Invariants, Frames and Postconditions: a Comparison of the VDM and B Notations

Juan Bicarregui and Brian Ritchie

Abstract—
VDM and B are two "model-oriented" formal methods. Each gives a notation for the specification of systems as state machines in terms of a set of states with operations defined as relations on that set. Each has a notion of refinement of data and operations based on the principles of reduction of non-determinism and increase in definedness.

This paper makes a comparison of the two notations through an example of a communications protocol previously formalised in [1]. Two abstractions and two reifications of the original specification are given.

Particular attention is paid to three areas where the notations differ: the use of postconditions that assume the invariant as opposed to postconditions that enforce it; the explicit "framing" of operations as opposed to the "minimal frame" approach; and the use of relational postconditions as opposed to generalised substitutions.

Keywords—**VDM, B Abstract Machines, Formal Methods.**

I. Introduction

In [1], Bruns and Anderson describe a communications protocol in CCS with value-passing. A data model for the values is given which is, in effect, a model of the state of the device. This model is defined in terms of the usual data constructors of model-oriented specification, but without the use of invariants.

The part of the protocol described is a mechanism for manipulating a series of flags that indicate the status of some shared-memory buffers. These flags are used to ensure that there is no "data-tearing" as multiple processors simultaneously read and write to the buffers. For the operations that update these flags, semaphores are used to ensure that each operation has uninterrupted access to the flags. Thus this part of the behaviour can be described as a purely sequential system.

This paper considers some alternative data-models for the specification (and reification) of these status flags. In particular, attention is paid to the use of invariants in the data model and frames of reference in the operation definitions, neither of which are available in the data modelling language of [1]. It is argued that these features can play a key role in describing the system in a "natural" fashion and can thus help to deepen our understanding of the model.

VDM [2] and B [3] are used for the analysis, and particular attention is paid to some areas where the notations differ: the use of postconditions that assume the invariant as opposed to postconditions that enforce it; the explicit "framing" of operations as opposed to the "minimal frame" approach; and the use of relational postconditions as opposed to generalised substitutions. In this small example, there is little scope for the effective use of structuring of specifications that is one of the major features of the B method. Familiarity with the basic concepts and notation of VDM and B is assumed.

The remainder of this first section is an informal description of the application and desired protocol. The second through fifth sections present the development in VDM. Section two presents a formal specification of the system at a level of abstraction similar to the "abstract" description of [1]. Motivated by an analysis of the invariant of that specification, section three describes two further abstractions that can be made. Section four provides an alternative model of the system that makes it possible to write more useful framing information about the operations, the fifth section extends this model to the "improved" protocol of [1]. The sixth section considers the development again using B, presenting those elements of the development that highlight the differences in the notations. The last section is a discussion of some of the points arising from the example and their treatment in the two notations.

A. *The Multiprocessor Shared-Memory Information Exchange*

The Multiprocessor Shared-Memory Information Exchange (MSMIE), is a protocol that addresses intra-subsystem communications with "several features which make it ideally suited to inter-processor communications in distributed, microprocessor based nuclear safety systems" [4]. It has been used in the embedded software of Westinghouse nuclear systems designs.

The protocol uses multiple buffering to ensure that no "data-tearing" occurs, that is, it ensures that data is never overwritten by one process whilst it is being read by another. One important requirement is that neither writing nor reading processes should have to wait for a buffer to become available; another is that "recent" information should be passed, via the buffers, from writers to readers. In the simplification considered in [1] it is assumed that information is being passed from a single writing "slave" processor, to several reading "master" processors.

The information exchange is realised by a system with three buffers. Very roughly, at any time, one buffer is available for writing, one for reading and the third is either in between a write and a read and hence contains the most recently written information, or between a read and a write and so is idle.

The status of each buffer is recorded by a flag which can take one of four values:

s -"assigned to slave." This buffer is reserved for writing, it may actually be being written

Both authors are with the Rutherford Appleton Laboratory, Oxfordshire, UK. E-mail: jcb@inf.rl.ac.uk and br@inf.rl.ac.uk

at the moment or just marked as available for writing.
n -“newest.” This buffer has just been written and contains the latest information. It is not being read at the moment.
m -“assigned to master.” This buffer is being read by one or more processors.
i -“idle.” This buffer is idle, not being read or written and not containing the latest data.

The names of the master processors that are currently reading are also stored in the state.

As mentioned earlier, neither the slave and master processors that access the buffers in parallel nor the actual transference of data are modelled here. This analysis concerns only the operations that modify the buffer status flags. These operations are protected by a system of semaphores which allow each operation uninterrupted access to the state and thus their behaviour is purely sequential.

There are three of these operations:

slave This operation is executed when a write finishes. *slave* sets the status of the buffer that was being written to “newest” thus replacing any other buffer with this status.
acquire This is executed when a read begins. The new reader name (passed as a parameter) is added to the set of readers and status flags are updated as appropriate.
release Executed when a read ends, this removes a reader from the set and updates flags as appropriate.

The details of the behaviour of these operations are quite intricate and their precise description is left to the formal specification in the following section.

It should be noted however that, as it stands, the protocol could have the undesirable property that information flow from slave to master can be held up indefinitely. This possibility is ruled out in the original system [4] via timing constraints whereas [1] suggests an improvement to the protocol (using a fourth buffer) that eliminates the possibility without recourse to timing arguments. This improved protocol is also examined in later sections.

II. A VDM Specification of MSMIE

The state in [1] is defined as

> “a set of three pairs (a, l) where a is the buffer status, drawn from $\{\mathsf{i}, \mathsf{s}, \mathsf{n}, \mathsf{m}\}$, and l is the buffer identification, drawn from $\{1, 2, 3\}$. The buffers are given as a set rather than a tuple to enable pattern matching rules in the description of the protocol.”

The pattern matching rules do indeed give a concise description of the transitions of the system, in particular, the associative and commutative properties of sets are used to good effect in order to avoid much repetitive case analysis. However, the present authors found that considerable effort was required to check that the patterns given were exhaustive and that the effects of overlaps between patterns were sensible. This difficulty is exacerbated by the fact that many of the states in the model are unreachable but no invariant on the state type is given to exclude them.

The specification given here makes the choice of a sequence of three buffers for the state description. In addition, an invariant is used to exclude unwanted values from the state type.

