

Enhancing the Tractability of Rely/Guarantee Specifications in the Development of Interfering Operations *

P. Collette and C. B. Jones

July 5, 2003

Abstract

Various forms of assumption/commitment specifications have been used to specify and reason about the interference that comes from concurrent execution; in particular, consistent and complete proof rules relating to shared state operation specifications –with rely and guarantee conditions– have been published elsewhere. This paper discusses some issues about the formulation of such specifications and the way to record design decisions so as to make the use of rely/guarantee conditions more tractable.

*Please cite the original publication details of this paper *Enhancing the Tractability of Rely/Guarantee Specifications in the Development of Interfering Operations*, Pierre Collette, Cliff B. Jones, in *Proof, Language and Interaction*, (eds.) Gordon Plotkin, Colin Stirling, Mads Tofte, Chapter 10, pp277–307, MIT Press, 2000

Contents

1	Introduction	3
2	Introduction to the Case Study	4
2.1	Sequential Operations	5
2.2	Interfering Operations	6
2.3	Data Reification	7
2.4	Operation Refinement	8
3	Visible Steps	8
4	Invariants for Interfering Operations	9
4.1	Data Invariants	9
4.2	Evolution Invariant	10
5	Writing Specifications	11
5.1	Usefulness of the Invariants	12
5.2	Enriched Mode Restrictions	13
5.3	Predominance of the Post Condition	14
5.4	Interference and Post Conditions	14
5.5	Reasoning about Specifications	15
5.6	Transitivity	16
6	Towards Code	16
6.1	Control over interference	16
6.2	Introduction of Code	18
6.3	Verification of the Invariants	19
7	Conclusion	20
	APPENDIX	22
A	Technical Summary	23

1 Introduction

Formal methods based on model-oriented specifications like VDM or B are applicable to the development of sequential operations. In such approaches, state components can be common to several operations but only one operation is executed at a time. A sequential operation can then be interpreted as a binary relation on the state space and specified with pre and post conditions; examples are given below but readers are assumed to be familiar with pre/post specifications in the style of VDM. In [Jon81], rely and guarantee conditions are proposed as an extension to cope with the specification and development of *concurrent* operations, a situation that occurs when operations sharing state components have overlapping executions. The necessary background about rely/guarantee specifications is recalled in this paper – detailed expositions (including sound and complete proof systems) can be found in [Stø91]. The new insights in this paper come from an emphasis on methodological issues. Theoretical aspects of rely/guarantee specifications are intentionally omitted in favour of suggestions that improve their practicability in the development of concurrent operations.

This paper is concerned with imperative programs whose meaning can be discussed with respect to a set of states – say $s_i \in \Sigma$. The additional complexity of concurrent versus sequential operations is due to the presence of *interference*: operations access state components that can be modified by the execution of other operations during their own execution. This difference with sequential operations can be emphasized by looking at computations. A computation of a sequential operation can be viewed as a single transition

$$s_0 \xrightarrow{\pi} s_1$$

from a starting state s_0 to a final state s_1 . Of course, there might be many intermediate states between s_0 and s_1 but only the initial and final states can be accessed by other operations. The (superfluous) label π indicates that this transition is performed by the operation. In the presence of interference, a computation not only includes steps of the operation, but also steps from its environment (other operations). If the latter are labelled with ϵ , a computation can be represented by a sequence of transitions

$$s_0 \xrightarrow{l_0} s_1 \xrightarrow{l_1} s_2 \dots$$

where each label l_i is either π or ϵ ; computations that terminate have a finite number of π -labelled steps.¹

Usually, termination in an acceptable state can only be ensured under assumptions about the initial state. In specifications of sequential operations, such *assumptions* are recorded in a pre condition, whereas the *commitments* of the operation (definition of acceptability) are recorded in a post condition. It is understood that the commitments are to be fulfilled only when the assumptions are satisfied. Termination of an interfering operation in an acceptable state also requires assumptions about the initial state but this is not sufficient: assumptions about the ϵ -labelled steps are essential. Indeed, nothing reasonable can be expected from an operation whose environment modifies the state in an arbitrary

¹Whether these are finer or coarser grained steps is a key issue that is discussed further below; meanwhile the steps are referred to as the visible steps of an operation.

way.

The use of assumption/commitment specifications in the development of concurrent systems is not restricted to the formalism discussed in this paper: other examples are [AL93, BK85, Col94, JT95, KR92, MC81, PJ91, Sta86, ZdBdR84]. Some of the methodological issues raised in this paper hopefully spread across examples and formalisms but the case study is only representative of one specific class of shared-state operations. In general, operations have both an *input/output* behaviour and a *reactive* behaviour. The former determines the result of an operation in terms of its inputs whereas the latter describes the way it interacts with other operations. This paper focuses on operations for which the input/output behaviour is more important than the reactive behaviour; what really matters for these operations is their final result. This can be contrasted with the development of components like schedulers or senders/receivers whose reactive behaviour is prominent and which can probably be better specified in formalisms like temporal logic. Any classification is highly debatable but a possible characterisation is that, in the absence of interference, the same operations should be specifiable with pre and post conditions. The case study illustrates this. Section 2 gives specifications for both the sequential and the interfering versions of the same operations; the latter are inevitably more sophisticated than the former but in both cases, what really matters is the input/output behaviour.

Section 2 illustrates the use of rely and guarantee conditions with top-level specifications from the case study; a brief sketch of the development is also given. No novelty appears in Section 2; in particular, the specifications are subject to substantial improvement—in accordance with the suggestions made—in the remainder of the paper. Visible steps are defined in Section 3. Next, the use of data invariants and other useful invariant properties is advocated in Section 4. Finally, recommendations on writing specifications and on the refinement of operations towards code are proposed in Sections 5 and 6 respectively.

2 Introduction to the Case Study

The problem of recording equivalence classes over a (finite) set T of elements occurs in a variety of contexts from controlling equivalent part numbers in manufacturing applications to tracking equivalence classes in cryptography. The two basic operations are $\text{TEST}(a, b)$ that tests if a and b are elements of the same class and $\text{EQUATE}(a, b)$ that merges the equivalence classes of a and b into a single class. An efficient implementation—both in terms of space and time—can be based on a representation which records equivalence classes as trees (one tree per class); there are various proposals for keeping the path lengths short within trees, here a new operation $\text{CLEANUP}(a)$ that shortens the path from a to its root in the tree is added.

The representation by trees and the introduction of the CLEANUP operation are clearly insights in the development; but this case study has been carried out without other insights than these. This deliberate ignorance of previous solutions and the hunger for interfering operations actually led to tackling a more general problem. The complexity of this problem indeed increases with the degree of concurrency. In the absence of interference, algorithms for the operations can be designed quite easily; the task becomes more complex when CLEANUP interferes with EQUATE or TEST , by modifying the inner structure

of trees. This development goes further in that it also permits the concurrent execution of TEST and EQUATE, and also the concurrent execution of *several instances* thereof. Although the usefulness of such a high degree of concurrency is debatable, it provides a sufficiently complex problem to raise issues about the practicability of a development method.

In model-oriented developments, operations are first specified over an abstract state space, which is then progressively reified into more concrete ones until all operations are specified in terms of implementable data structures. This section gives specifications of TEST and EQUATE on abstract states. The unique state component is p : *T-partition* where

$$T\text{-partition} = \{p \in (T\text{-set})\text{-set} \mid is\text{-partition}(p)\}.$$

The type invariant *is-partition*(p) indicates that the union of the sets in p is T and that sets are disjoint;² each set in p records an equivalence class on T . VDM specifications for the sequential version of TEST and EQUATE are given first; rely and guarantee conditions are then introduced to cope with the concurrent execution of several instances of TEST with several instances of EQUATE.

2.1 Sequential Operations

Sequential operations can be specified by pre and post conditions; hooked variables in post conditions refer to the initial state. Those elements in the same class as a in p are denoted by $P\text{-class}(a, p)$; the predicate $P\text{-equiv}(a, b, p)$ stands for $P\text{-class}(a, p) = P\text{-class}(b, p)$. Specifications P-TEST₀ and P-EQUATE₀ should not require further explanation³.

P-TEST₀ ($a: T, b: T$) $t: \mathbb{B}$
rd $p : T\text{-partition}$
post $t \Leftrightarrow P\text{-equiv}(a, b, p)$

P-EQUATE₀ ($a: T, b: T$)
wr $p : T\text{-partition}$
post $p = (\overleftarrow{p} \setminus \{P\text{-class}(a, \overleftarrow{p}), P\text{-class}(b, \overleftarrow{p})\}) \cup \{P\text{-class}(a, \overleftarrow{p}) \cup P\text{-class}(b, \overleftarrow{p})\}$

Note on data invariants. At this abstraction level, a state *is* a partition of T and thus there exists no ‘state’ in which two sets in p have a non-empty intersection. Nevertheless, this does not exclude an (inefficient) implementation of EQUATE that first copies all elements of one set into the other and then destroys the first set, thus creating intermediate sets with non-empty intersection. The correctness of this implementation can be formally justified by a reification of the state space that removes the data invariant *is-partition*(p) and adds it as a conjunct to the pre and post conditions of EQUATE. A ‘representation’

²Omitted definitions can be found in the appendix.

