be attached to them; this will also be done later.

- (4) The hotel reservation - visitor subsystem HR-V, and the transport reservation - visitor subsystem TR-V, both have the property that each resource is occupied by only one user and vice versa. Therefore both these systems are instances of $R \times U$.

In fact this model is sufficient to define HR-V and TR-V completely, by *instantiation*, as we now show.

6.3 The hotel reservation subsystem - HR-V

We define HR-V as follows:

$$\text{HR-V} \triangleq \text{REU}_{\text{HR-V}}[\text{Date}, \text{HR}, \text{V}]$$

This object is a decorated instance of the REU schema, with its parameter sets instantiated by the sets Date, HR and V introduced in section 2. To be more explicit, the definition above is shorthand for the following:

HR-V	
$ru_{\text{HR-V}}$	$: \text{Date} \rightarrow (\text{HR} \leftrightarrow \text{V})$
$in\text{-}use_{\text{HR-V}}$	$: \text{Date} \rightarrow \mathbb{P} \text{HR}$
$users_{\text{HR-V}}$	$: \text{Date} \rightarrow \mathbb{P} \text{V}$
$ur_{\text{HR-V}}$	$: \text{Date} \rightarrow (\text{V} \leftrightarrow \text{HR})$
<hr/>	
$\text{rng}(ru_{\text{HR-V}}) \subseteq \text{HR} \leftrightarrow \text{V} \wedge$	
$\text{rng}(ur_{\text{HR-V}}) \subseteq \text{V} \leftrightarrow \text{HR} \wedge$	
$(\forall t: \text{Date}; r: \text{HR} \cdot$	
$r \in in\text{-}use_{\text{HR-V}}(t) \Leftrightarrow r \in \text{dom}(ru_{\text{HR-V}}(t)) \wedge$	
$(\forall t: \text{Date}; u: \text{V} \cdot$	
$u \in users_{\text{HR-V}}(t) \Leftrightarrow u \in \text{ran}(ru_{\text{HR-V}}(t)) \wedge$	
$(\forall t: \text{Date} \cdot ur_{\text{HR-V}}(t) = (ru_{\text{HR-V}}(t))^{-1})$	

Thus each component of the schema is given the decoration in the definition, and each occurrence of the parametrised sets is instantiated as shown above. From now on we shall use such decoration without further comment.

The initial state of HR-V is given by

$$\text{Init-HR-V} \triangleq \text{Init-REU}_{\text{HR-V}}[\text{Date}, \text{HR}, \text{V}]$$

and the operations are given by

$$\text{Book-Hotel-Room}_0 \triangleq \text{REU-Book}_{\text{HR-V}}[\text{Date}, \text{HR}, \text{V}]$$

and

$$\text{Cancel-Hotel-Room}_0 \triangleq \text{REU-Cancel}_{\text{HR-V}}[\text{Date}, \text{HR}, \text{V}]$$

6.4 The transport reservation subsystem - TR-V

This subsystem is essentially the same as the HR-V subsystem except for the parametrisation. The instances of the parameters are denoted respectively Time , TR and V , where once again the sets TR and V are as in section 2. We shall not specify the set Time further, except to repeat that it contains a Date component (see section 2). Thus we have

$$\text{TR-V} \hat{=} \text{REU}_{\text{TR-V}}[\text{Time}, \text{TR}, \text{V}]$$

with initial state given by

$$\text{Init-TR-V} \hat{=} \text{Init-REU}_{\text{TR-V}}[\text{Time}, \text{TR}, \text{V}]$$

and operations given by

$$\text{Book-Transport}_0 \hat{=} \text{REU-Book}_{\text{TR-V}}[\text{Time}, \text{TR}, \text{V}]$$

and

$$\text{Cancel-Transport}_0 \hat{=} \text{REU-Cancel}_{\text{TR-V}}[\text{Time}, \text{TR}, \text{V}]$$

7. The Meeting Attendance Subsystem

We now turn our attention to what is necessary in order to complete a model for M-V. Booking and cancelling operations have been defined already but so far we have not taken account of the fact that before bookings can be made the system has to "create" meetings. The question of exactly which objects are "currently defined" at any particular time is important because in several cases only those objects known to the system (i. e., those objects that have been created but not yet destroyed) can book resources, etc.

7.1 A pool system

We can model this situation with a simple structure which we term a Pool. This schema is parametrised over the set T and an arbitrary set X . There are only two operations to be defined; namely those that add an object to, and delete an object from, the pool, over a specified time period.

Formally we have

$$\begin{array}{c}
 [T, X] \\
 \text{Pool} \\
 \text{exists : } T \rightarrow P X
 \end{array}$$

with initial state given by

$$\text{Init-Pool} \hat{=} [\text{Pool} \mid \text{rng}(\text{exists}) = \{ \{ \} \}]$$

For later use we define

$$\equiv \text{Pool} \hat{=} [\Delta \text{Pool} \mid \text{Pool}' = \text{Pool}]$$

Given

$$\Delta \text{Pool} \hat{=} \text{Pool} \wedge \text{Pool}'$$

The operations are given by

Create Δ Pool $t? : P \ T$ $x? : X$
$exists' = exists \bullet (\lambda t : t? \bullet exists(t) \cup \{ x? \})$

and

Destroy Δ Pool $t? : P \ T$ $x? : X$
$exists' = exists \bullet (\lambda t : t? \bullet exists(t) - \{ x? \})$

We could have included in the Create operation the pre-condition that the object $x?$ not already exist for any of the times in $t?$. However we make a deliberate decision here to omit this - having in mind the situation where an object may already exist for some of the times in $t?$ and its existence needs to be extended to all of $t?$. A similar remark applies to the Destroy operation.

7.2 The meeting - visitor subsystem

To construct the model for the M-V system we combine the Pool and R<U structures.

M-V $R\langle U_{M-V} [Session, M, V]$ $Pool_M [Session, M]$
$\forall t : T \bullet$ $in-use_{M-V}(t) \subseteq exists_M(t)$

Thus we have combined an M-V instance of an R \times U system and a meeting instantiation of a Pool system (with the parameter sets as shown). The predicate assures that visitors can only attend existing meetings.