A. The State

Possible values of the status flags are given via an enumerated type; the type of the names of master (reading) processors is deferred.

types

$$Status = \{\mathsf{s}, \mathsf{m}, \mathsf{n}, \mathsf{i}\}$$

$$MName = \mathsf{token}$$

The state is composed of three buffer status flags and a set of the names of the currently reading masters. The invariant captures the fact that only certain states are reachable by the operations. It gives restrictions as to the possible combinations of status flags, namely that there is always exactly one buffer assigned to the writing slave; there is at most one currently being read and at most one with newest data that is not being read; and the set of reader names is empty precisely when there is no buffer being read. The initial state assigns one buffer to the slave and records that the other two buffers are idle.

state Σ of
$b : Status^*$
$ms : MName\text{-set}$
inv $mk\text{-}\Sigma(b, ms) \triangleq \mathsf{len}\, b = 3 \wedge count(\mathsf{s}, b) = 1 \wedge count(\mathsf{m}, b) \in \{0,1\} \wedge count(\mathsf{n}, b) \in \{0,1\} \wedge (count(\mathsf{m}, b) = 0 \Leftrightarrow ms = \{\,\})$
init $mk\text{-}\Sigma(b, ms) \triangleq b = [\mathsf{s}, \mathsf{i}, \mathsf{i}] \wedge ms = \{\,\}$
end

where[1]

$$count : Status \times Status^* \to \mathbf{N}$$

$$count(status, l) \triangleq \mathsf{len}\,(l \rhd status)$$

A.1 A validation condition on the state.

We observe that only four combinations of buffers are allowed by the invariant:

$$\forall mk\text{-}\Sigma(b, ms) \colon \Sigma \cdot \{b(1), b(2), b(3)\}_m \in \{\{\mathsf{s,i,i}\}_m, \{\mathsf{s,i,n}\}_m, \{\mathsf{s,i,m}\}_m, \{\mathsf{s,n,m}\}_m\}$$

where we have used $\{\ldots\}_m$ as a notation for bags (multisets), for example $\{\mathsf{s,i,i}\}_m$ is the bag containing one ‘s’ and two ‘i’s.

Thus the invariant has captured, and brought to the fore, properties that would otherwise have to be deduced

[1] Here, range restriction is used on sequences, viewing them as maps from natural numbers to elements.

by looking in detail at the definitions of the operations. It makes it possible to build quickly our intuition of the workings of the specified machine. We know immediately that there is always one buffer reserved for writing, at most one being read, and at most one with newest data not being read.

B. The Operations

B.1 Slave.

The first operation, *slave*, is executed when a write completes. It reassigns the status of the buffer just written, previously s, to n, thus replacing any other n buffer. It also non-deterministically chooses another available buffer which is to be the new buffer reserved for writing and assigns to it status s.

$$
\begin{array}{l}
\mathit{slave}\ () \\
\textsf{ext wr}\ b\ :\ \mathit{Status}^* \\
\textsf{pre true} \\
\textsf{post}\ \forall i \in \{1,2,3\} \cdot \\
\qquad (\overleftarrow{b}(i) = \textsf{s} \Rightarrow b(i) = \textsf{n}) \wedge \\
\qquad (\overleftarrow{b}(i) = \textsf{m} \Rightarrow b(i) = \textsf{m})
\end{array}
$$

The postcondition may, at first sight, seem to be to liberal: what should happen to any buffer that had status n or i? However, in conjunction with the invariant and the frame, it ensures that no other n buffer remains, that exactly one new s buffer is chosen, and that no new m buffers are added. Thus for example we can write the following validation property for *slave* which can be proved in order to increase confidence in the correctness of the postcondition:

$$\overleftarrow{b}(i) \in \{\textsf{n},\textsf{i}\} \Rightarrow b(i) \in \{\textsf{i},\textsf{s}\}$$

Note that all three implications could have been equivalences without changing the operation.

B.2 Acquire.

The second operation, acquire, is executed when a read is about to start. It adds the new reader's name, passed as a parameter, to the record of active readers and reassigns status flags as necessary.

If there is a buffer currently being read then the new read also begins to read that same buffer and no status change is required. Otherwise the new read starts on the buffer with newest data, status n, and reassigns the status of that buffer to m.

The operation can only be executed in these two situations and this information is recorded in the precondition which requires that there is either a status m or status n buffer. The precondition also records the fact that the operation is only required to function when the new reader is not already in the set of readers.

Note that, in selecting which buffer is to be read, it is not always possible to choose the buffer with newest data. This situation occurs when there are currently buffers with both status m and n, which arises when the data in the n buffer has become available since the start of an ongoing read, that is, when there has been a *slave* since an *acquire* for which there has not yet been a corresponding *release*. In this situation, were the new master to begin reading the n buffer, there would then be two buffers reserved for reading. Consequently, should another *slave* now occur, attempting to preserve this new data would leave no buffer being available for another write to start, thus contradicting one of the fundamental requirements of the protocol: that processors should never have to wait to gain access to buffers. The invariant is designed to prevent this possibility, by insisting that there is always one (and precisely one) buffer with status s.

$$
\begin{array}{l}
\mathit{acq}\ (l\colon \mathit{MName}) \\
\textsf{ext wr}\ b\ \ :\ \mathit{Status}^* \\
\quad \textsf{wr}\ \mathit{ms}\ :\ \mathit{MName}\textsf{-set} \\
\textsf{pre}\ l \notin \mathit{ms} \wedge \\
\quad \exists i \in \{1,2,3\} \cdot b(i) = \textsf{n} \vee b(i) = \textsf{m} \\
\textsf{post}\ \mathit{ms} = \overleftarrow{\mathit{ms}} \cup \{l\} \wedge \\
\quad \forall i \in \{1,2,3\} \cdot \\
\qquad \textsf{if}\ \overleftarrow{b}(i) = \textsf{n} \wedge \overleftarrow{\mathit{ms}} = \{\,\}\ \textsf{then}\ b(i) = \textsf{m} \\
\qquad \textsf{else}\ b(i) = \overleftarrow{b}(i)
\end{array}
$$

It is worth observing that the last line of the postcondition could have been written as

$$\textsf{if}\ \overleftarrow{b}(i) = \textsf{n}\ \textsf{then}\ b(i) \in \{\textsf{n},\textsf{m}\}\ \textsf{else}\ b(i) = \overleftarrow{b}(i)$$

or simply as

$$\overleftarrow{b}(i) \neq \textsf{n} \wedge \mathit{ms} \neq \{\,\} \Rightarrow b(i) = \overleftarrow{b}(i).$$

The apparent non-determinism in the alternatives is illusory as the invariant will ensure that there is no real choice as to what status to assign to any buffer that previously had status n. However, the longer and apparently stronger postcondition is preferred as the shorter versions seem to be more cryptic.