³Specifications are prefixed and subscripted; undecorated names are reserved for informal references to operations; collections of states and operations could be collected into VDM modules.

state is then just a set of sets in T but those sets must form a partition when the operation terminates. Data invariants can thus be considered as pre and post conditions.

2.2 Interfering Operations

Because of the (potential) concurrent execution of EQUATE, equivalence classes might be merged *during* the execution of TEST, and also during the execution of another instance of EQUATE. The high degree of interference is more apparent when equivalence classes are represented by trees (roots change, the inner structure of trees change), but interference can already be specified, hence better understood, at this abstraction level.

Consider the specification P-TEST₁ below, where the keyword **ext** indicates that p can be modified by other operations. Its *rely* condition asserts that classes only grow. This rely condition is thus an assumption about the interference of the environment (other operations) during the execution of TEST. It is interpreted as a reflexive and transitive binary relation that characterises any uninterrupted sequence of ϵ -labelled steps in a computation.

$$P\text{-grows}(p_1, p_2) \triangleq \forall a: T \cdot P\text{-class}(a, p_1) \subseteq P\text{-class}(a, p_2)$$

P-TEST₁ ($a: T, b: T$) $t: \mathbb{B}$

ext rd $p : T\text{-partition}$

rely $P\text{-grows}(\overleftarrow{p}, p)$

post $(P\text{-equiv}(a, b, \overleftarrow{p}) \Rightarrow t) \wedge (t \Rightarrow P\text{-equiv}(a, b, p))$

An operation is only required to terminate and satisfy the post condition if the pre condition holds initially *and* the rely condition holds for all ϵ -labelled steps. Interference is thus explicitly specified but, as for the specification P-TEST₀ (without interference), the important part is the input/output behaviour. The post condition now consists of

- a sufficient condition that forces t to be **true** if a and b were members of the same class when the operation started; and
- a necessary condition that allows t to be **true** only if a and b are members of the same class when the operation terminates.

The result of the operation cannot be determined in the case where the classes of a and b are initially disjoint and then merged by a concurrent execution of EQUATE. Note that, in the presence of arbitrary interference (i.e. no rely condition), the result would be absolutely unpredictable. In contrast, if no interference is allowed (i.e. a rely condition of $p = \overleftarrow{p}$), the post condition of P-TEST₁ reduces to the post condition of P-TEST₀.

The counterpart of the rely condition is the *guarantee* condition which specifies the interference to others caused by an operation; this is what other operations may rely upon. This guarantee condition is interpreted as a reflexive binary relation that holds for all π -labelled steps in a computation. Together with the post condition, it forms the commitments of the operation to its environment. There is no explicit guarantee condition

in P-TEST₁ above but the mode restriction **rd** p guarantees that no step of the TEST operation modifies p . A guarantee condition appears in P-EQUATE₁: it asserts that classes only grow and no other classes than those of a and b can be modified by steps of the operation.

```

P-EQUATE1 ( $a: T, b: T$ )
  ext wr  $p : T$ -partition
  rely  $P$ -grows( $\overleftarrow{p}, p$ )
  guar  $P$ -grows( $\overleftarrow{p}, p$ )  $\wedge$ 
    let  $rest = T \setminus (P\text{-class}(a, \overleftarrow{p}) \cup P\text{-class}(b, \overleftarrow{p}))$  in
     $\forall e \in rest \cdot P\text{-class}(e, p) = P\text{-class}(e, \overleftarrow{p})$ 
  post  $P$ -equiv( $a, b, p$ )

```

Here again, this specification can be shown –in the absence of interference– to specialise to the non-interfering case (P-EQUATE₀): the argument relies on the guarantee condition as well as the post condition.

Coexistence. Whenever two operations are intended to be executed in parallel, there is a coexistence proof obligation on their specifications by which it is verified that the interference caused by one operation is allowed by the other. In this case, several instances of TEST can be executed in parallel with several instances of EQUATE because TEST is a read-only operation and the guarantee condition of P-EQUATE₁ implies the rely conditions of P-EQUATE₁ and P-TEST₁.

2.3 Data Reification

The recording of equivalence classes in a representation based on trees is captured by the reification of partitions into forests. The concrete state contains a single component $f: T$ -forest where⁴

$$T\text{-forest} = \{f \in T \xrightarrow{m} T \mid is\text{-forest}(f)\}.$$

The type invariant $is\text{-forest}(f)$ prevents f from containing cycles; each tree in f represents an equivalence class. Those elements in the same tree as a are denoted by $F\text{-class}(a, f)$; $F\text{-equiv}(a, b, f)$ stands for $F\text{-class}(a, f) = F\text{-class}(b, f)$; $is\text{-root}(a, f)$ indicates that a is a root in f ; $ancestors(a, f)$ is the set of elements on the path from a to its root (a not included).

Partitions can be easily retrieved from forests:

$$p = \{F\text{-class}(a, f) \mid a \in T \wedge is\text{-root}(a, f)\}.$$

In a second stage of reification, the forest is implemented by an array m from T to T ; $m(a) = a$ indicates that a is a root; the set of roots in m is $rts(m)$. The new data invariant is then

⁴The VDM notation $m: A \xrightarrow{m} B$ indicates a finite map m with domain type A and range type B ; \triangleleft is the operator for domain subtraction ($\mathbf{dom} s \triangleleft m = \mathbf{dom} m \setminus s$).

```

local  $x, y: T; t: \mathbb{B}$  in
   $x, y := a, b$ 
  repeat
    F-ROOT1( $x$ ) || F-ROOT1( $y$ );
    protect  $f$  in F-TEST-AND-CONNECT1( $x, y, t$ )
  until  $t$ 
end

```

Figure 1: Decomposition of F-EQUATE₁

$$m\text{-is-forest}(m) \triangleq \text{is-forest}(\text{rts}(m) \triangleleft m)$$

and the forest retrieved from m is $fr(m)$.

$$fr : T\text{-array} \rightarrow T\text{-forest}$$

$$fr(m) \triangleq \text{rts}(m) \triangleleft m$$

pre $m\text{-is-forest}(m)$

2.4 Operation Refinement

First, the specifications of TEST and EQUATE over partitions are refined into operations over forests and a new operation CLEANUP is added. These operations are then refined into more elementary operations. The decomposition of the EQUATE operation on the forest f is shown in Figure 1; variables x , y , and t are local. F-ROOT₁(z) is an operation specified to compute the root of z in f (with result in z); F-TEST-AND-CONNECT₁(c, d, t) is an operation specified to first check (with result in t) if c and d are roots *at the same time*; in case $t = \mathbf{true}$ the operation connects c and d , otherwise it does nothing; these specifications are given in detail in Section 5. The latter is embedded in a section protected from interference, otherwise trees could be merged by a concurrent instance of EQUATE between the test of $\text{is-root}(c, f)$ and the test of $\text{is-root}(d, f)$.

3 Visible Steps

The question of granularity arises as soon as interference is discussed; a detailed introduction to this problem with examples can be found in [MP92]. In this context, the question amounts to what are the π -labelled steps in a computation. This directly affects the interpretation of the guarantee condition of an operation (and of course the rely conditions of others). As discussed in next section, this also affects the interpretation of invariants.

A visible step of an operation is one that produces values relevant to other operations. These include the initial and the final values of its shared (non-local) variables (relevant for sequential composition) but must also encompass every *public* intermediate value of its shared variables. Each occurrence of a variable in the code of an operation can be classified

as either *public* or *private* [MP92]. An occurrence of a shared variable is private if the variable cannot be accessed by a concurrently executed operation, e.g. when it appears inside mutually exclusive code sections.

In this case study, all occurrences of the array m in the code are public and thus each assignment to m is a visible step. This however does not mean that all assignment statements are executed atomically. For example, a crucial assignment statement for detecting the termination of the computation of roots is $r := (m(z) = z)$. This statement has been safely introduced in the development with the assumption that other operations may interfere (and change the truth value of $m(z) = z$) between the read and write memory accesses; r and z are local variables. Imposing atomicity for all assignment statements would generate a lot of synchronisation overhead to implement them (see e.g. [And91]). Such overheads should only be incurred when required and specified by the designer (e.g. using atomic brackets). As discussed later in this paper, the evaluation of expressions is not assumed to be atomic either: the assignment statement $t := (m(x) = x \wedge m(y) = y)$ in the actual code of EQUATE is not supposed to be executed atomically but enough synchronisation has been introduced during the design to ensure that other operations may read but not modify m when this statement is executed.