The initial state is given by

$$\text{Init-M-V} \hat{=} \text{Init-R}\times\text{U}_{\text{M-V}}[\text{Session}, \text{M}, \text{V}] \wedge \text{Init-Pool}_{\text{M}}[\text{Session}, \text{M}]$$

We now define the operations on M-V in terms of

$$\Delta\text{M-V} \hat{=} \text{M-V} \wedge \text{M-V}'$$

The first operation is concerned with adding a visitor to a meeting.

$$\begin{aligned} \text{Add-Visitor-to-Meeting}_0 \hat{=} \\ \Delta\text{M-V} \wedge \exists \text{Pool}_{\text{M}}[\text{Session}, \text{M}] \wedge \text{R}\times\text{U-Book}_{\text{M-V}}[\text{Session}, \text{M}, \text{V}] \end{aligned}$$

When an operation is "promoted" in this way, its new pre-condition is determined as follows: the "old" pre-condition (i. e., that arising from its definition) must be conjoined with a further predicate which arises from the new invariant of the larger state. Here, for example, the pre-condition for the earlier booking operation is given in section 5.2: namely

$$\forall t : t?_{\text{M-V}} \cdot u?_{\text{M-V}} \notin \text{rng}(ru_{\text{M-V}}(t))$$

and this must be conjoined with

$$\forall t : t?_{\text{M-V}} \cdot r?_{\text{M-V}} \in \text{exists}_{\text{M}}(t).$$

This second predicate is a consequence of the M-V invariant.

Thus the complete pre-condition for the Add-Visitor-to-Meeting operation is given by

$$\forall t : t?_{\text{M-V}} \cdot u?_{\text{M-V}} \notin \text{rng}(ru_{\text{M-V}}(t)) \wedge r?_{\text{M-V}} \in \text{exists}_{\text{M}}(t)$$

which states that the visitor ($u?_{\text{M-V}}$) is not already attending a meeting at that time and that the meeting he is going to attend actually exists.

The second operation removes a visitor from a meeting.

$$\text{Remove-Visitor-from-Meeting}_0 \triangleq \\ \Delta M-V \wedge \exists \text{Pool}_H[\text{Session}, M] \wedge R\langle U\text{-Cancel}_{M-V}[\text{Session}, M, V]$$

It is easy to check that the pre-condition for the Remove-Visitor-from-Meeting operation is simply the predicate which is inherited from the initial R-U-Cancel operation; namely

$$\forall t : t?_{M-V} \cdot (r?_{M-V}, u?_{M-V}) \in ru_{M-V}(t)$$

We now define the operations which create and cancel meetings as follows:

$$\text{Create-Meeting}_0 \triangleq \\ \Delta M-V \wedge \exists R\langle U_{M-V}[\text{Session}, M, V] \wedge \text{Create}_H[\text{Session}, M]$$

For the creation there is no pre-condition.

Cancel-Meeting_0 $\Delta M-V$ $R\langle U\text{-Del-Res}_{M-V}[\text{Session}, M, V]$ $\text{Destroy}_H[\text{Session}, M]$
$t?_H = t?_{M-V} \wedge$ $x?_H = r?_{M-V}$

The pre-conditions for cancelling a meeting arise from the original R-U-Del-Res operation, i. e., that

$$\forall t : t?_{M-V} \cdot r?_{M-V} \in \text{dom}(ru_{M-V}(t))$$

and secondly from the identifications required for the input parameters.

8. The Meeting Resource Subsystems

We are left with the systems CR-M and DR-M to define. We observe that both of these have further information associated with the resource-user relationship, so in order to capture this facet in our model we introduce the concept of a *diary system*.

8.1 A diary system

The diary is to record information about some elements of a set. We denote the set in question by X and the associated information by I_X . For each t , the set of elements of X for which we have information is defined as $\text{recorded}(t)$. Once again this system is dependent on time, T .

$$\begin{array}{c}
 [T, X, I_X] \\
 \text{Diary} \text{ ---} \\
 \boxed{
 \begin{array}{l}
 \text{info} \quad : T \rightarrow (X \leftrightarrow I_X) \\
 \text{recorded} : T \rightarrow \mathcal{P} X \\
 \\
 \forall t : T \bullet \text{recorded}(t) = \text{dom}(\text{info}(t))
 \end{array}
 }
 \end{array}$$

with initial state given by

$$\text{Init-Diary} \hat{=} [\text{Diary} \mid \text{rng}(\text{info}) = \{ \{ \} \}]$$

The two operations to be defined both involve a change over a particular time period. Note that we are motivated to make this definition in order to maintain compatibility with existing systems. Formally we define

$$\Delta \text{Diary} \hat{=} \text{Diary} \wedge \text{Diary}' \wedge [t? : \mathcal{P} T]$$

$$\begin{array}{c}
 \text{Add} \text{ ---} \\
 \boxed{
 \begin{array}{l}
 \Delta \text{Diary} \\
 x? : X \\
 i? : I_X \\
 \\
 (\forall t : t? \bullet x? \notin \text{recorded}(t)) \wedge \\
 \text{info}' = \text{info} \bullet (\lambda t : t? \bullet \text{info}(t) \bullet \{ x? \mapsto i? \})
 \end{array}
 }
 \end{array}$$

The complementary erasure operation would remove one element (and the information associated with it) from $\text{info}(t)$. However we note that this is a special case of the following more powerful operation.

$\frac{\text{Erase} \quad \Delta \text{Diary} \quad x? : T \rightarrow P X}{\text{dom}(x?) = t? \wedge (\forall t : t? \cdot x?(t) \subseteq \text{recorded}(t?)) \wedge \text{info}' = \text{info} \oplus (\lambda t : t? \cdot x?(t) \triangleleft \text{info}(t))}$
--

8.2 The conference room booking subsystem

We are now in a position to fully specify the subsystem CR-M, by instantiation as follows:

$\frac{\text{CR-M} \quad R \times U_{\text{CR-M}}[\text{Session}, \text{CR}, \text{M}] \quad \text{Diary}_{\text{CR}}[\text{Session}, \text{CR}, \text{SI}]}{\text{in-use}_{\text{CR-M}} = \text{recorded}_{\text{CR}}}$
--

with initial state given by

$$\text{Init-CR-M} \hat{=} \text{Init-R} \times U_{\text{CR-M}}[\text{Session}, \text{CR}, \text{M}] \wedge \text{Init-Diary}_{\text{CR}}[\text{Session}, \text{CR}, \text{SI}]$$

It would be more correct to regard the session information SI as being related to a meeting rather than a conference room. The reason for associating SI with conference rooms is that it contains information which is issued to the department supplying equipment for meetings, and they are concerned with the venue rather than what is to take place there.