B.3 Release.

The release operation is executed when a reading, master processor finishes its read. The name of the processor is removed from the set of readers and again, status flags reassigned as required.

If this master is not the last one currently reading, then no change is required to the status flags. However, if this is the last master currently reading the m buffer, then this buffer must have its flag reassigned. There are two possibilities. On the one hand, should there be another buffer with status n available at this time, that is if a write has been completed since the current "chain of reads" began on this buffer, then the m buffer no longer contains the most recent data and so should now be set to i. On the other hand, if there has been no write since the chain of reads began, and hence there is no n buffer available, the m buffer contains the most recent data and its status should be reset to n.

$$
\begin{array}{l}
\mathit{rel}\ (l\colon \mathit{MName}) \\
\textsf{ext wr}\ b\ \ :\ \mathit{Status}^* \\
\quad \textsf{wr}\ \mathit{ms}\ :\ \mathit{MName}\textsf{-set} \\
\textsf{pre}\ l \in \mathit{ms}
\end{array}
$$

post $ms = \overleftarrow{ms} - \{l\} \wedge$
$\forall i \in \{1, 2, 3\} \cdot$
if $ms = \{\,\} \wedge \overleftarrow{b}(i) = \mathsf{m}$
then $b(i) \in \{\mathsf{n}, \mathsf{i}\} \wedge count(\mathsf{n}, b) = 1$
else $b(i) = \overleftarrow{b}(i)$

Again there is some choice as to how much of the information that is deducible from the invariant should be made explicit in the postcondition. For example the first conjunct of the 'then' clause $b(i) \in \{\mathsf{n}, \mathsf{i}\}$ could have been omitted as no other possibilities are permitted by the invariant, or alternatively, the whole 'then' clause could be replaced by a more explicit form

if $\exists \mathsf{j} \in \{1, 2, 3\} \cdot \overleftarrow{b}(j) = \mathsf{n}$ then $b(i) = \mathsf{i}$ else $b(i) = \mathsf{n}$

It is debatable which gives the clearer specification.

This specification has given a fairly algorithmic description of which buffers are assigned to what status by each operation. This is a good level of abstraction at which to reason about whole system safety properties such as the freshness of the data transferred from slave to masters which is the focus of [1]. Much of the detail of this specification, however, is undesirable clutter for other purposes and it is interesting to give more "external" views of the system, as is done in the next section.

III. Two More-Abstract Specifications

In this section we give two formal abstractions of the above specification. The new specifications maintain the same external behaviour, however the abstract states are progressively simpler than the one just given. The abstractions arise by ignoring detail in the state model that is unnecessary to capture the external behaviour. Retrieve functions from concrete to abstract states are also given which are many-to-one thus demonstrating "implementation bias" in the concrete specification.

As it is usual to give more concrete specifications successively higher numbers, from now on we will use Σ_2 to refer to the state of the specification given earlier.

A. A First Abstraction: Ignoring the Identity of Buffers

Taking inspiration from the validation condition on the state of the above specification, we can give a more abstract specification where, rather than explicitly giving the status of each individual buffer, the state only records which of the four possible *combinations* of buffer the machine is in.

types

$Status_1 = \{\mathsf{sii}, \mathsf{sin}, \mathsf{sim}, \mathsf{snm}\}$

state Σ_1 of
$bs : Status_1$
$ms : MName\text{-set}$
inv $mk\text{-}\Sigma_1(bs, ms) \;\triangleq\; ms = \{\,\} \Leftrightarrow bs \in \{\mathsf{sii}, \mathsf{sin}\}$
init $mk\text{-}\Sigma_1(bs, ms) \;\triangleq\; bs = \mathsf{sii} \wedge ms = \{\,\}$
end

operations

$slave\ ()$
ext wr $bs : Status_1$
pre true
post $(\overleftarrow{bs} \in \{\mathsf{sii}, \mathsf{sin}\} \Rightarrow bs = \mathsf{sin}) \wedge$
$(\overleftarrow{bs} \in \{\mathsf{sim}, \mathsf{snm}\} \Rightarrow bs = \mathsf{snm})$

As in the earlier specification of *slave*, there is no change to the readers of the m buffer, thus there is no need to access ms.

$acq\ (l{:}\ MName)$
ext wr $bs : Status_1$
wr $ms : MName\text{-set}$
pre $l \notin ms \wedge bs \neq \mathsf{sii}$
post $ms = \overleftarrow{ms} \cup \{l\} \wedge$
if $\overleftarrow{ms} = \{\,\}$ then $bs = \mathsf{sim}$ else $bs = \overleftarrow{bs}$

The definition of *rel* is similar to *acq* and for reasons of brevity has been omitted from this and all subsequent specifications. The complete specifications can be found in [5].

The retrieve function from the first, more concrete, specification to this one is simple to define by cases.

$retr_{2\text{-}1} : \Sigma_2 \rightarrow \Sigma_1$
$retr_{2\text{-}1}(mk\text{-}\Sigma(b_1, b_2, b_3, ms)) \;\triangleq$
cases $(count(\mathsf{n}, [b_1, b_2, b_3]), count(\mathsf{m}, [b_1, b_2, b_3]))$ of
$(0, 0) \rightarrow mk\text{-}\Sigma_1(\mathsf{sii}, ms)$
$(1, 0) \rightarrow mk\text{-}\Sigma_1(\mathsf{sin}, ms)$
$(0, 1) \rightarrow mk\text{-}\Sigma_1(\mathsf{sim}, ms)$
$(1, 1) \rightarrow mk\text{-}\Sigma_1(\mathsf{snm}, ms)$
end

This specification abstracts away from the behaviour of the individual buffers and so it does not help us to reason about the algorithm for updating them. However, it does exhibit a useful congruence on the original state space and makes the property of not returning to the sii states very clear. This observation motivates the following further abstraction.