4 Invariants for Interfering Operations

A development method is helpful only if it helps master the inherent complexity of a problem. The use of rely and guarantee conditions favours local reasoning but the first versions of the specifications of TEST and EQUATE (exposed later in this paper) were still too complex, hence their subsequent modification. Essentially, the gratuitous complexity was due to the lack of invariant properties. This section first discusses the application of data invariants (in the style of VDM) to interfering operations; it then introduces a new kind of invariant property.

4.1 Data Invariants

In a first attempt, the state component was basically the array m from T to T described in Section 2 but the predicate *m-is-forest* appeared almost everywhere in the specifications of operations and in the pre conditions of auxiliary functions (*is-root*, *ancestors*, ...). Without doubt, it is much easier to view *m-is-forest* as a data invariant. This motivates the introduction of the type *T-forest*. Data invariants are helpful in the development of sequential operations and remain so in the development of interfering operations. However, their interpretation appears to differ in the two cases. In the specification of sequential operations, data invariants can be considered as implicit pre and post conditions on all operations on the state space. Since the initial and final values are the only visible values of a sequential operation, this means that data invariants are required to hold for all visible values of an operation. Remarkably, this is conceptually the same for interfering operations. Data invariants are still required to hold for all visible values; there are just more visible values than the initial and final ones.

Preservation of an invariant by visible steps can thus be considered as an implicit

guarantee condition on all operations (hence a rely condition as well). In this case study, the preservation of the invariant is ultimately verified for the assignments to m , one in EQUATE, and one in CLEANUP.

4.2 Evolution Invariant

Although helpful, data invariants are not enough. The complexity of a development can be further reduced by the use of another invariant property. It should be clear from the examples in Section 2 that the relation $P\text{-grows}(p, p')$ holds for *any* pair of states where p' follows p in a computation (no matter whether the intermediate steps are operation or environment steps). This relation between computation states can be recorded by an *evolution* invariant, that should appear just next to the data invariant, in the data part of a specification.

$$ev\text{-}T\text{-partition}(p_1, p_2) \triangleq P\text{-grows}(p_1, p_2)$$

This evolution invariant can be viewed as an implicit guarantee condition on all operations, and thus an implicit rely condition as well. Specifications then simplify. For instance, there is no explicit rely condition in P-EQUATE₂ and the guarantee condition is simpler than in P-EQUATE₁.

```

P-EQUATE2 (a: T, b: T)
  ext wr p : T-partition
  guar let rest = T \ (P-class(a,  $\overleftarrow{p}$ )  $\cup$  P-class(b,  $\overleftarrow{p}$ )) in
     $\forall e \in \text{rest} \cdot P\text{-class}(e, p) = P\text{-class}(e, \overleftarrow{p})$ 
  post P-equiv(a, b, p)

```

At the forest representation level, the evolution invariant records the fact that not only do classes grow but also that no new root is ever created. It also prevents the situation where an ancestor of an element later becomes one of its descendants; the case $\text{ancestors}(a, f_2) \not\subseteq F\text{-class}(a, f_1)$ is due to the possible merging of trees. The proof obligation that the evolution invariant on forests implies the one on partitions is easily discharged.

$$ev\text{-}T\text{-forest}(f_1, f_2) \triangleq F\text{-grows}(f_1, f_2) \wedge \text{no-new-roots}(f_1, f_2) \wedge \text{no-loop}(f_1, f_2)$$

where

$$\begin{aligned}
F\text{-grows}(f_1, f_2) &\triangleq \forall a: T \cdot F\text{-class}(a, f_1) \subseteq F\text{-class}(a, f_2) \\
\text{no-new-roots}(f_1, f_2) &\triangleq \forall a: T \cdot \text{is-root}(a, f_2) \Rightarrow \text{is-root}(a, f_1) \\
\text{no-loop}(f_1, f_2) &\triangleq \forall a: T \cdot \text{ancestors}(a, f_2) \cap F\text{-class}(a, f_1) \subseteq \text{ancestors}(a, f_1)
\end{aligned}$$

With computations restricted by that evolution invariant, the three operations on forests are specified below. The guarantee condition in F-EQUATE₁ ensures that equivalence classes are merged only by connecting the root of a tree to another tree. The rely condition of F-CLEANUP₁ is not an evolution invariant because some steps of CLEANUP are obviously intended to modify the inner structure of trees. The same applies to its guarantee condition (equivalence classes are untouched and nothing but $f(a)$ changes) because f can be modified by the environment in other ways.

$$\begin{aligned} \text{bodyunch}(f_1, f_2) &\triangleq \forall a: T \cdot \neg \text{is-root}(a, f_1) \Rightarrow \neg \text{is-root}(a, f_2) \wedge f_2(a) = f_1(a) \\ \text{rootunch}(f_1, f_2) &\triangleq \forall a: T \cdot \text{is-root}(a, f_2) \Leftrightarrow \text{is-root}(a, f_1) \end{aligned}$$

F-TEST₁ ($a: T, b: T$) $t: \mathbb{B}$

ext rd $f : T\text{-forest}$

post ($F\text{-equiv}(a, b, \overleftarrow{f}) \Rightarrow t$) \wedge ($t \Rightarrow F\text{-equiv}(a, b, f)$)

F-EQUATE₁ ($a: T, b: T$)

ext wr $f : T\text{-forest}$

guar $\text{bodyunch}(\overleftarrow{f}, f) \wedge$

let $\text{rest} = T \setminus (F\text{-class}(a, \overleftarrow{f}) \cup F\text{-class}(b, \overleftarrow{f}))$ **in**

$\forall e \in \text{rest} \cdot F\text{-class}(e, f) = F\text{-class}(e, \overleftarrow{f})$

post $F\text{-equiv}(a, b, f)$

F-CLEANUP₁ ($a: T$)

ext wr $f : T\text{-forest}$

rely $\text{bodyunch}(\overleftarrow{f}, f)$

guar $\text{rootunch}(\overleftarrow{f}, f) \wedge \{a\} \triangleleft f = \{a\} \triangleleft \overleftarrow{f}$

post $\neg \text{is-root}(a, \overleftarrow{f}) \wedge \neg \text{is-root}(\overleftarrow{f}(a), \overleftarrow{f}) \Rightarrow f(a) \neq \overleftarrow{f}(a)$

The evolution invariant at the array representation level just mimics the previous one:

$$\text{ev-}T\text{-array}(m_1, m_2) \triangleq \text{ev-}T\text{-forest}(\text{fr}(m_1), \text{fr}(m_2))$$

Evolution invariants are not a novelty /em per se. Predicates that appear in the rely and guarantee conditions of all operations were already emphasized in [Stø91] (called there binary invariants). As explained in Section 5, there are advantages in moving them from the specifications of individual operations into the specification of the shared state. But the idea that properties of all computations can be attached to the definition of a state is not new either. The state specification modules of [Mid93] include a dynamic constraint which is a temporal formula. Interestingly—in the detailed case study of [Mid93]—the temporal formula has precisely the form of an evolution invariant.

5 Writing Specifications

Based on lessons learned from the case study, this section presents a few guidelines on writing specifications. To understand their impact, there is some incentive to include a ‘bad’ specification, very similar to one of the very first ones written during this case study. This is a specification of the ROOT operation. The first mistake was that the level of abstraction was wrong: ROOT was immediately viewed as an operation on the array m .

5.1 Usefulness of the Invariants

Data and evolution invariants do not increase the expressive power of specifications because they can be otherwise incorporated into specifications. Indeed, the data invariant holds initially (pre condition), is preserved by visible steps (rely and guar condition) and thus holds upon termination (post condition); the evolution invariant holds for every pair of visible steps (rely and guar condition) and by transitivity holds upon termination (post condition).