The operations that we require for CR-M are given below. Information is recorded about each resource when it is booked, and must be erased when a cancellation takes place. The definitions use

$$\Delta CR-M \triangleq CR-M \wedge CR-M'$$

Book-Conf-Room_0 $\Delta CR-M$ $R \triangleright U\text{-Book}_{CR-M}[\text{Session}, CR, M]$ $\text{Add}_{CR}[\text{Session}, CR, SI]$
<hr style="width: 50%; margin: 0 auto;"/> $t?_{CR-M} = t?_{CR} \wedge$ $r?_{CR-M} = x?_{CR}$

$\text{Cancel-Conf-Rooms}_0$ $\Delta CR-M$ $R \triangleright U\text{-Del-User}_{CR-M}[\text{Session}, CR, M]$ $\text{Erase}_{CR}[\text{Session}, CR, SI]$
<hr style="width: 50%; margin: 0 auto;"/> $t?_{CR-M} = t?_{CR} \wedge$ $(\forall t: t?_{CR-M} \cdot x?_{CR}(t) = u\Gamma_{CR-M}(t) (\{u?_{CR-M}\}))$

The cancellation operation here deletes all conference rooms associated with a particular meeting over the specified time period. This is the operation which is most compatible with the Cancel-Meeting operation defined for M-V. However, if required, we could also define the operation that cancels just one conference room-meeting pairing.

8.3 The dining room booking subsystem

The final subsystem that we need to consider is DR-M.

The analysis so far does not take account of the fact that dining rooms have a finite capacity, so we need to extend our model. We suppose that we have been given a function

$$\text{max-no} : \text{DR} \rightarrow \mathbf{N}$$

which records this capacity and we record the number of seats in each dining room which have been reserved already.

The DR-M system is defined formally as follows:

$\begin{array}{l} \text{DR-M} \\ \hline \text{R}\langle\text{U}\rangle_{\text{DR-M}}[\text{Date}, \text{DR}, \text{M}] \\ \text{Diary}_{\text{DR}}[\text{Date}, \text{M}, \text{LI}] \\ \text{rsvd} : \text{T} \rightarrow (\text{DR} \rightarrow \mathbf{N}) \\ \hline \text{users}_{\text{DR-M}} = \text{recorded}_{\text{DR}} \wedge \\ (\forall \text{t:Date} \bullet \text{dom}(\text{rsvd}(\text{t})) = \text{in-use}_{\text{DR-M}}(\text{t}) \wedge \\ \quad (\forall \text{r:in-use}_{\text{DR-M}}(\text{t}) \bullet \text{rsvd}(\text{t})(\text{r}) \leq \text{max-no}(\text{r})) \\) \end{array}$
--

Observe that in this case information is associated with each user, and therefore the diary system takes M as its main parameter. Dining rooms that are in use have a number of seats reserved, and this number has to be within the dining room's capacity.

The initial state of DR-M is given by

$$\text{Init-DR-M} \hat{=} \text{Init-R}\langle\text{U}\rangle_{\text{DR-M}}[\text{Date}, \text{DR}, \text{M}] \wedge \text{Init-Diary}_{\text{DR}}[\text{Date}, \text{M}, \text{LI}]$$

The two operations that we require for this structure are *booking* a (number of seats in a) dining room and *cancelling* a lunch booking for a particular meeting. In normal circumstances, a resource (dining room) will not be subject to being taken out of service (although this occurrence is clearly easy to model if required).

Both these operations leave $rsvd$ unchanged for time values outside the period in question; we make this part of the operation invariant.

$\Delta DR-M$ $\Delta R \times U_{DR-M}[Date, DR, M]$ $\Delta Diery_{DR}[Date, M, LI]$ $amount? : T \rightarrow N$
$t?_{DR-M} = t?_{DR} \wedge$ $dom(amount?) = t?_{DR-M} \wedge$ $t?_{DR-M} \triangleleft rsvd' = t?_{DR-M} \triangleleft rsvd$

$Book-Dining-Room_0$ $\Delta DR-M$ $R \times U-Book_{DR-M}[Date, DR, M]$ $Add_{DR}[Date, M, LI]$
$x?_{DR} = u?_{DR-M} \wedge$ $(\forall t : t?_{DR-M} \cdot$ $\quad rsvd(t)(r?_{DR-M}) + amount?(t) \leq max-no(r?_{DR-M}) \wedge$ $\quad rsvd'(t) = rsvd(t)$ $\quad \bullet \{ r?_{DR-M} \mapsto rsvd(t)(r?_{DR-M}) + amount?(t) \}$ $)$

$Cancel-Dining-Room_0$ $\Delta DR-M$ $R \times U-Cancel_{DR-M}[Date, DR, M]$ $Erase_{DR}[Date, M, LI]$
$(\forall t : t?_{DR-M} \cdot$ $\quad x?_{DR}(t) = \{ u?_{DR-M} \} \wedge$ $\quad rsvd'(t) = rsvd(t)$ $\quad \bullet \{ r?_{DR-M} \mapsto rsvd(t)(r?_{DR-M}) - amount?(t) \}$ $)$

8.4 The visitor pool - V-P

From the informal requirements we find that visitors must be "legitimate" before they are allowed to attend meetings or have resources booked on their behalf. This requirement is easily met by introducing a visitor Pool structure, with actual parameters Date and V. Thus we define V-P as

$$V-P \hat{=} Pool_V[Date, V]$$

with initial state given by

$$Init-V-P \hat{=} Init-Pool_V[Date, V]$$

The operations that we require on this structure are simply those of creation and destruction of visitors. Formally we have

$$Create-Visitor_0 \hat{=} Create_V[Date, V]$$


and

$$Destroy-Visitor_0 \hat{=} Destroy_V[Date, V]$$

8.5 The construction process

In this section we summarise the constructions we have used to build the individual CAVIAR components.

In sections 7 and 8 we added *pool* and *diary* components to our basic library in section 5.4. We now have a library which consists of the 6 components R-U, R>U, R<U, R=U, Pool and Diary. We indicate in the following diagram how each subsystem has been constructed using components from the library.




9. The Complete CAVIAR System

We have now achieved our first goal of specifying all constituent subsystems of CAVIAR. We have yet to combine the subsystems into a coherent whole. This is now a comparatively easy task, once we have observed a few extra constraints.

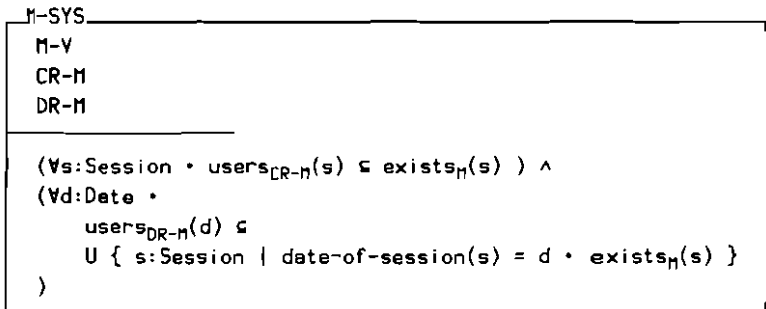
9.1 Combining subsystems to form the system state

We define the visitor part of the system as follows:



The invariant states that visitors that have hotel or transport reservations must be known.