B. A Further Abstraction

In this specification we abstract away from the buffers entirely: their place being taken by a single boolean flag that records whether a write has ever occurred. Although this specification is consequently extremely simple, it still exhibits the same external behaviour as the original.

state Σ_0 of
$b : \mathbf{B}$
$ms : MName\text{-set}$
inv $mk\text{-}\Sigma_0(b, ms) \;\triangleq\; b = \mathsf{false} \Rightarrow ms = \{\,\}$
init $mk\text{-}\Sigma_0(b, ms) \;\triangleq\; b = \mathsf{false} \wedge ms = \{\,\}$
end

The operations specifications are now very simple:

operations

$slave\ ()$
ext wr b : $\mathbf{B}$
pre true
post $b =$ true

$acq\ (l: MName)$
ext rd b : $\mathbf{B}$
wr ms : $MName$-set
pre $b =$ true $\wedge l \notin ms$
post $ms = \overleftarrow{ms} \cup \{l\}$

The retrieve function is straightforward.

$$retr_{1\text{-}0} : \Sigma_1 \to \Sigma_0$$

$$retr_{1\text{-}0}(mk\text{-}\Sigma_1(bs, ms)) \quad \triangleq \quad mk\text{-}\Sigma_0(bs \neq \text{sii}, ms)$$

IV. An Alternative View of MSMIE

The above specifications are based on the state recording the status of each buffer. Effectively, the state is a map from each buffer to its status. Returning to the original specification, we observe that there is always exactly one buffer with status s and at most one with status m or n. This makes it possible to invert the map and think of the state as mapping each status to a buffer.

This leads to a specification that is equivalent to the first one, but might yield a more efficient basis for an implementation. This change also makes it possible to specify the access constraints more closely.

A. The State

types

$BName = \{1, 2, 3\}$

$MName =$ token

state Σ_3 of
s : $BName$
n : $[BName]$
m : $[BName]$
ms : $MName$-set
inv $mk\text{-}\Sigma_3(s, n, m, ms) \quad \triangleq \quad (\mathsf{m} = nil \ \Leftrightarrow\ ms = \{\,\}) \wedge nil\text{-}or\text{-}different([s, m, n])$
init $mk\text{-}\Sigma_3(s, n, m, ms) \quad \triangleq \quad mk\text{-}\Sigma(1, \text{nil}, \text{nil}, \{\,\})$
end

where $nil\text{-}or\text{-}different([s, n, m])$ is true if and only if each of s, n and m are each mapped to distinct $BNAME$s or nil:

$$nil\text{-}or\text{-}different : [BNAME]^* \to \mathbf{B}$$

$$nil\text{-}or\text{-}different(l) \ \triangleq \ \forall i \in \text{inds}\ l \cdot l(i) = nil \vee l(i) \notin \text{elems}\,(i \lhd l)$$

It is perhaps worth noting that an alternative data model would consist of a single map from *Status* to *BNAME*. However, we have chosen the above because this allows us to narrow the read and write frames of some of the operations.

A.1 The Retrieve Function.

In this case we give the retrieve function implicitly, noting however that it is fully determined (and implementable):

$retr_{3\text{-}2}\ (mk\text{-}\Sigma_3(s, n, m, ms): \Sigma_3)\ \sigma_2: \Sigma_2$
pre true
post let $mk\text{-}\Sigma_2(bs, ms_2) = \Sigma_2$ in
len $bs = 3 \wedge$
$\forall i \in \{1, 2, 3\} \cdot (s = i \ \Rightarrow\ b_i = \mathsf{s}) \wedge (n = i \ \Rightarrow\ b_i = \mathsf{n}) \wedge (m = i \ \Rightarrow\ b_i = \mathsf{m}) \wedge (i \notin \{s, n, m\} \ \Rightarrow\ b_i = \mathsf{i})$
$\wedge\ ms_2 = ms$

B. The Operations

B.1 Slave.

$slave\ ()$
ext rd m : $[BName]$
wr n : $[BName]$
wr s : $BName$
pre true
post $n = \overleftarrow{s}$

The interaction between invariant and externals is interesting. Here, read access to m is required although m is not referred to in the specification. This is because m is linked to s via the invariant and the value of s which is not fully determined by the post-condition: any implementation will need to read m in order to ascertain what value it is valid to assign to s.

Thus rather than think of the externals clauses as giving information about the variables mentioned in the *specification*, we see them as giving details of what access to state variables an *implementation* of that operation can be allowed to make. This distinction separates their semantic role giving information about access to state variables from the syntactic role they play in binding the free variables of the pre- and post-condition.

B.2 Acquire.

$acq\ (l: MName)$
ext wr ms : $MName$-set
wr n, m : $[BName]$
pre $l \notin ms \wedge \neg(n = \text{nil} \wedge m = \text{nil})$
post $ms = \overleftarrow{ms} \cup \{l\} \wedge$
$(\overleftarrow{ms} \neq \{\,\} \ \Rightarrow\ m = \overleftarrow{m} \wedge n = \overleftarrow{n}) \wedge$
$(\overleftarrow{ms} = \{\,\} \ \Rightarrow\ m = \overleftarrow{n} \wedge n = \text{nil})$

Interestingly, the last conjunct of this postcondition could be considered to be redundant. When $\overleftarrow{ms} = \{\,\}$ and thus $\overleftarrow{m}$ is nil, then $ms = \{l\}$ and so m must be assigned a non nil value. Now, as read access to s is prohibited, the only buffer that we can be sure is not already in use is that previously assigned to n. So any implementation that respects the frames of reference must assign this buffer to m. Then the only remaining possible value for n is nil. However, to hide so much information in the externals clause seems to be counter-productive.

V. The Improved MSMIE

As mentioned earlier, Bruns and Anderson observe that, as it stands, the three buffer MSMIE can exhibit an undesirable behaviour. That is, it is possible for a series of overlapping reads, each beginning before the last ends, to lock-out indefinitely the latest data. They suggest an improved protocol that uses a fourth buffer to eliminate this possibility.

Surprisingly, although this new protocol exhibits the same external behaviour as the earlier one, there is no formal refinement relationship between them. To understand why this is, we recall that the part of the system modelled only concerns itself with the assignment of processors to buffers and so does not model the actual transfer of information from slave to masters. Thus, the values assigned to the status flags have no externally visible effect and all the machinations of the state can be seen as purely an implementation bias in the model.

However, the four-buffer version is a refinement of the most abstract specification given earlier, which gives another important reason for considering those abstractions. In particular, validations of the abstract model will carry over to both the three and four buffer versions.

Of course, in this case, it is the internal properties of the model itself that are of interest, as it is these properties that influence the "freshness" of the data read by the masters. In this respect, the four buffer protocol is indeed better behaved as it would lead to a system where the delay in information transfer is at worst equal to that of the three buffer version.