Yet, data and evolution invariants are not just syntactic sugar. They both bring insight into the problem. Having those invariants in mind helps the process of writing specifications. Interesting properties can also be deduced from the invariant. A typical example is the irreversibility of the transformation of forests:

$$\frac{\begin{array}{c} f_1, f_2, f_3: T\text{-forest} \\ ev\text{-}T\text{-forest}(f_1, f_2), ev\text{-}T\text{-forest}(f_2, f_3) \\ f_3 = f_1 \end{array}}{f_2 = f_1}$$

The many roles of the evolution invariant (rely, guar, and post conditions) are especially useful in proofs. The first premise of most proofs is in fact a list of state components, e.g. $f_0, f_1, f_2: T\text{-forest}$. The states in consideration can be the initial state, the intermediate state in a sequential composition, the states before and after a visible step of the operation (proof obligations for guarantee condition), a potential final state and a new one due to interference, In all cases, they represent successive states in a computation and this means that

$$ev\text{-}T\text{-forest}(f_0, f_1), ev\text{-}T\text{-forest}(f_1, f_2), ev\text{-}T\text{-forest}(f_0, f_2)$$

can be freely used anywhere in the proof, just like

$$is\text{-forest}(f_0), is\text{-forest}(f_1), is\text{-forest}(f_2)$$

can be. Automatic inheritance of those predicates is convenient in proofs. A typical example from this development is when $F\text{-equiv}(a, b, f_1)$ holds after the execution of a suboperation and $F\text{-equiv}(a, b, f_2)$ is required to hold after the execution of some other suboperation which, together with the environment, transforms f_1 into f_2 . This easily follows from $F\text{-grows}(f_1, f_2)$; another frequent case is the deduction of $(is\text{-root}(a, f_2) \Rightarrow t)$ from $(is\text{-root}(a, f_1) \Rightarrow t)$ which follows from $no\text{-new-roots}(f_1, f_2)$. Without explicit invariants, those predicates would have had to have been reconstructed separately from the guarantee conditions of the sub-operations and from the overall rely condition.

In conclusion, although data and evolution invariants could be incorporated in the individual specifications of the operations, what eases the development process is precisely *avoiding* thinking about them in terms of assumptions and commitments. Invariants should be considered as *given* and available for free use in writing and reasoning about specifications. The same philosophy is adopted in [MV94]: the use of invariants in the design should be separate from their ultimate verification. How the latter is carried out is addressed in Section 6

5.2 Enriched Mode Restrictions

Write-mode restrictions on variables can be understood as commitments of the operation: no other variables can be modified. Read-mode can be interpreted in several ways [Bic94]; in this paper, all variables that can be accessed but not modified by the operation are required to appear with read-mode; non-mentioned variables cannot be accessed by the operation. The mode restrictions also play a syntactic role: only the variables in write-mode can be hooked in post conditions of sequential operations. However, in the presence of interference, it makes sense to use the hooked version of read-mode variables in post conditions because these might have been modified by the environment during the execution; P-TEST₁ in Section 2 is a typical example. This in fact reveals an asymmetry in the use of mode restrictions: they give commitments of the operations but no assumptions on the environment. To compensate for this, the **rd** and **wr** mode restrictions are enriched with:

- the keyword **ext** (external) if the variable can be modified by the environment;
- the keyword **ptc** (protected) if the variable can be accessed but not modified by the environment;
- the keyword **prv** (private) if the variable cannot be accessed by the environment.

The result variables of an operation are implicitly of mode **prv wr**. The use of **ext** and **ptc** mode restrictions was already advocated in [Stø91]; the novelty is the explicit distinction between protected and private variables.

The specifications F-ROOT₁ and F-TEST-AND-CONNECT₁ below illustrate the use of mode restrictions. The former must be used in a context where z is private and the latter must be used in a context where f is protected; the decomposition of F-EQUATE₁ in Figure 1 provides such a context. In this case, there are concurrent instances of ROOT but the variables x and y match a **prv** mode because each of the two concurrent instances of ROOT manipulates only one of these variables.

F-ROOT₁

ext rd $f : T\text{-forest}$

prv wr $z : T$

post $F\text{-equiv}(\overleftarrow{z}, z, f) \wedge \text{is-root}(z, \overleftarrow{f})$

F-TEST-AND-CONNECT₁ ($c, d : T$) $t : \mathbb{B}$

ptc rd $f : T\text{-forest}$

guar $\text{bodyunch}(\overleftarrow{f}, f) \wedge$

let $\text{rest} = T \setminus (F\text{-class}(c, \overleftarrow{f}) \cup F\text{-class}(d, \overleftarrow{f}))$ **in**

$\forall e \in \text{rest} \cdot F\text{-class}(e, f) = F\text{-class}(e, \overleftarrow{f})$

post $(t \Leftrightarrow \text{is-root}(c, \overleftarrow{f}) \wedge \text{is-root}(d, \overleftarrow{f})) \wedge (t \Rightarrow F\text{-equiv}(c, d, f))$

With richer mode restrictions, information on interference can be better organised. First of all, only external variables have to be taken into account when the effect of environment steps has to be considered (e.g. in writing post conditions or in the proof obligations related to interference). Tool-supported proof obligations also become simpler because mode restrictions identify which variables are kept unchanged by environment and/or operation steps and automatic substitution of equals simplifies proofs significantly.

Mode restrictions also play a syntactic role by restricting the set of variables whose names may occur free in the various parts of a specification. Far from being exhaustive, it is first observed that protected variables should not appear hooked in rely conditions because none can be modified by the environment. Private variables should not appear in the rely nor the guarantee conditions because those characterise visible steps and the intermediate value of private variables are invisible to other operations. As one would hope, this implies that operations on private variables have pre and post conditions only; these are indeed sequential operations and thus sequential reasoning should be the standard.

5.3 Predominance of the Post Condition

Since both the guarantee and the post condition are commitments of the operations, there can be a debate about where to put some information. For the considered class of problems (when the input/output behaviour is more important than the reactive behaviour), preference should be given to the post condition. In other words, the guarantee condition should be used for what it is intended, i.e the commitments of the operation to interference, nothing else. This is partially enforced by the syntactic constraints due to mode restrictions (no private variables in the guarantee condition).

5.4 Interference and Post Conditions

Because of interference from other operations, the post condition of an operation which has to tolerate interference is often weaker and more sophisticated than that for the sequential version. This is illustrated by P-TEST₁ in Section 2. Post conditions expressed with the same pattern as in P-TEST₁ (necessary and sufficient conditions) often occurred in the case study. An example is F-TEST-ROOT₁: interfering operations that destroy roots (as allowed by the evolution invariant) influence the result of this operation.

```
F-TEST-ROOT1 (a: T) t:ℬ
  ext rd f : T-forest
  post (is-root(a, f) ⇒ t) ∧ (t ⇒ is-root(a,  $\overleftarrow{f}$ ))
```

For the same reason, the ROOT operation can only ‘approximate’ the root of an element. The post condition in F-ROOT₁ above only requires z to be a root when the computation started. Indeed, even if the process of going up in the tree stops because z is tested to be a root, it might not be a root when the operation terminates.

The preservation of the post condition by interference is probably the most important single proof obligation on specifications. Most significantly, it reveals potential errors in specifications. In the case of the ROOT operation, the failure of the proof obligation

$$\frac{f_1: T\text{-forest}, f_2: T\text{-forest} \quad \begin{array}{l} ev\text{-}T\text{-forest}(f_1, f_2) \\ is\text{-root}(z, f_1) \end{array}}{is\text{-root}(z, f_2)} \quad (\text{fail})$$

reveals that the post condition $is\text{-root}(z, f)$ would be erroneous (in this case, there is no rely condition and interference is fully specified by the evolution invariant). Evolution invariants are thus often very helpful in writing correct specifications, and consequently in detecting wrong ones.

5.5 Reasoning about Specifications

The proof that the post-condition is preserved by interference creates confidence in the specification. But, as for the specifications of sequential operations, more confidence can be gained by establishing further properties of specifications. A typical check for interfering operations is to consider how the post condition simplifies in the case of less interference. For instance, the fact that z is the root of c in f easily follows from the post condition of F-ROOT₁ if f is not subject to interference. A less trivial example is given by the specification F-CLEANUP₁. In case the class of a is merged with another class during its execution, CLEANUP might connect a to an element in that new class. But suppose that the equivalence class of a is preserved throughout the computation (*ii*); then, one may verify that the operation effectively shortens the path from a to its root (v), if possible (*iii*). Premise (*iv*) is the post condition of F-CLEANUP₁. This validation thus additionally shows that the evolution invariant (*i*) can also be thought of as a post condition.

$$\frac{\begin{array}{l} (i) \quad a: T; f_0, f_1: T, ev\text{-}T\text{-forest}(f_0, f_1) \\ (ii) \quad F\text{-class}(a, f_1) = F\text{-class}(a, f_0) \\ (iii) \quad \neg is\text{-root}(a, f_0) \wedge \neg is\text{-root}(f_0(a), f_0) \\ (iv) \quad \neg is\text{-root}(a, f_0) \wedge \neg is\text{-root}(f_0(a), f_0) \Rightarrow f_1(a) \neq f_0(a) \end{array}}{(v) \quad f_1(a) \in ancestors(f_0(a), f_0)}$$

The proof is as follows:

(1) $ancestors(a, f_1) \subseteq F\text{-class}(a, f_1)$	by (<i>i</i>), def(s).
(2) $no\text{-loop}(f_0, f_1) \wedge no\text{-new-roots}(f_0, f_1)$	by (<i>i</i>), def(s).
(3) $ancestors(a, f_1) \cap F\text{-class}(a, f_0) \subseteq ancestors(a, f_0)$	by (2), def(s).
(4) $ancestors(a, f_1) \cap F\text{-class}(a, f_1) \subseteq ancestors(a, f_0)$	by (<i>ii</i>), (3)
(5) $ancestors(a, f_1) \subseteq ancestors(a, f_0)$	by (1), (3)
(6) $\neg is\text{-root}(a, f_1)$	by (<i>iii</i>), (2)
(7) $f_1(a) \in ancestors(a, f_1)$	by (<i>i</i>), (6), def(s).
(8) $f_1(a) \in ancestors(a, f_0)$	by (5), (7)
(9) $f_1(a) \neq f_0(a)$	by (<i>iii</i>), (<i>iv</i>)
(10) $ancestors(a, f_0) = \{f_0(a)\} \cup ancestors(f_0(a), f_0)$	by (<i>i</i>), (<i>iii</i>), def(s).
(v) $f_1(a) \in ancestors(f_0(a), f_0)$	by (8), (9), (10)

All specifications make an intensive use of auxiliary functions (*no-loop*, *ancestors*, *bodyunch*, ...). It is recommended [Jon79] to use them to develop a ‘theory’ of the data types involved. This not only simplifies proofs but also improves the designer’s understanding of the problem.