The meeting part of the system is defined by



The invariant states that meetings which are occupying conference rooms or dining rooms must be known to the system at that time.

These two subsystems are now combined to form the CAVIAR system.

CAVIAR
V-SYS
M-SYS
$\forall s : \text{Session} \cdot \text{users}_{M-V}(s) \subseteq \text{exists}_V(\text{date-of-session}(s))$

Informally, the invariant states that all visitors who are attending meetings must be known to the system.

The initial state of the system is given by the conjunction of all the initialisations. It is easy to verify that this conjunction satisfies the invariant.

$$\text{Init-CAVIAR} \triangleq \text{Init-HR-V} \wedge \text{Init-TR-V} \wedge \text{Init-M-V} \wedge \\ \text{Init-CR-M} \wedge \text{Init-DR-M} \wedge \text{Init-V-P}$$

9.2 Operations on CAVIAR

The operations on CAVIAR may be divided naturally into three groups.

9.2.1 Operations which involve meetings only

These operations are concerned with M-SYS only and leave V-SYS unchanged. We denote this by

$$M\text{-OP} \triangleq \Delta\text{CAVIAR} \wedge \equiv V\text{-SYS}$$

where

$$\Delta\text{CAVIAR} \triangleq \text{CAVIAR} \wedge \text{CAVIAR}'$$

and

$$\equiv V\text{-SYS} \triangleq [V\text{-SYS} \wedge V\text{-SYS}' \mid V\text{-SYS} = V\text{-SYS}']$$

(Note: in the following similar definitions of $\equiv\text{CR-M}$, $\equiv\text{DR-M}$, etc. are omitted.)

The first operation is to construct a meeting

$$\text{Create-Meeting} \triangleq \text{M-OP} \wedge \text{Create-Meeting}_0 \wedge \equiv \text{CR-M} \wedge \equiv \text{DR-M}$$

This operation has no pre-condition (there is no pre-condition for Create-Meeting_0), so it is total. The next operation is to cancel a meeting.

$$\text{Cancel-Meeting}_1 \triangleq \text{M-OP} \wedge \text{Cancel-Meeting}_0 \wedge \equiv \text{CR-M} \wedge \equiv \text{DR-M}$$

We can determine the pre-condition for this operation as follows: first we establish the constraint arising from the system invariant. The operation removes an element from exists_M so this element cannot be a user in CR-M or DR-M during the period $t?_M$. Formally, we require that

$$\forall t : t?_M \cdot r?_{M-U} \notin \text{users}_{\text{CR-M}}(t) \cup \text{users}_{\text{DR-M}}(\text{date-of-session}(t))$$

The second part of the pre-condition arises from the earlier pre-condition for Cancel-Meeting_0 . This is precisely

$$t?_M = t?_{M-U} \wedge x?_M = r?_{M-U} \wedge (\forall t : t?_{M-U} \cdot r?_{M-U} \in \text{dom}(r_{U-M}(t))).$$

We shall at this point fulfil the promise made in section 4.1: indicating how to define the corresponding total operation. This is formed by the *disjunct* of the successful operation with the schema which takes as its qualifying predicate the *negation* of the pre-condition established above.

Cancel-Meeting-Fail
$\equiv \text{CAVIAR}$ $t?_{M-U} : \mathbf{P} \text{ Session}$ $t?_M : \mathbf{P} \text{ Session}$ $x?_M : \mathbf{M}$ $r?_{M-U} : \mathbf{M}$
$(\exists t : t?_M \cdot$ $\quad r?_{M-U} \in \text{users}_{\text{CR-M}}(t) \cup \text{users}_{\text{DR-M}}(\text{date-of-session}(t)))$ $\vee t?_M \neq t?_{M-U}$ $\vee x?_M \neq r?_{M-U}$ $\vee (\exists t : t?_{M-U} \cdot r?_{M-U} \notin \text{dom}(r_{U-M}(t)))$

$$\text{Cancel-Meeting} \hat{=} \text{Cancel-Meeting}_1 \vee \text{Cancel-Meeting-Fail}$$

Informally, if the required pre-condition for the meeting cancellation is not satisfied, the system is unchanged. In practice we would require an appropriate error message to be output.

For the sake of brevity, we shall present the remainder of the operations without going through this process.

The next two operations add visitors to, and delete visitors from, a meeting.

$$\begin{aligned} \text{Add-Visitor-to-Meeting} \hat{=} \\ M\text{-OP} \wedge \text{Add-Visitor-to-Meeting}_0 \wedge \equiv \text{CR-M} \wedge \equiv \text{DR-M} \end{aligned}$$

$$\begin{aligned} \text{Remove-Visitor-from-Meeting} \hat{=} \\ M\text{-OP} \wedge \text{Remove-Visitor-from-Meeting}_0 \wedge \equiv \text{CR-M} \wedge \equiv \text{DR-M} \end{aligned}$$

The pre-conditions for these operations are straightforward to determine in the usual way and we shall omit them and also those for the remaining operations.

The next two operations deal with conference rooms.

$$\text{Book-Conf-Room} \hat{=} M\text{-OP} \wedge \equiv M\text{-V} \wedge \text{Book-Conf-Room}_0 \wedge \equiv \text{DR-M}$$

$$\text{Cancel-Conf-Room} \hat{=} M\text{-OP} \wedge \equiv M\text{-V} \wedge \text{Cancel-Conf-Room}_0 \wedge \equiv \text{DR-M}$$

We now have the two operations concerning dining rooms.

$$\text{Book-Dining-Room} \hat{=} M\text{-OP} \wedge \equiv M\text{-V} \wedge \equiv \text{CR-M} \wedge \text{Book-Dining-Room}_0$$

$$\text{Cancel-Dining-Room} \hat{=} M\text{-op} \wedge \equiv M\text{-V} \wedge \equiv \text{CR-M} \wedge \text{Cancel-Dining-Room}_0$$

There is one final operation to be defined in this section: namely the cancellation of both dining room and conference room(s) associated with a particular meeting. This is not the conjunct of the two cancellation operations already given because each of these leaves the components it is not acting on fixed. Hence we need a different operation defined by

$$\begin{aligned} \text{Cancel-Meeting-Arrangements} \hat{=} \\ M\text{-OP} \wedge \equiv M\text{-V} \wedge \text{Cancel-Conf-Room}_0 \wedge \text{Cancel-Dining-Room}_0 \end{aligned}$$

9.2.2 Operations which involve visitors only

This section contains operations which involve V-SYS only and leave M-SYS unchanged. We denote this group by

$$V-OP \triangleq \Delta CAVIAR \wedge \equiv M-SYS$$

The first pair of operations introduce visitors to and remove visitors from the visitor system.