In the four buffer version of MSMIE, there is also an extra status possible for buffers. o is used to denote a buffer that is still being read but no longer contains most up-to-date information. Thus:

- s as before, is a buffer that is reserved for writing
- n as before, is a buffer that contains the latest data but is not being read (waiting for read)
- m is a buffer being read, (and the newest such)
- o is a buffer being read (but there is also a newer one being read)
- ms is the set of masters reading m
- os is the set of masters reading o.

New masters are always assigned to the n or the m buffer. m buffers are "demoted" to o status in a way that ensures that the o buffer will periodically become idle. In this way the protocol avoids the "refresh" problems of the three-buffer version. Again detailed descriptions of the mechanisms used to achieve this is given accompanying the formal text.

It might help to think of the status transitions $i \rightarrow s \rightarrow n$ as the write phase of a buffer and the transitions $n \rightarrow m \rightarrow o \rightarrow i$ as the read phase. We will see that this variant of MSMIE always has two buffers in write phase and two buffers in read phase.

A. The State

types

$$BName = \{1, 2, 3, 4\}$$

state Σ_4 of
- $s : BName$
- $n : [BName]$
- $m : [BName]$
- $o : [BName]$
- $ms : MName\text{-set}$
- $os : MName\text{-set}$

inv $mk\text{-}\Sigma_4(s, n, m, o, ms, os) \ \triangleq$
$$(m = nil \ \Leftrightarrow \ ms = \{\,\}) \wedge$$
$$(o = nil \ \Leftrightarrow \ os = \{\,\}) \wedge$$
$$(ms \cap os = \{\,\}) \wedge$$
$$(nil\text{-}or\text{-}different([s, n, m, o])) \wedge$$
$$(m = nil \wedge n = nil \ \Rightarrow \ o = nil)$$

init $\sigma_4 \ \triangleq \ \sigma_4 = mk\text{-}\Sigma_4(1, \mathsf{nil}, \mathsf{nil}, \mathsf{nil}, \{\,\}, \{\,\})$

end

The last conjunct in the invariant, which rules out the states corresponding to $\{\mathsf{s,o,i,i}\}_m$, is the result of the way that readers of m are released which, as in the earlier specifications, ensures that there is always an m or an n buffer remaining.

A.1 A Validation Property for the State.

The invariant only allows the following states corresponding to the following 7 combinations of buffer status:

$$\{\mathsf{s,i,i,i}\}_m, \{\mathsf{s,i,i,n}\}_m, \{\mathsf{s,i,i,m}\}_m,$$
$$\{\mathsf{s,i,m,n}\}_m, \{\mathsf{s,i,m,o}\}_m,$$
$$\{\mathsf{s,i,n,o}\}_m, \{\mathsf{s,m,n,o}\}_m$$

A.2 Retrieve Function.

As stated earlier this version is a data refinement of the most abstract model. The retrieve function is straightforward:

$$retr_{4\text{-}0} : \Sigma_4 \rightarrow \Sigma_0$$

$$retr_{4\text{-}0}(mk\text{-}\Sigma_4(s, n, m, o, ms, os)) \ \triangleq$$
$$mk\text{-}\Sigma_0(n = \mathsf{nil} \wedge m = \mathsf{nil} \wedge o = \mathsf{nil}, ms \cup os)$$

B. The Operations

B.1 Slave.

$slave\ ()$
ext rd $m, o \ : \ [BName]$
wr $n \ : \ [BName]$
wr $s \ : \ BName$
pre true
post $n = \overleftarrow{s}$

As before the implementation will require access to m and o in order to be able to set a valid s. That is:
$s \in BName - \{n, m, o\}$
This access requirement is recorded in the externals even though the predicates do not mention m and o.

B.2 Acquire.

The descriptions of *acquire* and *release*, given via case analysis, are rather unwieldy. As different variables change in the different cases, the operations have to have wide write access and hence require clauses saying which variables do not change in that case. Thus we introduce a notational shorthand used in postconditions to say that certain state components are unchanged[2]:

$Id : A^* \to \mathbf{B}$

$Id(l) \ \ \triangleq \ \ \forall i \in \mathsf{inds}\, l \cdot l(i) = \overleftarrow{l(i)}$

$acq\ (l\colon MName)$
ext wr $ms, os \ : \ MName$-set
wr $n, m, o \ : \ [BName]$
pre $l \notin ms \cup os \wedge \neg (n = \mathsf{nil} \wedge m = \mathsf{nil})$
post $(ms \cup os = \overleftarrow{ms} \cup \overleftarrow{os} \cup \{l\}) \wedge$
$(\overleftarrow{m} = \mathsf{nil} \ \Rightarrow \ m = \overleftarrow{n} \wedge n = \mathsf{nil} \wedge Id([o, os])) \wedge$
$(\overleftarrow{m} \neq \mathsf{nil} \wedge (\overleftarrow{o} \neq \mathsf{nil} \vee n = \mathsf{nil})$
$\Rightarrow \ Id([m, n, o, os])) \wedge$
$(\overleftarrow{m} \neq \mathsf{nil} \wedge \overleftarrow{o} = \mathsf{nil} \wedge n \neq \mathsf{nil}$

$\Rightarrow \ o = \overleftarrow{m} \wedge m = \overleftarrow{n} \wedge n = \mathsf{nil} \wedge os = \overleftarrow{ms} \wedge ms = \{l\})$

Acquire behaves in a manner similar to before: a reader is assigned to the n or the m buffer as appropriate. The only extra consideration is in the case where there is an n buffer waiting, and an m buffer being read, but no o buffer. In this case, where previously the new reader would have been assigned to the m buffer, it is now possible to begin the read on the n buffer, hence the improvement to the freshness of the data exchanged. The buffer that was already being read is marked as o, and correspondingly the record of processors reading that buffer, ms, is moved to os; and the new read begins on the buffer that was n, thus making it into a new m and the new reader is recorded in ms. No more masters will be assigned to the o buffer until it has been through the write cycle again.

We have seen five specifications which exhibit the same external behaviour. All except for the most abstract incorporate some degree of implementation bias. However, it is this very bias that is the subject under investigation. In [1] validation conditions describing some desirable global properties of the protocol are expressed in the modal μ-calculus. For the purposes of comparison of the two notations considered in this paper it is sufficient to note that neither provides a formalism to express such conditions.

[2] This should be seen as a syntactic "macro", rather than an auxiliary function.

VI. The Specification Using B

A similar series of specifications and refinements was constructed in B. In preference to presenting this material in full, we present only those parts that highlight the notational and stylistic differences between VDM and B which arose in this example.

