

Kent Academic Repository

Full text document (pdf)

Citation for published version

Boiten, Eerke Albert and Derrick, John and Schellhorn, Gerhard (2008) Relational Concurrent Refinement II: Internal Operations and Outputs. *Formal Aspects of Computing*, 21 (1-2). pp. 65-102. ISSN 0934-5043.

DOI

<https://doi.org/10.1007/s00165-007-0066-z>

Link to record in KAR

<http://kar.kent.ac.uk/14524/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Relational Concurrent Refinement Part II: Internal Operations and Outputs

Eerke Boiten¹, John Derrick² and Gerhard Schellhorn³

¹ Computing Laboratory, University of Kent, Canterbury, Kent, CT2 7NF, UK

² Department of Computer Science, University of Sheffield, Sheffield, S1 4DP, UK

³ Department of Computer Science, Institute of Software Engineering and Programming Languages, University of Augsburg, 86135 Augsburg, Germany.

Abstract. Two styles of description arise naturally in formal specification: state-based and behavioural. In state-based notations, a system is characterised by a collection of variables, and their values determine which actions may occur throughout a system history. Behavioural specifications describe the chronologies of actions – interactions between a system and its environment. The exact nature of such interactions is captured in a variety of semantic models with corresponding notions of refinement; refinement in state based systems is based on the semantics of sequential programs and is modelled relationally. Acknowledging that these viewpoints are complementary, substantial research has gone into combining the paradigms.

The purpose of this paper is to do three things. First, we survey recent results linking the relational model of refinement to the process algebraic models. Specifically, we detail how variations in the relational framework lead to relational data refinement being in correspondence with traces-divergences, singleton failures and failures-divergences refinement in a process semantics.

Second, we generalise these results by providing a general flexible scheme for incorporating the two main “erroneous” concurrent behaviours: deadlock and divergence, into relational refinement. This is shown to subsume previous characterisations. In doing this we derive relational refinement rules for specifications containing both internal operations and outputs that corresponds to failures-divergences refinement. Thirdly, the theory has been formally specified and verified using the interactive theorem prover KIV.

Keywords: Data refinement, Z, simulations, process algebraic semantics, failures-divergences refinement, deadlock, internal operations, outputs, mechanisation, KIV.

1. Introduction

Two styles of description arise naturally in formal specification: state-based and behavioural specification. In state-based notations, such as Z [43], a system is characterised by a collection of variables. The system's operations are defined by their effects on the variables, and conditions on their values determine when operations are applicable.

In behavioural notations, e.g., process algebras such as CSP [27, 36], the allowable chronologies of actions are specified explicitly. Moreover, actions actually represent *interactions* between a system and its environment. The exact way in which the environment is allowed to interact with the system varies between notations, or to be more precise: between concurrency semantics. Typical semantics are set-based, associating one or more sets with each process, for example traces, refusals, divergences. Refinement is then defined in terms of set inclusions and equalities between the corresponding sets for different processes. For example, in CSP one could use trace refinement, failures refinement or failures-divergences refinement [36]. In CCS, bisimulation is typically used [33], whereas in LOTOS reduction, extension and conformance are defined [5]. A survey of many of the most prominent refinement relations is given in [45].

On the other hand, in state based systems specifications are considered to define abstract data types (ADTs), consisting of an initialisation, a collection of operations and a finalisation, where a program over an ADT is a sequential composition of these elements. Refinement in this context is taken to be the subset relation over program behaviours, where what is deemed visible (i.e., the domain of the initialisation and the range of the finalisation) is the input/output relation. Thus an ADT C refines an ADT A if for every program and sequence of inputs, the outputs that C produces are outputs that A could also have produced. This definition of refinement quantifies over program behaviour and *simulations* have become the accepted approach to make verification of refinements tractable [16]. Two different forms of simulations are needed to provide a complete method: downward and upward simulations. Theoretical background is given in [16], and examples of their use in Z are given in [46, 18].

Motivated by both theoretical comparisons of refinement and integrations of specification languages, there has been recent interest in relating these two viewpoints of refinement. That is, in order to understand the nature and structure of refinement, as well as provide a means to combine languages and their development methodologies, it is necessary to understand the correspondence between data refinement and process refinement relations.

The purpose of this paper is three-fold. First, we survey this existing work linking relational models of refinement to their process algebraic counterparts. Second, we extend these results, and thirdly we verify the theory mechanically. In doing so, and as a particular byproduct of the theory, we derive simulation rules for relational data refinement of specifications containing both internal operations and outputs.

Work relating the two paradigms includes Josephs [30], He [29], Woodcock and Morgan [47], Bolton and Davies [7, 8], Derrick and Boiten [4, 19] and Schneider [40]. That due to Josephs [30], He [29], Woodcock and Morgan [47] defines a basic correspondence between simulation rules and failures-divergences refinement. The more recent work of Bolton and Davies [7, 8], Derrick and Boiten [4, 19] and Schneider [40] investigates a direct correspondence between the relational model and process semantics, and includes specific consideration of input and output which introduces some subtleties.