5.6 Transitivity

The verification that the evolution invariant and the rely conditions are transitive is another useful proof obligation. An error in the development was spotted quite late because that proof obligation had been postponed. Indeed, the predicate *no-loop* prevents the situation where the computation of roots does not terminate because of interfering operations that, for example, first connect an element a to an element b , then connect b to a . In a first development without evolution invariants, the rely condition was

$$F\text{-grows}(\overleftarrow{f}, f) \wedge \text{no-new-roots}(\overleftarrow{f}, f) \wedge \\ (\forall a, b: T \cdot a \in \text{ancestors}(b, \overleftarrow{f}) \Rightarrow b \notin \text{ancestors}(a, f))$$

but this fails to prevent that situation because it is not transitive.

6 Towards Code

The previous section was devoted to guidelines on writing specifications. How a specification is written obviously influences its subsequent development in that the sub-operations are often designed from it. This section is devoted to further comments on the development of specifications towards code. This can only be subjective and incomplete, if only because comments that are not specific to the development of interfering operations are not included.

6.1 Control over interference

As illustrated by the examples in previous sections, specification of interference is part of the design method. But not only can interference be specified; it can also be *controlled*. The search for the most adequate mechanisms to control interference is out of the scope of this paper but some are of course needed in the examples. This case study uses the **protect** mechanism that prevents the environment of an operation from modifying state components (no ϵ -labelled step modifies them). This mechanism is not assumed to be part of the programming language, and the decision of how to implement it has in fact been postponed.

The protected section of Figure 1 (Section 2) is developed into the pseudo-code of Figure 2. Protection prevents other operations from modifying m and this ensures that

1. the two accesses to m return the same value in the expression $m(x) = x \wedge m(y) = y$;
2. x and y are still roots in m when the connection occurs.

```

protect  $m$  in
   $t := (m(x) = x \wedge m(y) = y);$ 
  if  $t \wedge x \neq y$  then  $m(x) := y$  endif;
end

```

Figure 2: Pseudo-code with critical sections

Nevertheless, m can still be accessed (but not modified) by other operations (e.g. TEST), even between the two accesses to m in the Boolean expression. Thus, the assignment statements in Figure 2 are not assumed to be executed atomically.

Critical sections are well known in concurrent programming (e.g. [And91]). The key issue is that such critical sections do not appear all of a sudden in the final code. They can be introduced *during the design*. This **protect** mechanism has been introduced (cf. Figure 1) in the early refinement of EQUATE(a, b) before the specifications F-ROOT₁ and F-TEST-AND-CONNECT₁ were further developed. This control information is recorded by the mode restrictions introduced in Section 5 and thus propagates through the design to the final code. Whether the occurrence of a variable in the code is protected or not thus follows from the design.

Control over Granularity. The **protect** mechanism does not enforce mutual exclusion in that other operations have read-access to the shared state components. If mutual exclusion (or atomic execution of an assignment statement) was required, then this should also be introduced explicitly during the design. Such a mechanism was in fact introduced in a first attempt to implement TEST-AND-CONNECT but this appeared to be a bad design decision. Indeed, if m appears in any section where read access is forbidden, implementation of that critical section will require synchronisation overhead to be added before and after *every* access to m , including in the much executed ROOT operation.

Easiness versus efficiency. Control over interference can be necessary: roots should not be connected by other operations between the ‘test’ and ‘connect’ parts in Figure 2: protecting each part separately is not sufficient. At the other extreme, the development of EQUATE would have been easier if the whole body of the operation was under the scope of a **protect** mechanism. This would however drastically restrict concurrency! In this development, the computation of roots, which is probably the most time-consuming part of the execution of EQUATE, can be executed concurrently with any other operation.

Suppose that **protect** is implemented by a readers and writers protocol⁵. Then the only synchronisation overhead is: a reader protocol around one test in TEST (after the computation of roots), a writer protocol around the code for TEST-AND-CONNECT inside EQUATE, and a writer protocol around the only assignment statement of CLEANUP. There is no synchronisation overhead in the computation of roots.

The writer protocol around the assignment in CLEANUP is of special interest. Its presence is due to the implementation of the **protect** mechanism in other operations.

⁵Details in the appendix.

This mechanism made the development of TEST-AND-CONNECT easier, but the loss of efficiency in CLEANUP is excessive. In fact, only roots need to be protected and the guarantee condition in F-CLEANUP₁ tells us that roots are unchanged. Thus, on the one hand, the current formal development gives enough confidence for a safe removal of the synchronisation overhead in CLEANUP. But, on the other hand, it is unclear how to do it formally, in a cost-effective way.

Synchronisation and Compositionality. When the concurrent execution of several instances of EQUATE was first considered, it seemed that the addition of *explicit* synchronisation variables between the operations might be required. A fully compositional development indeed requires each operation to be developed independently down to machine code. But an attempt to add explicit synchronisation variables was quickly abandoned, first because it was unclear how to choose the variables, and second because this would have implied adding all ‘protocol information’ in specifications and carry all those complications through the development. An easier development that ends up with (perhaps less efficient) pseudo-code like that in Figure 2 is preferred.

6.2 Introduction of Code

As illustrated in Figure 1, language constructs (loop, ‘;’, assignment statements) appear early in the development. This of course biases the development towards imperative programming languages, but those are the target languages, at least for the code of the individual operations. How those operations are actually activated (procedure call, message passing, ...) is not considered in this development.

But the most interesting feature is the introduction of assignment statements. Most often, it is much easier to introduce an assignment statement than to describe it by a specification. A description of $x, y := a, b$ in Figure 1 with guarantee and post conditions is unnecessarily opaque. A similar remark holds for the development of the specifications M-CONNECT-TO-ANCESTOR₁ and M-CONNECT-ROOTS₁ below into the assignment statement $m(a) := b$. There is no need for any intermediate specification that would try to mimic the effect of the assignment statement in the guarantee condition. This is in accordance with the suggestion of Section 5 that the effect of an operation should be specified in the post condition rather than in the guarantee condition.

```
M-CONNECT-TO-ANCESTOR1 ( $a, b: T$ )
  ext wr  $m : T$ -array
  pre  $a \notin rts(m) \wedge b \in ancestors(a, fr(m))$ 
  rely  $bodyunch(fr(\overline{m}), fr(m))$ 
  guar  $rts(m) = rts(\overline{m}) \wedge \{a\} \triangleleft m = \{a\} \triangleleft \overline{m}$ 
  post  $m(a) = b$ 
```

```
M-CONNECT-ROOTS1 ( $a, b: T$ )
  ptc wr  $m : T$ -array
```

pre $a \neq b \wedge a \in rts(m) \wedge b \in rts(m)$
guar $bodyunch(fr(\overleftarrow{m}), fr(m)) \wedge$
 $\text{let } rest = T \setminus (F\text{-class}(a, fr(\overleftarrow{m})) \cup F\text{-class}(b, fr(\overleftarrow{m}))) \text{ in}$
 $\forall e \in rest \cdot F\text{-class}(e, fr(m)) = F\text{-class}(e, fr(\overleftarrow{m}))$
post $m = \overleftarrow{m} \dagger \{a \mapsto b\} \vee m = \overleftarrow{m} \dagger \{b \mapsto a\}$

The proof that the implementation of those specifications by $m(a) := b$ is correct proceeds by taking into account interference from the environment before m is assigned to; the interference after termination of the assignment statement has already been captured by the proof obligation on the post condition. Three values of m can then be identified: the initial value m_0 , the value m_1 just before m is assigned to, and the value m_2 just after it is assigned to. The pre condition characterises m_0 , the rely condition characterises the transitions from m_0 to m_1 , and the transition from m_1 to m_2 is characterised by $m_2 = m_1 \dagger \{a \mapsto b\}$. As usual, all transitions are also characterised by the evolution invariant and $m\text{-is-forest}(m_i)$ can be assumed for each i .