This development was carried out using the current alpha-release of the B Toolkit [6]. Although this has meant that the specifications have been required to conform exactly to the language supported by the machine[3], it has given us the advantages of automatic consistency checking that the toolkit provides. In this paper, we present the B machines as they were entered in the toolkit, though in some places syntactic sugaring may have made them more readily digestible.

A. The Two 'More-Abstract' Specifications

Figures 1 and 2 gives the B text for the two abstract specifications.

The first specification gives the most abstract specification of MSMIE as a machine $b0$. This corresponds to the VDM specification with state Σ_0. The machine is parameterised by a set $MNAME$ of master names which is assumed to be non-empty. The state consists of two variables $b0$ and ms. Unlike VDM, the typing information for the state variables is given in the invariant. Here, the first two clauses of the invariant give "static" (decidable) typing, and the third clause gives subtyping information. The initialisation and operations are given as generalised substitutions. At this level, the operations are very similar to the VDM ones presented earlier.

The major syntactic differences from VDM are that the types of the arguments are given explicitly as predicates. Also, the read and write frames of the machine operations are implicit. The read frames are the full state of the machine and the variables written are determined by the generalised substitution, for example, those that appear on the left side of a simple substitution. Thus, the operation *slave* writes $b0$, and *acquire* writes ms. Framing is addressed further in the closing discussion.

The first reification $b1$ of $b0$ is presented in Figure 2 as a B refinement. It corresponds to the VDM specification with state Σ_1. Note that the B method makes a notational distinction between refinements and other specifications.

The concrete state also has two variables. By repeating the ms variable name, we are saying that this variable is the same as the one in the abstract specification. Technically, the new state variables include those of the abstract state

[3] In the VDM specifications, we have tried to follow the draft BSI standard as far as possible, but have allowed at least one notational extension in the Id function in the previous section.

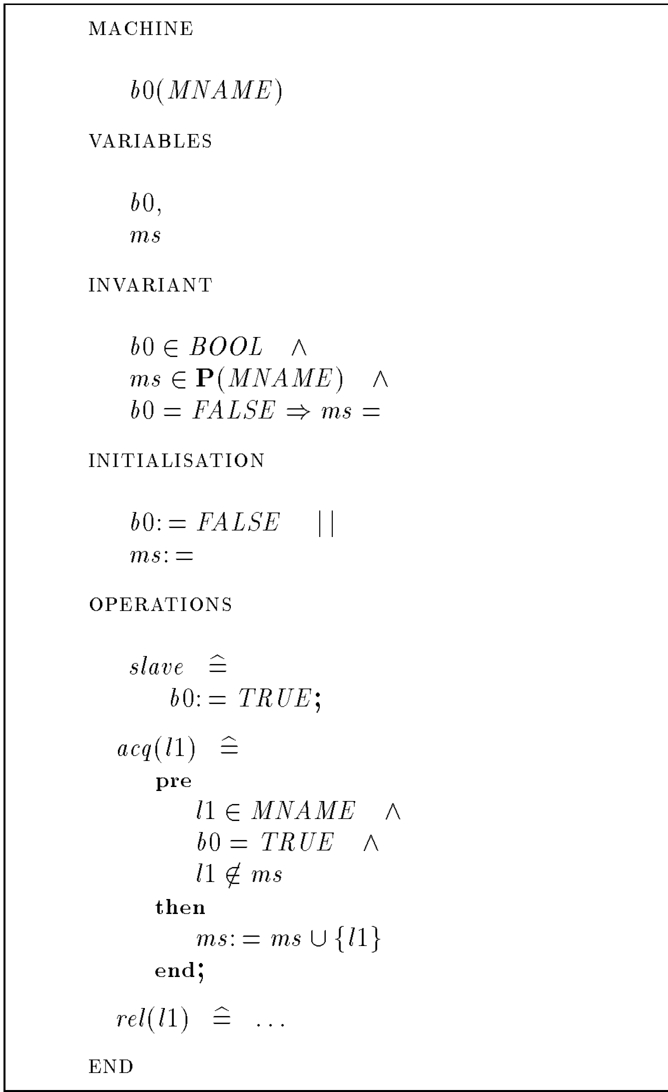

Fig. 1. The Abstract B machine $b0$

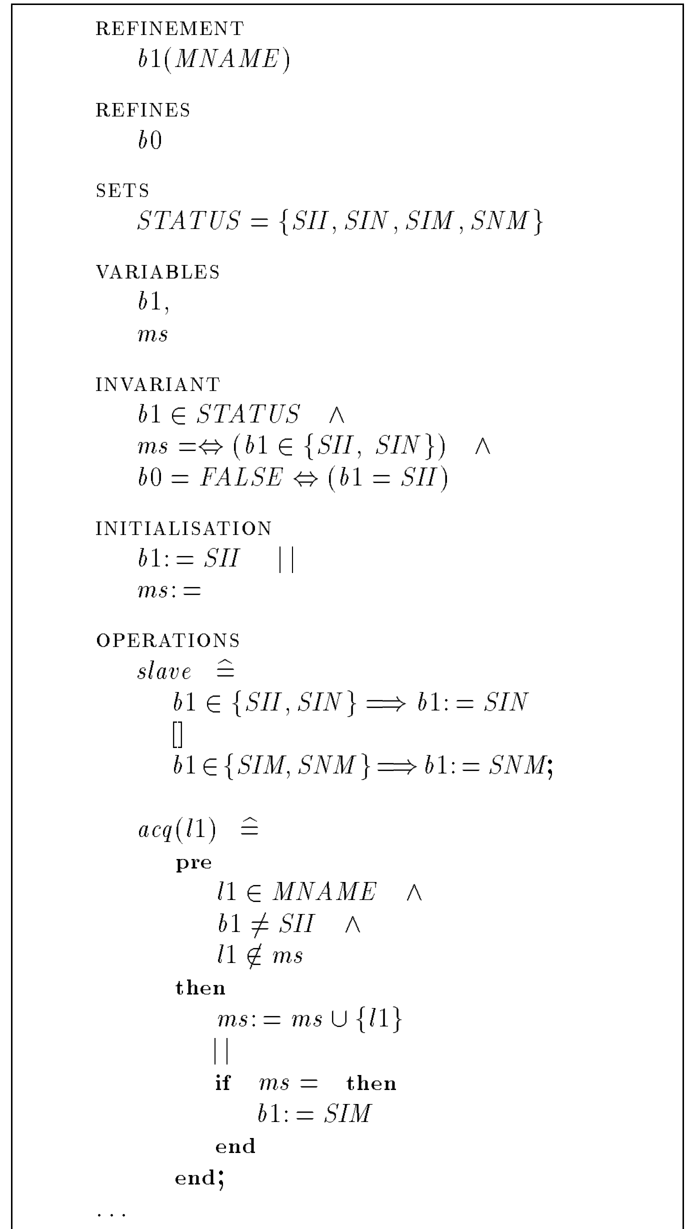

Fig. 2. The B refinement $b1$

as well any added here, however, the variables from the abstract state are subject to full hiding and cannot appear in the definition of operations. The relationship between abstract and concrete variables is given as part of the invariant – a coupling relation, as for example in [7]. In this example, the coupling relation appears as the last conjunct of the invariant. As in $b0$, the operation definitions are similar to those of the corresponding VDM specification.