$$\text{Create-Visitor} \triangleq V-OP \wedge \text{Create-Visitor}_0 \wedge \equiv HR-V \wedge \equiv TR-V$$

$$\text{Destroy-Visitor} \triangleq V-OP \wedge \text{Destroy-Visitor}_0 \wedge \equiv HR-V \wedge \equiv TR-V$$

The Caviar invariant induces the following pre-condition for the Destroy operation.

$$\begin{aligned} \forall t : t?_V \cdot x?_V \notin \text{users}_{HR-V}(t) \\ \cup \cup \{ t : \text{date-of-time}^{-1}(t?_V) \cdot \text{users}_{TR-V}(t) \} \\ \cup \cup \{ s : \text{date-of-session}^{-1}(t?_V) \cdot \text{users}_{M-V}(s) \} \end{aligned}$$

The two operations concerned with hotel rooms are as follows:

$$\text{Book-Hotel-Room} \triangleq V-OP \wedge \equiv V-P \wedge \text{Book-Hotel-Room}_0 \wedge \equiv TR-V$$

$$\text{Cancel-Hotel-Room} \triangleq V-OP \wedge \equiv V-P \wedge \text{Cancel-Hotel-Room}_0 \wedge \equiv TR-V$$

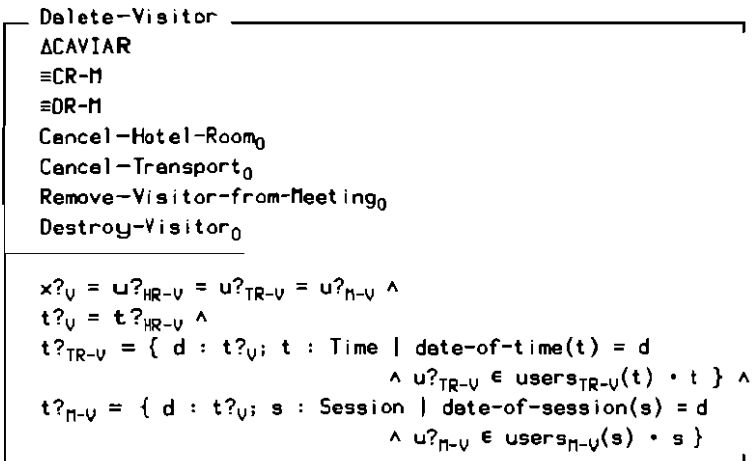
The two operations concerned with transport reservations are

$$\text{Book-Transport} \triangleq V-OP \wedge \equiv V-P \wedge \equiv HR-V \wedge \text{Book-Transport}_0$$

$$\text{Cancel-Transport} \triangleq V-OP \wedge \equiv V-P \wedge \equiv HR-V \wedge \text{Cancel-Transport}_0$$

9.2.3 A general visitor removal operation

Finally we define an operation which removes a visitor entirely from the system for a particular set of dates.



10. Conclusion

This specification has created a conceptual model for the CAVIAR system which provides a precise description of the system state and its external interface, together with an exact functional specification of every operation. The subtle inter-relationships between constituent subsystems are described in the predicates which constrain the combination of these subsystems, and these have been taken into account in the specification of the operations. The system designer can now concentrate on the important parts of the design task: namely selecting appropriate data structures and algorithms, without having to be simultaneously concerned with the complexity of subsystem interactions. This reflects the classical principle of *separation of concerns*.

It may be argued that a specification such as we have given above is a long way from an actual software product. Experience shows however that minimal effort is required to develop software once such a specification has been constructed. For example, in the case of CAVIAR, a Pascal implementation was constructed directly and quickly from the specification.

11. Acknowledgements

A formal specification of CAVIAR was given in 1981 by J.-R. Abrial. This work was carried out at the Programming Research Group at Oxford University in collaboration with B. Sufrin, T. Clement and one of the co-authors. T. Clement implemented a prototype version of the specification on a ITT-2020 computer in UCSD Pascal. J.-R. Abrial's original specification document listed most of the properties of the system that appear in this document, though the style of the presentation, the notation, and the conventions used in this paper have since been developed by members of the Programming Research Group.

We would like to thank J.-R. Abrial for his original contribution, I. Hayes for editing this paper and all those involved in helping with the project, particularly the personnel in the Visitor Services Department of STL, who willingly provided the team with information about the current manual system in operation at that time.

We would also like to thank Bernie Cohen, Tim Denvir and Tom Cox for their initial effort in setting up this collaborative effort between STL and the Programming Research Group and their continuing interest.

12. References and Related Work

1. Abrial, J.-R. The specification language Z: Basic library. *Oxford University Programming Research Group internal report*, (April 1980).
2. Morgan, C. C. Schemas in Z: A preliminary reference manual. *Oxford University Programming Research Group Distributed Computing Project report*, (March 1984).
3. Sufrin, B. A., Sørensen, I. H., Morgan, C. C., and Hayes, I. J. Notes for a Z Handbook. *Oxford University Programming Research Group internal report*, (July 1985).
4. Morgan, C. C., and Sufrin, B. A. Specification of the UNIX file system. *IEEE Transactions on Software Engineering*, Vol. 10, No. 2, (March 1984), pp. 128-142.
5. Hayes, I. J. Specification Case Studies. *Oxford University Programming Research Group Monograph, PRG-46*, (July 1985).

Z Reference Card
Mathematical Notation
Version 2.2

Programming Research Group
Oxford University

1. Definitions and declarations.

Let x, x_k be identifiers and T, T_k sets.

LHS $\hat{=}$ RHS Definition of LHS as
 syntactically equivalent to RHS.

$x : T$ Declaration of x as type T .

$x_1 : T_1; x_2 : T_2; \dots; x_n : T_n$
 List of declarations.

$x_1, x_2, \dots, x_n : T$
 $\hat{=}$ $x_1 : T; x_2 : T; \dots; x_n : T$.

$[A, B]$ Introduction of generic sets.

2. Logic.

Let P, Q be predicates and D declarations.

true, false Logical constants.

$\neg P$ Negation: "not P ".

$P \wedge Q$ Conjunction: " P and Q ".

$P \vee Q$ Disjunction: " P or Q ".

$P \Rightarrow Q$ Implication: " P implies Q " or
 "if P then Q ".

$P \Leftrightarrow Q$ Equivalence: " P is logically
 equivalent to Q ".

$\forall x : T \cdot P$
 Universal quantification:
 "for all x of type T, P holds".

$\exists x : T \cdot P$
 Existential quantification: "there
 exists an x of type T such that P ".

$\exists! x : T \cdot P_x$
 Unique existence: "there exists a
 unique x of type T such that P ".
 $\hat{=}$ $(\exists x : T \cdot P_x \wedge$
 $\neg(\exists y : T \mid y \neq x \cdot P_y))$

$\forall x_1 : T_1; x_2 : T_2; \dots; x_n : T_n \cdot P$
 "For all x_1 of type T_1 ,
 x_2 of type T_2, \dots , and
 x_n of type T_n, P holds.