6.3 Verification of the Invariants

As illustrated by M-CONNECT-TO-ANCESTOR₂, invariants could be expanded into the individual specifications before assignment statements are introduced.

M-CONNECT-TO-ANCESTOR₂ ($a, b: T$)
ptc wr $m : T\text{-array}$
pre $m\text{-is-forest}(m) \wedge a \notin rts(m) \wedge b \in ancestors(a, fr(m))$
rely $m\text{-is-forest}(\overleftarrow{m}) \Rightarrow m\text{-is-forest}(m) \wedge ev\text{-}T\text{-array}(\overleftarrow{m}, m) \wedge bodyunch(fr(\overleftarrow{m}), fr(m))$
guar $m\text{-is-forest}(\overleftarrow{m}) \Rightarrow$
 $m\text{-is-forest}(m) \wedge ev\text{-}T\text{-array}(\overleftarrow{m}, m) \wedge rts(m) = rts(\overleftarrow{m}) \wedge \{a\} \triangleleft m = \{a\} \triangleleft m$
post $m(a) = b$

But this does not help. Keeping the invariants outside the individual specifications until code is introduced seems as easy. The preservation of invariants (between m_1 and m_2) by the assignment statement is then to be verified first. There are only two such proof obligations in this case study. The one for the implementation of M-CONNECT-TO-ANCESTOR₁ is:

$$\frac{
\begin{array}{l}
m_0, m_1, m_2: T\text{-array} \\
m\text{-is-forest}(m_0) \wedge m\text{-is-forest}(m_1) \\
ev\text{-}T\text{-array}(m_0, m_1) \\
a \notin rts(m_0) \wedge b \in ancestors(a, fr(m_0)) \\
bodyunch(fr(m_0), fr(m_1)) \\
m_2 = m_1 \dagger \{a \mapsto b\}
\end{array}
}{
m\text{-is-forest}(m_2) \wedge ev\text{-}T\text{-array}(m_1, m_2)
}$$

Once this proof obligation is discharged, the invariants can be freely used in the proof obligations for the guarantee and post conditions of M-CONNECT-TO-ANCESTOR₁. Notice that a common pattern to all proofs related to assignment statements is to first show

that the pre condition is preserved by interference, that is to show $a \notin \text{rts}(m_1) \wedge b \in \text{ancestors}(a, \text{fr}(m_1))$ in this case.

7 Conclusion

Rely and guarantee conditions have been proposed to handle concurrency while preserving local reasoning in the development. Designed for the specification of interference, these conditions can also be used in an anarchic way, by encoding as much information as possible into them. This quickly leads to intractable specifications. In contrast, despite the high level of concurrency, this development makes a rather economic use of rely and guarantee conditions: out of 11 specifications at the forest level, only 5 have an explicit guarantee condition, and only 3 have an explicit rely condition. A development that tends to generate many sophisticated rely and guarantee conditions is probably poorly organised. Of course, this remark is based on a single case study but the failure to present an elegant development in other cases would probably indicate that the specified operations fall outside the considered class of problems. In particular, the specification style in this paper does not work well with operations whose reactive behaviour is the most important feature; the use of other styles of rely/guarantee specifications for the development of a non-trivial reactive system is illustrated in [KR92].

Although rely and guarantee conditions favour local reasoning, this paper emphasizes the role of the invariants (data invariant and evolution invariant), which by nature record global information. Therefore, local reasoning is not totally enforced because each operation is not developed independently down to code: a data reification step (with strengthening of the invariants) concern all operations. But this is already the case for data reification steps in the development of sequential operations in VDM [Jon90] or B [Abr96]. The methodological importance of invariants in concurrency is not new; detailed developments based on invariants can be found –for example– in [CM88, Gri93].

As mentioned in the introduction of this paper, theoretical aspects have been intentionally neglected. Expressiveness is one of them. An attentive reader should have noticed that the only restriction to concurrency in this paper is the execution of at most one instance of CLEANUP at a time. Concurrent execution of that operation not only further complicates the development but also raises expressiveness problems: it seems that the formulation of an adequate evolution invariant then requires the use of history determined auxiliary variables. Use of auxiliary variables with rely/guarantee specifications is detailed in [GNL91, Stø91]. Auxiliary variables lead to clearer specifications than nested temporal operators, but inappropriate use can also lead to cumbersome specifications. At worse, rely and guarantee conditions could be reduced to an update of a history variable that records all transitions in a computation and the post condition be then expressed as a predicate on that history variable; guidelines for auxiliary variables are thus required.

The design of appropriate proof rules for data reification with rely/guarantee conditions is another theoretical aspect that deserves further work. Thanks to the evolution invariant, the problem of the appearance of new rely conditions with data reification [WD88] does not occur in this case study but might appear in others.

Acknowledgements

This work has been supported by funding from the UK EPSRC. We thank Ketil Stølen for his helpful comments on a draft of this paper.

References

- [Abr96] J.-R. Abrial. *The B-Book: Assigning programs to meanings*. Cambridge University Press, 1996.
- [AL93] Martin Abadi and Leslie Lamport. Composing specifications. *ACM Transactions on Programming Languages and Systems*, 15:73–132, 1993.
- [And91] Gregory R. Andrews. *Concurrent Programming: Principles and Practice*. The Benjamin/Cummings Publishing Company Inc., 1991.
- [Bic94] Juan Bicarregui. Operation semantics with read and write frames. In D. Till, editor, *Proceedings of the 6th Refinement Workshop*, pages 260–278. Springer-Verlag, 1994.
- [BJM88] R. Bloomfield, R. B. Jones, and L. S. Marshall, editors. *VDM'88: VDM – The Way Ahead*, volume 328 of *Lecture Notes in Computer Science*. Springer-Verlag, 1988.
- [BK85] Howard Barringer and Ruud Kuiper. Hierarchical development of concurrent systems in a temporal logic framework. In S.D. Brookes, A.W. Roscoe, and G. Winskel, editors, *Seminar on Concurrency*, volume 197 of *Lecture Notes in Computer Science*, pages 35–61. Springer-Verlag, 1985.
- [CM88] K. M. Chandy and J. Misra. *Parallel Program Design: A Foundation*. Addison-Wesley, 1988.
- [Col94] Pierre Collette. Composition of assumption-commitment specifications in a UNITY style. *Science of Computer Programming*, 23:107–126, 1994.
- [GNL91] Peter Grønning, Thomas Qvist Nielsen, and Hans Henrik Løvengreen. Refinement and composition of transition-based rely-guarantee specifications with auxiliary variables. In K.V. Nori and C.E. Veni Madhavan, editors, *Foundations of Software Technology and Theoretical Computer Science*, volume 472 of *Lecture Notes in Computer Science*, pages 332–348. Springer-Verlag, 1991.
- [Gri93] E. P. Gribomont. Concurrency without toil: a systematic method for parallel program design. *Science of Computer Programming*, 21:1–56, 1993.
- [Jon79] C. B. Jones. Constructing a theory of a data structure as an aid to program development. *Acta Informatica*, 11:119–137, 1979.
- [Jon81] C. B. Jones. *Development Methods for Computer Programs including a Notion of Interference*. PhD thesis, Oxford University, June 1981. Printed as: Programming Research Group, Technical Monograph 25.

- [Jon90] C. B. Jones. *Systematic Software Development using VDM*. Prentice Hall International, second edition, 1990. ISBN 0-13-880733-7.
- [JT95] Bengt Jonsson and Yih-Kuen Tsay. Assumption/guarantee specifications in linear time temporal logic. In P.D. Mosses, M. Nielsen, and M.I. Schwartzbach, editors, *TAPSOFT'95: Theory and Practice of Software Development*, volume 915 of *Lecture Notes in Computer Science*, pages 262–276. Springer-Verlag, 1995.
- [KR92] A. Kay and J. N. Reed. A rely and guarantee method for timed CSP: A specification and design of a telephone exchange. *IEEE, Transactions on Software Engineering*, 19(6):625–639, 1992.
- [MC81] J. Misra and K. M. Chandy. Proofs of networks of processes. *IEEE Transactions on Software Engineering*, 7:417–426, 1981.
- [Mid93] Cornelius A. Middelburg. *Logic and Specification: Extending VDM-SL for advanced formal specification*. Chapman and Hall, 1993.
- [MP92] Zohar Manna and Amir Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, 1992.
- [MV94] C. C. Morgan and T. Vickers. *On the Refinement Calculus*. Formal Approaches to Computing and Information Technology series (FACET). Springer-Verlag, 1994.
- [PJ91] Paritosh K. Pandya and Mathai Joseph. P-A logic — a compositional proof system for distributed programs. *Distributed Computing*, 5:27–54, 1991.
- [PT91] S. Prehn and W. J. Toetenel, editors. *VDM'91 – Formal Software Development Methods. Proceedings of the 4th International Symposium of VDM Europe, Noordwijkerhout, The Netherlands, October 1991. Vol.1: Conference Contributions*, volume 551 of *Lecture Notes in Computer Science*. Springer-Verlag, 1991.
- [Sta86] Eugene W. Stark. A proof technique for rely/guarantee properties. In S.N. Maheshwari, editor, *Foundations of Software Technology and Theoretical Computer Science*, volume 206 of *Lecture Notes in Computer Science*, pages 369–391. Springer-Verlag, 1986.
- [Stø91] K. Stølen. An Attempt to Reason About Shared-State Concurrency in the Style of VDM. In [PT91], pages 324–342, 1991.
- [WD88] J. C. P. Woodcock and B. Dickinson. Using VDM with rely and guarantee-conditions: Experiences of a real project. In [BJM88], pages 434–458, 1988.
- [ZdBdR84] Job Zwiers, Arie de Bruin, and Willem-Paul de Roever. A proof system for partial correctness of dynamic networks of processes. In E. Clarke and D. Kozen, editors, *Logics of Programs*, volume 164 of *Lecture Notes in Computer Science*, pages 513–527. Springer-Verlag, 1984.