B. Three Buffer Status

Now we proceed with a refinement $b2$ of $b1$ (Figure 3). This corresponds to the first VDM specification with state Σ_2. The data model (including the invariant) is similar to that of the VDM specification, and the coupling relation is similar to the retrieve function $retr_{2\text{-}1}$. Some informal comments have been added to the formal text.

Non-deterministic choice in a substitution is given by the notation "$@(v).S$", where v is a variable and S a generalised substitution. Here it is used in conjunction with guarded substitutions, in the form "$@(v).(P(v) \implies S)$" which can be read operationally as, "if there is any v such that $P(v)$, then apply the substitution S for one such v".

In this specification, we see a significant difference from the VDM in the presentation of the *slave* operation. Here, we have given a definition of *slave* which provides more explicit algorithmic information than its VDM counterpart. In particular, it breaks down the operation into seperate substitutions executed sequentially and could take one buffer through two changes of status. This style was found to be convenient here as the one "choice" that might be available, namely whether to choose the new slave to be the old newest or an old idle buffer, affects the outcome of two buffers. There is also a syntactic restriction prohibiting parallel substitutions to the same variable. The resulting operation definition has a far more programatic feel.

A similar change in style is reflected in the definitions of *acquire* and *release*, which for brevity are not included

REFINEMENT
$b2(MNAME)$

REFINES
$b1$

SETS
$STATUSII = \{S2,\ I2,\ N2,\ M2\}$;

VARIABLES
$b2$,
ms

INVARIANT
/* typing */
$b2 \in \mathsf{seq}(STATUSII) \quad \wedge$
$ms \in \mathbf{P}(MNAME) \quad \wedge$
/* subtyping */
$\mathsf{size}(b2) = 3 \quad \wedge$
$\mathsf{card}(b2 \rhd \{S2\}) = 1 \quad \wedge$
$\mathsf{card}(b2 \rhd \{M2\}) \in \{0,\ 1\} \quad \wedge$
$\mathsf{card}(b2 \rhd \{N2\}) \in \{0,\ 1\} \quad \wedge$
$\mathsf{card}(b2 \rhd \{M2\}) = 0 \Leftrightarrow (ms =) \quad \wedge$
/* coupling */
$((\mathsf{card}(b2 \rhd \{N2\}) = 0 \ \wedge\ \mathsf{card}(b2 \rhd \{M2\}) = 0)$
$\Leftrightarrow (b1 = SII)) \quad \wedge$
$((\mathsf{card}(b2 \rhd \{N2\}) = 1 \ \wedge\ \mathsf{card}(b2 \rhd \{M2\}) = 0)$
$\Leftrightarrow (b1 = SIN)) \quad \wedge$
$((\mathsf{card}(b2 \rhd \{N2\}) = 0 \ \wedge\ \mathsf{card}(b2 \rhd \{M2\}) = 1)$
$\Leftrightarrow (b1 = SIM)) \quad \wedge$
$((\mathsf{card}(b2 \rhd \{N2\}) = 1 \ \wedge\ \mathsf{card}(b2 \rhd \{M2\}) = 1)$
$\Leftrightarrow (b1 = SNM))$

INITIALISATION
$b2 := [S2,\ I2,\ I2] \quad ||$
$ms :=$

OPERATIONS
$slave \ \widehat{=}$
(/* If there's a buffer that's set to N,
find it and set it to I */
$\exists\, z1.(z1 \in \{1,2,3\} \wedge b2(z1) = N2) \Longrightarrow$
$@(z1).(z1 \in \{1,2,3\} \wedge b2(z1) = N2 \Longrightarrow$
$b2 := (b2 \lhd\!\!\!+ \{z1 \mapsto I2\}))$
[]
/* but if you can't find one that is N,
then don't worry */
$\forall\, z2.(z2 \in \{1,2,3\} \Rightarrow b2(z2) \neq N2) \Longrightarrow skip$
);
/* Then, find a buffer that was S,
and set it to N (It's unique) */
$@(z3).(z3 \in \{1,2,3\} \ \wedge\ b2(z3) = S2 \Longrightarrow$
$b2 \ := \ (b2 \lhd\!\!\!+ \{z3 \mapsto N2\}))$
;
/* then find an I and set it to S
(there are one or two of these) */
$@(z4).(z4 \in \{1,2,3\} \ \wedge\ b2(z4) = I2 \Longrightarrow$
$b2 \ := \ (b2 \lhd\!\!\!+ \{z4 \mapsto S2\}))$
;

Fig. 3. The original three buffer machine specified in B.

here.

C. The "Improved MSMIE" Version

The B machine *snmo* (Figures 4 and 5) corresponds to the VDM specification of the 4-buffer MSMIE. (The 3-buffer "inverted map" version has been omitted from this paper, because the same issues arise with this machine.)

At this level, the two notations are once again fairly similar. One difference from the VDM version is that because the generalised substitutions explicitly indicate which variables are to be substituted there is no need for the *Id* clauses that appear in the VDM version.

VII. Discussion

As stated in the introduction, this example highlights three areas where the notations encourage different approaches. This closing section gives a brief discussion of some of the points that arose from our study of the MSMIE protocol.

A. Invariants.

In both notations, the invariant is useful for quickly conveying an understanding of the reachable values of the state. However the use of invariants in operation definitions differs. In B, postconditions (in the form of generalised substitutions) have to be written so as to ensure the maintenance of the invariant. In VDM the state invariant is effectively part of the state typing information, and as such is assumed to be maintained in addition to the postcondition.