$\exists x_1 : T_1; x_2 : T_2; \dots; x_n : T_n \cdot P$
 Similar to \forall .

$\exists! x_1 : T_1; x_2 : T_2; \dots; x_n : T_n \cdot P$
 Similar to \forall .

$\forall D \mid P \cdot Q \hat{=} (\forall D \cdot P \Rightarrow Q)$.

$\exists D \mid P \cdot Q \hat{=} (\exists D \cdot P \wedge Q)$.

$t_1 = t_2$ Equality between terms.

$t_1 \neq t_2 \hat{=} \neg(t_1 = t_2)$.

3. Sets.

Let S, T and X be sets; t, t_k terms; P a
 predicate and D declarations.

$t \in S$ Set membership: " t is an element
 of S ".

$t \notin S \hat{=} \neg(t \in S)$.

$S \subseteq T$ Set inclusion:
 $\hat{=} (\forall x : S \cdot x \in T)$.

$S \subset T$ Strict set inclusion:
 $\hat{=} S \subseteq T \wedge S \neq T$.

$\{\}$ The empty set.

$\{t_1, t_2, \dots, t_n\}$ The set
 containing t_1, t_2, \dots and t_n .

$\{x : T \mid P\}$
 The set containing exactly those
 x of type T for which P holds.

(t_1, t_2, \dots, t_n) Ordered n -tuple
 of t_1, t_2, \dots and t_n .

$T_1 \times T_2 \times \dots \times T_n$ Cartesian product:
 the set of all n -tuples such that
 the k th component is of type T_k .

$\{x_1 : T_1; x_2 : T_2; \dots; x_n : T_n \mid P\}$
 The set of n -tuples
 (x_1, x_2, \dots, x_n) with each
 x_k of type T_k such that P holds.

$\{D \mid P \cdot t\}$ The set of t 's such that given
 the declarations D, P holds.

$\{D \cdot t\}$	$\hat{=} \{D \mid \text{true} \cdot t\}$.
$\mathbf{P} S$	Powerset: the set of all subsets of S .
$\mathbf{F} S$	Set of finite subsets of S : $\hat{=} \{T : \mathbf{P} S \mid T \text{ is finite}\}$.
$S \cap T$	Set intersection: given $S, T : \mathbf{P} X$, $\hat{=} \{x : X \mid x \in S \wedge x \in T\}$.
$S \cup T$	Set union: given $S, T : \mathbf{P} X$, $\hat{=} \{x : X \mid x \in S \vee x \in T\}$.
$S - T$	Set difference: given $S, T : \mathbf{P} X$, $\hat{=} \{x : X \mid x \in S \wedge x \notin T\}$.
$\cap SS$	Distributed set intersection: given $SS : \mathbf{P} (\mathbf{P} X)$, $\hat{=} \{x : X \mid (\forall S : SS \cdot x \in S)\}$.
$\cup SS$	Distributed set union: given $SS : \mathbf{P} (\mathbf{P} X)$, $\hat{=} \{x : X \mid (\exists S : SS \cdot x \in S)\}$.
$ S $	Size (number of distinct elements) of a finite set.
$\#S$	$\hat{=} S $.

4. Numbers.

\mathbf{N}	The set of natural numbers (non-negative integers).
\mathbf{N}^+	The set of strictly positive natural numbers: $\hat{=} \mathbf{N} - \{0\}$.
\mathbf{Z}	The set of integers (positive, zero and negative).
$m..n$	The set of integers between m and n inclusive: $\hat{=} \{k : \mathbf{Z} \mid m \leq k \wedge k \leq n\}$.
$\min S$	Minimum of a set, $S : \mathbf{F} \mathbf{N}$. $\min S \in S \wedge$ $(\forall x : S \cdot x \geq \min S)$.
$\max S$	Maximum of a set, $S : \mathbf{F} \mathbf{N}$. $\max S \in S \wedge$ $(\forall x : S \cdot x \leq \max S)$.

5. Relations.

A relation is modelled by a set of ordered pairs hence operators defined for sets can be used on relations.

Let X, Y , and Z be sets; $x : X$; $y : Y$; and $R : X \leftrightarrow Y$.

$X \leftrightarrow Y$ The set of relations from X to Y :
 $\hat{=} \mathbf{P} (X \times Y)$.

$x R y$ x is related by R to y :
 $\hat{=} \langle x, y \rangle \in R$.

$x \mapsto y \hat{=} \langle x, y \rangle$

$\{x_1 \mapsto y_1, x_2 \mapsto y_2, \dots, x_n \mapsto y_n\}$

The relation

$\{\langle x_1, y_1 \rangle, \dots, \langle x_n, y_n \rangle\}$

relating x_1 to y_1, \dots , and x_n to y_n .

$\text{dom } R$ The domain of a relation:

$\hat{=} \{x : X \mid (\exists y : Y \cdot x R y)\}$.

$\text{rng } R$ The range of a relation:

$\hat{=} \{y : Y \mid (\exists x : X \cdot x R y)\}$.

$R_1 \# R_2$ Forward relational composition:

given $R_1 : X \leftrightarrow Y$; $R_2 : Y \leftrightarrow Z$,

$\hat{=} \{x : X; z : Z \mid (\exists y : Y \cdot$

$x R_1 y \wedge y R_2 z)\}$.

$R_1 \circ R_2$ Relational composition:

$\hat{=} R_2 \# R_1$.

R^{-1} Inverse of relation R :

$\hat{=} \{y : Y; x : X \mid x R y\}$.

$\text{id } X$ Identity function on the set X :

$\hat{=} \{x : X \cdot x \mapsto x\}$.

R^k The relation R composed with itself k times: given $R : X \leftrightarrow X$,

$R^0 \hat{=} \text{id } X$, $R^{k+1} \hat{=} R^k \circ R$.

R^* Reflexive transitive closure:

$\hat{=} \cup \{n : \mathbf{N} \cdot R^n\}$.

R^+ Non-reflexive transitive closure:

$\hat{=} \cup \{n : \mathbf{N}^+ \cdot R^n\}$.

$R(S)$ Image: given $S : \mathbf{P} X$,

$\hat{=} \{y : Y \mid (\exists x : S \cdot x R y)\}$.

$S \triangleleft R$ Domain restriction to S:
 given $S = P X$,
 $\hat{=} \{x: X; y: Y \mid x \in S \wedge x R y\}$.

$S \triangleleft R$ Domain subtraction:
 given $S = P X$,
 $\hat{=} (X - S) \triangleleft R$.