A Technical Summary

Types and auxiliary functions

$$is-disj : T\text{-set} \times T\text{-set} \rightarrow \mathbb{B}$$

$$is-disj(s_1, s_2) \triangleq s_1 \cap s_2 = \{\}$$

$$is-partition : (T\text{-set})\text{-set} \rightarrow \mathbb{B}$$

$$is-partition(p) \triangleq \bigcup p = T \wedge \{\} \notin p \wedge (\forall s_1, s_2 \in p \cdot s_1 = s_2 \vee is-disj(s_1, s_2))$$

$$T\text{-partition} = \{p \in (T\text{-set})\text{-set} \mid is-partition(p)\}.$$

$$P\text{-class} : T \times T\text{-partition} \rightarrow T\text{-set}$$

$$P\text{-class}(a, p) \triangleq \iota s \in p \cdot a \in s$$

$$P\text{-equiv} : T \times T \times T\text{-partition} \rightarrow \mathbb{B}$$

$$P\text{-equiv}(a, b, p) \triangleq P\text{-class}(a, p) = P\text{-class}(b, p)$$

$$P\text{-grows} : T\text{-partition} \times T\text{-partition} \rightarrow \mathbb{B}$$

$$P\text{-grows}(p_1, p_2) \triangleq \forall a: T \cdot P\text{-class}(a, p_1) \subseteq P\text{-class}(a, p_2)$$

$$in-cycles : (T \xrightarrow{m} T) \rightarrow (T\text{-set})\text{-set}$$

$$in-cycles(f) \triangleq \{c: T\text{-set} \mid c \subseteq \mathbf{dom} f \wedge \forall e \in c \cdot f(e) \in c\}$$

$$is-forest : (T \xrightarrow{m} T) \rightarrow \mathbb{B}$$

$$is-forest(f) \triangleq in-cycles(f) = \{\}$$

$$T\text{-forest} = \{f \in T \xrightarrow{m} T \mid is-forest(f)\}.$$

$$F\text{-class} : T \times T\text{-forest} \rightarrow T\text{-set}$$

$$F\text{-class}(a, f) \triangleq \{b \mid root(b, f) = root(a, f)\}$$

$$F\text{-equiv} : T \times T \times T\text{-forest} \rightarrow \mathbb{B}$$

$$F\text{-equiv}(a, b, f) \triangleq F\text{-class}(a, f) = F\text{-class}(b, f)$$

$is\text{-root} : T \times T\text{-forest} \rightarrow \mathbb{B}$

$is\text{-root}(a, f) \triangleq a \notin \mathbf{dom} f$

$ancestors : T \times T\text{-forest} \rightarrow T\text{-set}$

$ancestors(a, f) \triangleq \mathbf{if} \text{ } is\text{-root}(a, f) \mathbf{ then } \{\} \mathbf{ else } f(a) \cup ancestors(f(a), f)$

$F\text{-grows} : T\text{-forest} \times T\text{-forest} \rightarrow \mathbb{B}$

$F\text{-grows}(f_1, f_2) \triangleq \forall a: T \cdot F\text{-class}(a, f_1) \subseteq F\text{-class}(a, f_2)$

$no\text{-new-roots} : T\text{-forest} \times T\text{-forest} \rightarrow \mathbb{B}$

$no\text{-new-roots}(f_1, f_2) \triangleq \forall a: T \cdot is\text{-root}(a, f_2) \Rightarrow is\text{-root}(a, f_1)$

$no\text{-loop} : T\text{-forest} \times T\text{-forest} \rightarrow \mathbb{B}$

$no\text{-loop}(f_1, f_2) \triangleq \forall a: T \cdot ancestors(a, f_2) \cap F\text{-class}(a, f_1) \subseteq ancestors(a, f_1)$

$bodyunch : T\text{-forest} \times T\text{-forest} \rightarrow \mathbb{B}$

$bodyunch(f_1, f_2) \triangleq \forall a: T \cdot \neg is\text{-root}(a, f_1) \Rightarrow \neg is\text{-root}(a, f_2) \wedge f_2(a) = f_1(a)$

$rootunch : T\text{-forest} \times T\text{-forest} \rightarrow \mathbb{B}$

$rootunch(f_1, f_2) \triangleq \forall a: T \cdot is\text{-root}(a, f_1) \Leftrightarrow is\text{-root}(a, f_2)$

$T\text{-array} = \{m \in T \xrightarrow{m} T \mid \mathbf{dom} m = T\}.$

$rts : T\text{-array} \rightarrow T\text{-set}$

$rts(m) \triangleq \{a: T \mid m(a) = a\}$

$m\text{-is-forest} : T\text{-array} \rightarrow \mathbb{B}$

$m\text{-is-forest}(m) \triangleq is\text{-forest}(rts(m) \triangleleft m)$

$fr : T\text{-array} \rightarrow T\text{-forest}$

$fr(m) \triangleq rts(m) \triangleleft m$

pre $m\text{-is-forest}(m)$

Specifications

P-TEST₂ ($a: T, b: T$) $t: \mathbb{B}$

ext rd $p : T\text{-partition}$

post $(P\text{-equiv}(a, b, \overleftarrow{p}) \Rightarrow t) \wedge (t \Rightarrow P\text{-equiv}(a, b, p))$

P-EQUATE₂ ($a: T, b: T$)

ext wr $p : T\text{-partition}$

guar let $rest = T \setminus (P\text{-class}(a, \overleftarrow{p}) \cup P\text{-class}(b, \overleftarrow{p}))$ **in**

$\forall e \in rest \cdot P\text{-class}(e, p) = P\text{-class}(e, \overleftarrow{p})$

post $P\text{-equiv}(a, b, p)$

F-TEST₁ ($a: T, b: T$) $t: \mathbb{B}$

ext rd $f : T\text{-forest}$

post $(F\text{-equiv}(a, b, \overleftarrow{f}) \Rightarrow t) \wedge (t \Rightarrow F\text{-equiv}(a, b, f))$

F-EQUATE₁ ($a: T, b: T$)

ext wr $f : T\text{-forest}$

guar $bodyunch(\overleftarrow{f}, f) \wedge$

let $rest = T \setminus (F\text{-class}(a, \overleftarrow{f}) \cup F\text{-class}(b, \overleftarrow{f}))$ **in**

$\forall e \in rest \cdot F\text{-class}(e, f) = F\text{-class}(e, \overleftarrow{f})$

post $F\text{-equiv}(a, b, f)$

F-CLEANUP₁ ($a: T$)

ext wr $f : T\text{-forest}$

rely $bodyunch(\overleftarrow{f}, f)$

guar $rootunch(\overleftarrow{f}, f) \wedge \{a\} \triangleleft f = \{a\} \triangleleft \overleftarrow{f}$

post $\neg is\text{-root}(a, \overleftarrow{f}) \wedge \neg is\text{-root}(\overleftarrow{f}(a), \overleftarrow{f}) \Rightarrow f(a) \neq \overleftarrow{f}(a)$

F-ROOT₁

ext rd $f : T\text{-forest}$

prv wr $z : T$

post $F\text{-equiv}(\overleftarrow{z}, z, f) \wedge is\text{-root}(z, \overleftarrow{f})$

F-TEST-2-ROOTS₁ ($a, b: T$) $t: \mathbb{B}$

ptc rd $f : T\text{-forest}$

post $t \Leftrightarrow is\text{-root}(a, f) \wedge is\text{-root}(b, f)$

F-TEST-ROOT₁ ($a: T$) $t: \mathbb{B}$

ext rd $f : T\text{-forest}$

post $(is\text{-root}(a, f) \Rightarrow t) \wedge (t \Rightarrow is\text{-root}(a, \overleftarrow{f}))$