VDM's implicit maintenance of the invariant led to the choice discussed earlier of how much of the information in the invariant is repeated in a postcondition. There was often some tension between the most concise form that relied on properties of the invariant for its correctness, and a longer, but more explicit form, that included some redundant information. This choice can be seen as an opportunity to prove the stronger forms from the weaker. Which form is chosen may make a significant difference to the complexity of the proofs: the form that most clearly conveys the information may not be the form that will be most usable in proofs. Indeed, the stronger form is more likely to be helpful when the specification is being proved to be a reification of another, and the weaker form when it is itself being reified.

In the B notation, on the other hand, one writes operations so as to imply the preservation of the invariant. This can encourage a tendency to describe *how* the invariant is maintained, which may lead to less abstract specifications.

B. Operation Definitions.

The greater programmatic feel of the B notation is reinforced by the use of generalised substitutions, as opposed to VDM's relational post-conditions. Although the two forms have the same expressive power, in some cases (as for example in *slave* in the $b2$ machine) we found it convenient to give greater algorithmic detail in the B version. This

```
REFINEMENT
    snmo(MNAME)

REFINES
    b0

VARIABLES
    sb, nb, mb, ob,
    ms4, os4

INVARIANT
    /* typing */
    sb ∈ 1 .. 4   ∧
    nb ∈ seq(1 .. 4)   ∧
    mb ∈ seq(1 .. 4)   ∧
    ob ∈ seq(1 .. 4)   ∧
    ms4 ∈ P(MNAME)   ∧
    os4 ∈ P(MNAME)   ∧
    /* subtyping */
    size(nb) ∈ {0, 1}   ∧
    size(mb) ∈ {0, 1}   ∧
    size(ob) ∈ {0, 1}   ∧
    mb = [] ⇔ (ms4 =)   ∧
    ob = [] ⇔ (os4 =)   ∧
    [sb] ≠ nb   ∧
    [sb] ≠ mb   ∧
    [sb] ≠ ob   ∧
    nb = mb ⇔ (nb = [] ∧ mb = [])  ∧
    nb = ob ⇔ (nb = [] ∧ ob = [])  ∧
    mb = ob ⇔ (mb = [] ∧ ob = [])  ∧
    mb = [] ∧ nb = [] ⇒ ob = []   ∧
    /* coupling */
    ms = ms4 ∪ os4   ∧
    b0 = FALSE ⇔
        (nb = [] ∧ mb = [] ∧ ob = [])
```

Fig. 4. The four buffer MSMIE (state) as a B refinement

```
INITIALISATION
    sb := 1   ||
    nb := []   ||
    mb := []   ||
    ob := []   ||
    ms4 :=   ||
    os4 :=

OPERATIONS

    slave  ≙
        nb := [sb]   ||
        sb :∈ {1, 2, 3, 4} − {sb}
              −ran(mb) − ran(ob);

    acq(l1)  ≙
        pre
            l1 ∈ MNAME   ∧
            l1 ∉ ms4   ∧
            conc([nb, mb]) ≠ []
        then
            select   mb = []   then
                mb := nb   ||
                nb := []   ||
                ms4 := {l1}
            when   mb ≠ [] ∧ nb ≠ [] ∧
                      ob = []  then
                ob := mb   ||
                mb := nb   ||
                nb := []   ||
                os4 := ms4   ||
                ms4 := {l1}
            else
                ms4 := ms4 ∪ {l1}
            end
        end;
```

Fig. 5. The four buffer MSMIE (operations) as a B refinement

would appear to imply that the B notation is more useful for the development of algorithms. Indeed, the process of operation decomposition has been given greater attention in the B methodology than for VDM. By contrast, perhaps VDM's relational postconditions give a greater facility for non-algorithmic specifications of complex operations.

C. Framing.

As stated earlier, the read and write frames are given explicitly in a VDM operation, whereas in B the variable access and modification is implicit in the form of the generalised substitution.

In VDM operations, the semantic role of the read frame is often underplayed. Typically, it is interpreted as merely providing syntactic scoping for variables appearing in the precondition or postcondition. Alternatively, it could be interpreted as a constraint on implementations, restricting which state components can be read. Thus rather than think of the externals clauses as giving information about the variables mentioned in the *specification*, we see them as giving details of what access to state variables an *implementation* of that operation can be allowed to make. (See [8] and [9] for further discussion of this point.)

In B, similar restrictions can be given through the hiding principles inherent in the different forms of machine structuring. For instance in this example, where we were able to narrow the read frames in the later VDM specifications, in the B counterparts there is a potential to structure the overall machine as a B "implementation" in terms of simpler machines, one for each status flag.

In the above we have emphasised three areas where our experiments have suggested that the notations of VDM and B encourage different specification styles. Each style may have its own advantages at different stages of the development process. In this example we found that the process of developing implementation code was better addressed in B's abstract machine notation. However, we also found VDM's relational postconditions more convenient for expressing wholly implicit specifications of operations, par-

ticularly when the data model involved complex interdependencies.

References

[1] G. Bruns and S. Anderson, *The Formalization and Analysis of a Communications Protocol.* Formal Aspects of Computing (1994) 6: 92-112. (Previously released as *The Formalization of a Communications Protocol.* University of Edinburgh Tech. Rep. LFCS 91-137 (April 1991) and as LFCS/Adelard SCCS Tech. Rep. April 6, 1992.

[2] C.B. Jones, *Systematic Software Development Using VDM*, second edition. Prentice Hall, 1990.

[3] J.R. Abrial, *Assigning Programs to Meanings.* Prentice-Hall, to appear.

[4] L.L. Santoline *et al. Multiprocessor Shared-Memory Information Exchange.* IEEE Trans. on Nuclear Science. 36(**1**) 1989, pp.626-633.

[5] J.C. Bicarregui, and B. Ritchie. *Invariants, Frames and Postconditions: A Comparison of Two Formal Specification Notations*, Rutherford Appleton Laboratory Technical Report, RAL-93-100, 1993.

[6] J.R. Abrial, *Introducing B-Technologies.* Unpublished, May 1992.

[7] C. Morgan, *Programming from Specifications.* Prentice Hall, 1990.

[8] J.C. Bicarregui, *Algorithm refinement with read and write frames*, In J.C.P. Woodcock and P.G. Larsen (Eds), FME'93: Industrial Strength Formal Methods, Lecture Notes in Computer Science 670, Springer-Verlag 1993.

[9] J.C. Bicarregui, *Operation Semantics with read and write frames*, In Proc. of BCS FACS Sixth Refinement workshop, City University, London, 5-7 January 1994. To appear in Workshops in Computing, Springer-Verlag.