$R \triangleright T$ Range restriction to T:
 given $T = P Y$,
 $\hat{=} \{x: X; y: Y \mid x R y \wedge y \in T\}$.

$R \triangleright T$ Range subtraction of T:
 given $T = P Y$,
 $\hat{=} R \triangleright (Y - T)$.

$R_1 \oplus R_2$ Overriding: given $R_1, R_2 : X \leftrightarrow Y$,
 $\hat{=} (\text{dom } R_2 \triangleleft R_1) \cup R_2$.

6. Functions.

A function is a relation with the property that for each element in its domain there is a unique element in its range related to it. As functions are relations all the operators defined above for relations also apply to functions.

$X \rightarrow Y$ The set of partial functions from X to Y:
 $\hat{=} \{ f : X \leftrightarrow Y \mid$
 $(\forall x : \text{dom } f \cdot$
 $(\exists ! y : Y \cdot x f y)) \}$.

$X \rightarrow Y$ The set of total functions from X to Y:
 $\hat{=} \{ f : X \rightarrow Y \mid \text{dom } f = X \}$.

$X \twoheadrightarrow Y$ The set of one-to-one partial functions from X to Y:
 $\hat{=} \{ f : X \leftrightarrow Y \mid$
 $(\forall y : \text{rng } f \cdot$
 $(\exists ! x : X \cdot x f y)) \}$.

$X \twoheadrightarrow Y$ The set of one-to-one total functions from X to Y:
 $\hat{=} \{ f : X \twoheadrightarrow Y \mid \text{dom } f = X \}$.

$f \ t$ The function f applied to t.

$(\lambda x : X \mid P \cdot t)$
 Lambda-abstraction:
 the function that given an argument x of type X such that P holds the result is t.
 $\hat{=} \{ x : X \mid P \cdot x \mapsto t \}$.

$(\lambda x_1 : T_1; \dots ; x_n : T_n \mid P \cdot t)$
 $\hat{=} \{ x_1 : T_1; \dots ; x_n : T_n \mid P \cdot$
 $(x_1, \dots, x_n) \mapsto t \}$.

7. Orders.

partial_order X

The set of partial orders on X.
 $\hat{=} \{ R : X \leftrightarrow X \mid \forall x, y, z : X \cdot$
 $x R x \wedge$
 $x R y \wedge y R x \implies x = y \wedge$
 $x R y \wedge y R z \implies x R z$
 $\}$.

total_order X

The set of total orders on X.
 $\hat{=} \{ R : \text{partial_order } X \mid$
 $\forall x, y : X \cdot$
 $x R y \vee y R x$
 $\}$.

monotonic $X <_X$

The set of functions from X to X that are monotonic with respect to the order $<_X$ on X.
 $\hat{=} \{ f : X \rightarrow X \mid$
 $x <_X y \implies f(x) <_X f(y)$
 $\}$.

8. Sequences.

seq X The set of sequences whose elements are drawn from X:
 $\hat{=} \{ A: \mathbb{N}^+ \rightarrow X \mid \text{dom } A = 1..|A| \}$.

$|A|$ The length of sequence A.
 $[\]$ The empty sequence $\{\}$.
 $[a_1, \dots, a_n]$
 $\hat{=} \{ 1 \mapsto a_1, \dots, n \mapsto a_n \}$.
 $[a_1, \dots, a_n] \hat{\sim} [b_1, \dots, b_m]$
 Concatenation:
 $\hat{=} [a_1, \dots, a_n, b_1, \dots, b_m]$,
 $[\] \hat{\sim} A = A \hat{\sim} [\] = A$.

head A $\hat{=} A(1)$.
last A $\hat{=} A(|A|)$.
tail [x] $\hat{\sim} A \hat{=} A$.
front A $\hat{\sim} [x] \hat{=} A$.
rev [a₁, a₂, ..., a_n]
 Reverse:
 $\hat{=} [a_n, \dots, a_2, a_1]$,
 $\text{rev } [\] = [\]$.

$\hat{\sim} / AA$ Distributed concatenation:
 given $AA: \text{seq}(\text{seq}(X))$,
 $\hat{=} AA(1) \hat{\sim} \dots \hat{\sim} AA(|AA|)$,
 $\hat{\sim} / [\] = [\]$.

$\hat{\sim} / AR$ Distributed relational composition:
 given $AR: \text{seq}(X \leftrightarrow X)$,
 $\hat{=} AR(1) \hat{\sim} \dots \hat{\sim} AR(|AR|)$,
 $\hat{\sim} / [\] = \text{id } X$.

disjoint AS Pairwise disjoint:
 given $AS: \text{seq}(P X)$,
 $\hat{=} (\forall i, j: \text{dom } AS \cdot i \neq j \Rightarrow AS(i) \cap AS(j) = \{\})$.

AS partitions S
 $\hat{=} \text{disjoint } AS$
 $\wedge \cup \text{ran } AS = S$.

A in B Contiguous subsequence:

$\hat{=} (\exists C, D: \text{seq } X \cdot C \hat{\sim} A \hat{\sim} D = B)$.

squash f Convert a function, $f: \mathbb{N} \rightarrow X$, into a sequence by squashing its domain.

$\text{squash } \{\} = [\]$,
 and if $f \neq \{\}$ then

$\text{squash } f =$

$[f(i)] \hat{\sim} \text{squash}(\{i\} \triangleleft f)$

where $i = \min(\text{dom } f)$ e.g.

$\text{squash } \{2 \mapsto A, 27 \mapsto C, 4 \mapsto B\}$
 $= [A, B, C]$

S 1 A Restrict the sequence A to those items whose index is in the set S:
 $\hat{=} \text{squash}(S \triangleleft A)$

A 1 T Restrict the range of the sequence A to the set T:
 $\hat{=} \text{squash}(A \triangleright T)$.

9. Bags.

bag X The set of bags whose elements are drawn from X:

$\hat{=} X \rightarrow \mathbb{N}^+$

A bag is represented by a function that maps each element in the bag onto its frequency of occurrence in the bag.

$[\]$ The empty bag $\{\}$.

$[x_1, x_2, \dots, x_n]$ The bag containing x_1, x_2, \dots and x_n with the frequency they occur in the list.

items s The bag of items contained in the sequence s:

$\hat{=} \{ x: \text{rng } s \cdot$

$x \mapsto \{i: \text{dom } s \mid s(i) = x\} \}$

Z Reference Card

Schema Notation

[For details see "Schemas in Z"]

Programming Research Group
Oxford University

Schema definition: a schema groups together some declarations of variables and a predicate relating these variables. There are two ways of writing schemas: vertically, for example

S	_____
	x : N
	y : seq N

	x ≤ y

or horizontally, for the same example

$S \hat{=} [x : N; y : \text{seq } N \mid x \leq |y|]$.