F-GO-UP₁

ext rd $f : T\text{-forest}$

prv wr $x : T$

pre $\neg is\text{-root}(x, f)$

post $F\text{-equiv}(\overleftarrow{x}, x, f) \wedge (F\text{-equiv}(\overleftarrow{x}, x, \overleftarrow{f}) \Rightarrow x \in ancestors(\overleftarrow{x}, \overleftarrow{f}))$

F-FATHER₁ ($a: T$) $x: T$

ext rd $f : T\text{-forest}$

pre $\neg is\text{-root}(a, f)$

rely $bodyunch(\overleftarrow{f}, f)$

post $x = \overleftarrow{f}(a)$

F-CONNECT-TO-ANCESTOR₁ ($a, b: T$)

ext wr $f : T\text{-forest}$

pre $\neg is\text{-root}(a, f) \wedge b \in ancestors(a, f)$

rely $bodyunch(\overleftarrow{f}, f)$

guar $rootunch(\overleftarrow{f}, f) \wedge \{a\} \triangleleft f = \{a\} \triangleleft \overleftarrow{f}$

post $f(a) = b$

F-TEST-AND-CONNECT₁ ($c, d: T$) $t: \mathbb{B}$

ptc rd $f : T\text{-forest}$

guar $bodyunch(\overleftarrow{f}, f) \wedge$

let $rest = T \setminus (F\text{-class}(c, \overleftarrow{f}) \cup F\text{-class}(d, \overleftarrow{f}))$ **in**

$\forall e \in rest \cdot F\text{-class}(e, f) = F\text{-class}(e, \overleftarrow{f})$

post $(t \Leftrightarrow is\text{-root}(c, \overleftarrow{f}) \wedge is\text{-root}(d, \overleftarrow{f})) \wedge (t \Rightarrow F\text{-equiv}(c, d, f))$

F-CONNECT-ROOTS₁ ($a: T, b: T$)

ptc wr $f : T\text{-forest}$

pre $a \neq b \wedge is\text{-root}(a, f) \wedge is\text{-root}(b, f)$

guar $bodyunch(\overleftarrow{f}, f) \wedge$

let $rest = T \setminus (F\text{-class}(a, \overleftarrow{f}) \cup F\text{-class}(b, \overleftarrow{f}))$ **in**

$\forall e \in rest \cdot F\text{-class}(e, f) = F\text{-class}(e, \overleftarrow{f})$

post $f = \overleftarrow{f} \dagger \{a \mapsto b\} \vee f = \overleftarrow{f} \dagger \{b \mapsto a\}$

Refinements

Refinement of F-TEST₁(a, b, t):

```
local  $x, y: T$  in
   $x, y := a, b$ ;
  repeat
    F-ROOT1( $x$ ) || F-ROOT1( $y$ );
     $t := (x = y)$ 
  until  $t \vee (r \text{ from } (\text{protect } f \text{ in F-TEST-2-ROOTS}_1(x, y, r)))$ 
end
```

Refinement of F-EQUATE₁(a, b):

```
local  $x, y: T; t: \mathbb{B}$  in
   $x, y := a, b$ 
  repeat
    F-ROOT1( $x$ ) || F-ROOT1( $y$ );
    protect  $f$  in F-TEST-AND-CONNECT1( $x, y, t$ )
  until  $t$ 
end
```

Refinement of F-ROOT₁(x):

```
while  $\neg t$  from F-TEST-ROOT1( $x, t$ )
  do
    F-GO-UP1( $x$ )
  od
```

Refinement of F-TEST-AND-CONNECT₁(a, b, t):

```
F-TEST-2-ROOTS1( $a, b, t$ );
if  $t \wedge a \neq b$ 
  then F-CONNECT-ROOTS1( $a, b$ )
```

Refinement of F-CLEANUP₁(a):

```
local  $x: T$  in
  if  $\neg t$  from F-TEST-ROOT1( $a, t$ )
  then F-FATHER1( $a, x$ );
    if  $\neg t$  from F-TEST-ROOT1( $x, t$ )
    then F-FATHER1( $x, x$ );
      F-CONNECT-TO-ANCESTOR1( $a, x$ )
  end
```

Specifications

```
M-TEST-2-ROOTS1 ( $a, b: T$ )  $t: \mathbb{B}$ 
  ptc rd  $m : T\text{-array}$ 
  post  $t \Leftrightarrow a \in rts(m) \wedge b \in rts(m)$ 
```

M-TEST-ROOT₁ ($a: T$) **wr** $t: \mathbb{B}$
ext rd $m : T\text{-array}$
post $(a \in rts(m) \Rightarrow t) \wedge (t \Rightarrow a \in rts(m))$

M-GO-UP₁
ext rd $m : T\text{-array}$
prv wr $x : T$
pre $x \notin rts(m)$
post $F\text{-equiv}(\overleftarrow{x}, x, fr(m)) \wedge (F\text{-equiv}(\overleftarrow{x}, x, fr(\overleftarrow{m})) \Rightarrow x \in ancestors(\overleftarrow{x}, fr(\overleftarrow{m})))$

M-FATHER₁ ($a: T$) $x: T$
ext rd $m : T\text{-array}$
pre $a \notin rts(m)$
rely $bodyunch(fr(\overleftarrow{m}), fr(m))$
post $x = \overleftarrow{m}(a)$

M-CONNECT-TO-ANCESTOR₁ ($a, b: T$)
ext wr $m : T\text{-array}$
pre $a \notin rts(m) \wedge b \in ancestors(a, fr(m))$
rely $bodyunch(fr(\overleftarrow{m}), fr(m))$
guar $rts(m) = rts(\overleftarrow{m}) \wedge \{a\} \triangleleft m = \{a\} \triangleleft \overleftarrow{m}$
post $m(a) = b$

M-CONNECT-ROOTS₁ ($a, b: T$)
ptc wr $m : T\text{-array}$
pre $a \neq b \wedge a \in rts(m) \wedge b \in rts(m)$
guar $bodyunch(fr(\overleftarrow{m}), fr(m)) \wedge$
 $\mathbf{let} \text{ rest} = T \setminus (F\text{-class}(a, fr(\overleftarrow{m})) \cup F\text{-class}(b, fr(\overleftarrow{m}))) \mathbf{in}$
 $\forall e \in \text{rest} \cdot F\text{-class}(e, fr(m)) = F\text{-class}(e, fr(\overleftarrow{m}))$
post $m = \overleftarrow{m} \dagger \{a \mapsto b\} \vee m = \overleftarrow{m} \dagger \{b \mapsto a\}$

Refinements

Refinement of M-TEST-2-ROOT₁:

$$t := (m(a) = a) \wedge (m(b) = b)$$

Refinement of M-TEST-ROOT₁:

$$t := m(a) = a$$

Refinement of M-GO-UP₁:

$x := m(x)$

Refinement of M-FATHER₁:

$x := m(a)$

Refinement of M-CONNECT-TO-ANCESTOR₁:

$m(a) := b$

Refinement of M-CONNECT-ROOTS₁:

$m(a) := b$

Code for the operations

TEST(a, b): t

local $x, y: T, r: \mathbb{B}$ **in**

$x, y := a, b;$

repeat

 ROOT(x) || ROOT(y);

$t := (x = y);$

$\frac{\text{reader-entry-protocol}}{r := m(x) = x \wedge m(y) = y}$

$\frac{\text{reader-exit-protocol}}{r := m(x) = x \wedge m(y) = y}$

until $t \vee r$

end

CLEANUP(a)

local $x: T, t: \mathbb{B}$ **in**

$t := m(a) = a;$

if $\neg t$

then

$x := m(a);$

$t := m(x) = x;$

if $\neg t$

then

$x := m(x);$

$\frac{\text{writer-entry-protocol}}{m(a) := x}$

$m(a) := x$

$\frac{\text{writer-exit-protocol}}{m(a) := x}$

end

EQUATE(a, b)

local $x, y: T; t: \mathbb{B}$ **in**

$x, y := a, b;$

repeat

 ROOT(x) || ROOT(y);

$\frac{\text{writer-entry-protocol}}{t := (m(x) = x \wedge m(y) = y);}$

$t := (m(x) = x \wedge m(y) = y);$

if $t \wedge x \neq y$

then $m(x) := y$

$\frac{\text{writer-exit-protocol}}{t := (m(x) = x \wedge m(y) = y);}$

until t

end

ROOT($\text{var } z$)

local $t: \mathbb{B}$ **in**

$t := m(z) = z;$

while $\neg t$

do

$z := m(z);$

$t := m(z) = z$

od

end