Use in signatures after $\forall, \lambda, \{ \dots \}$, etc.:

$(\forall S \cdot y \neq []) \hat{=} (\forall x : N; y : \text{seq } N \mid x \leq |y| \cdot y \neq [])$.

tuple S The tuple formed of a schema's variables.

pred S The predicate part of a schema: e.g. pred S is $x \leq |y|$.

Inclusion A schema S may be included within the declarations of a schema T, in which case the declarations of S are merged with the other declarations of T (variables declared in both S and T must be the same type) and the predicates of S and T are conjoined. e.g.

T	_____
	S
	z : N

	z < x

is

x, z : N
y : seq N

x ≤ y ∧ z < x

S | P The schema S with P conjoined to its predicate part. e.g.

$(S \mid x > 0)$ is

$[x : N; y : \text{seq } N \mid x \leq |y| \wedge x > 0]$.

S ; D

The schema S with the declarations D merged with the declarations of S. e.g.

$(S ; z : N)$ is

$[x, z : N; y : \text{seq } N \mid x \leq |y|]$

S[new/old] Renaming of components:

the schema S with the component old renamed to new in its

declaration and every free use of that old within the predicate.

e.g. $S[z/x]$ is

$[z : N; y : \text{seq } N \mid z \leq |y|]$

and $S[y/x, x/y]$ is

$[y : N; x : \text{seq } N \mid y \leq |x|]$

Decoration

Decoration with subscript, superscript, prime, etc.;

systematic renaming of the

variables declared in the schema.

e.g. S' is

$[x' : N; y' : \text{seq } N \mid x' \leq |y'|]$

¬S

The schema S with its predicate part negated. e.g. ¬S is

$[x : N; y : \text{seq } N \mid \neg(x \leq |y|)]$

S ^ T

The schema formed from schemas S and T by merging their declarations (see inclusion above) and and'ing their predicates. Given

$T \hat{=} [x : N; z : P \ N \mid x \in z]$,

$S \wedge T$ is

$x : \mathbb{N}$ $y : \text{seq } \mathbb{N}$ $z : \mathbb{P} \mathbb{N}$
$x \leq y \wedge x \in z$

$S \vee T$ The schema formed from schemas S and T by merging their declarations and or'ing their predicates. e.g. $S \vee T$ is

$x : \mathbb{N}$ $y : \text{seq } \mathbb{N}$ $z : \mathbb{P} \mathbb{N}$
$x \leq y \vee x \in z$

$S \Rightarrow T$ The schema formed from schemas S and T by merging their declarations and taking $\text{pred } S \Rightarrow \text{pred } T$ as the predicate. e.g. $S \Rightarrow T$ is similar to $S \wedge T$ and $S \vee T$ except the predicate contains an " \Rightarrow " rather than an " \wedge " or an " \vee ".

$S \Leftrightarrow T$ The schema formed from schemas S and T by merging their declarations and taking $\text{pred } S \Leftrightarrow \text{pred } T$ as the predicate. e.g. $S \Leftrightarrow T$ the same as $S \wedge T$ with " \Leftrightarrow " in place of the " \wedge ".

$S \setminus (v_1, v_2, \dots, v_n)$
 Hiding: the schema S with the variables $v_1, v_2, \dots,$ and v_n hidden: the variables listed are removed from the declarations and are existentially quantified in the predicate. e.g. $S \setminus x$ is $[y:\text{seq } \mathbb{N} \mid (\exists x:\mathbb{N} \cdot x \leq |y|)]$

A schema may be specified instead of a list of variables; in this case the variables declared in that schema are hidden. e.g. $(S \wedge T) \setminus S$ is

$z : \mathbb{P} \mathbb{N}$
$(\exists x:\mathbb{N}; y:\text{seq } \mathbb{N} \cdot x \leq y \wedge x \in z)$

$S \uparrow (v_1, v_2, \dots, v_n)$
 Projection: The schema S with any variables that do not occur in the list v_1, v_2, \dots, v_n hidden: the variables removed from the declarations are existentially quantified in the predicate. e.g. $(S \wedge T) \uparrow (x, y)$ is

$x : \mathbb{N}$ $y : \text{seq } \mathbb{N}$
$(\exists z : \mathbb{P} \mathbb{N} \cdot x \leq y \wedge x \in z)$

The list of variables may be replaced by a schema as for hiding; the variables declared in the schema are used for the projection.

The following conventions are used for variable names in those schemas which represent operations:
 undashed state before the operation,
 dashed state after the operation,
 ending in "?" inputs to the operation, and
 ending in "!" outputs from the operation.

The following schema operations only apply to schemas following the above conventions.

pre S Precondition: all the state after components (dashed) and the outputs (ending in "!=") are hidden. e.g. given

$x?, s, s', y! : \mathbb{N}$
$s' = s - x? \wedge y! = s$

pre S is

$x?, s : \mathbb{N}$
$(\exists s', y! : \mathbb{N} \cdot s' = s - x? \wedge y! = s)$

post S Postcondition: this is similar to precondition except all the state before components (undashed) and inputs (ending in "?") are hidden.

S \bullet T Overriding:
 $\hat{=} (S \wedge \neg \text{pre } T) \vee T$.
 e.g. given S above and

$x?, s, s' : \mathbb{N}$
$s < x? \wedge s' = s$

S \bullet T is

$x?, s, s', y! : \mathbb{N}$
$(s' = s - x? \wedge y! = s \wedge \neg (\exists s' : \mathbb{N} \cdot s < x? \wedge s' = s))$ $\vee (s < x? \wedge s' = s)$

The predicate can be simplified:

$x?, s, s', y! : \mathbb{N}$
$(s' = s - x? \wedge y! = s \wedge s \geq x?)$ $\vee (s < x? \wedge s' = s)$

S \sharp T

Schema composition: if we consider an intermediate state that is both the final state of the operation S and the initial state of the operation T then the composition of S and T is the operation which relates the initial state of S to the final state of T through the intermediate state. To form the composition of S and T we take the state after components of S and the state before components of T that have a basename* in common, rename both to new variables, take the schema "and" (\wedge) of the resulting schemas, and hide the new variables.
 e.g. S \sharp T is

$x?, s, s', y! : \mathbb{N}$
$(\exists s_0 : \mathbb{N} \cdot s_0 = s - x? \wedge y! = s \wedge s_0 < x? \wedge s' = s_0)$

* basename is the name with any decoration ("?", "!", "?", etc.) removed.

S >> T Piping: this schema operation is similar to schema composition; the difference is that rather than identifying the state after components of S with the state before components of T, the output components of S (ending in "!") are identified with the input components of T (ending in "?") that have the same basename.