The correspondence between a relational model and a process model can be investigated either by defining a 'corresponding process' in, say, CSP for each ADT, and then deriving the process semantics, or by defining a process semantics directly for an ADT. Schneider [40] and Bolton and Davies [7, 8] do the former, whilst Derrick and Boiten the latter [19]. Either way, a process semantics $\llbracket A \rrbracket$ can be given for an ADT A . The central aim is to derive results of the following form:

In relational model X , $A \sqsubseteq_{data} C$ if and only if $\llbracket A \rrbracket \sqsubseteq_{ps} \llbracket C \rrbracket$.

where \sqsubseteq_{data} denotes relational data refinement, and \sqsubseteq_{ps} the refinement relation induced by the given process semantics. Varying X , and to some extent \sqsubseteq_{data} , gives different process semantics. In particular, we are

interested in the semantic models of CSP, for example, traces-divergences or failures-divergences and how they related to data refinement in a relational model.

The variations in the relational model include the interpretation of an operation given as a partial relation, and the observations made. Two possible interpretations are usually articulated for a partial operation: non-blocking and blocking. The former denotes a contract approach - outside a precondition anything may happen - the latter a behavioural approach - outside a precondition (guard) nothing may happen. The observations made in a relational model are usually restricted to the input/output of the ADT, however, these can be extended to include, for example, the refusals in a given state. We thus gain results such as:

In non-blocking relational model with standard observations, $A \sqsubseteq_{data} C$ if and only if $\llbracket A \rrbracket$ is traces-divergences refined by $\llbracket C \rrbracket$.

This particular result is due to Schneider [40]. A version in a blocking model is due to Bolton and Davies, where the process semantics induced is a singleton-failures model. Derrick and Boiten show [19] what additional observations are needed to induce failures-divergences refinement. The latter derive simulation rules for failures-divergences refinement in the blocking model, and recently [20] included the integration of internal operations into this model. The construction relies on a proof that the standard relational semantics is an abstraction of the concurrency semantics, that is, the former does not distinguish between processes that are considered equal by the latter.

Here we generalise these results by providing a general flexible scheme for incorporating the two main “erroneous” concurrent behaviours: deadlock and divergence, into relational refinement. This is shown to subsume previous results. We are also able to derive downward and upward simulation conditions for specifications containing both internal operations and outputs that correspond to failures-divergences refinement. We highlight the role of divergence due to unbounded internal evolution and the relation between outputs and refusals, and their effect on the simulation conditions.

The paper is structured as follows. In Section 2 we provide the basic definitions and background, and we survey existing work in Section 3. The new relational datatype is defined in Section 4, and refinement rules for it are derived through its embedding in the standard theory. Sections 5 and 6 instantiate this new data type with a number of varieties of data types with internal operations and outputs. We conclude in Section 7 which also discusses the formalisation of the theory in KIV.

2. Background

In this section we present the standard refinement theory [18] for abstract data types in a relational setting. We describe the relational model for abstract data types written as total relations over a global state in Section 2.1 and that for partial relations in Section 2.2. Section 2.3 introduces process algebraic variants of refinement that are relevant to our discussion. In Section 2.4 shows how the relational theory can be applied to a specification language such as Z, and Section 2.5 studies the global state in more detail.

2.1. The standard relational model

Relational abstract data types are centred around a hidden local state **State**. However, their semantics are defined in terms of observations on a visible “global” state **G**. These observations are induced by *programs* which are characterised by (sequences of) invocations of the ADT’s operations. The *initialisation* of the program takes a global state to a local state, on which the operations act, a *finalisation* translates back from local to global. The semantics of a program is then a relation on the global state: an initialisation, followed by operations on the local state, followed by a finalisation.

In order to distinguish between relational formulations (which use Z as a meta-language) and expressions

in terms of Z schemas etc., we introduce the convention that expressions and identifiers in the world of relational data types are typeset in a sans serif font.

Definition 1 (Basic data type; Total data type; Program controlled ADT).

A *basic data type* is a quadruple $(\text{State}, \text{Init}, \{\text{Op}_i\}_{i \in I}, \text{Fin})$. The operations $\{\text{Op}_i\}$, indexed by $i \in I$, are relations on the set State ; Init is a total relation from G to State ; Fin is a total relation from State to G . If all operations $\{\text{Op}_i\}$ are total relations, we call it a *total data type*.

A basic ADT with initialisation Init and global state G is *program controlled* if the initial global state is irrelevant for initialisation, i.e., $\text{Init} = G \times \text{ran Init}$. \square

All ADTs in this paper will be program controlled, that is, the global state before initialisation is irrelevant. This ensures that the output of an ADT run is fully determined by the program, not by any other global state information. This does not mean that initialisation is irrelevant: it still describes in which local state the ADT might start, but this choice is not influenced by outside information.

Definition 2 (Program; Data refinement).

A *program* over a data type $D = (\text{State}, \text{Init}, \{\text{Op}_i\}_{i \in I}, \text{Fin})$ is a sequence over the index set I , which is identified with the sequential composition of the corresponding operations. For a program p , the corresponding *complete program* for p in D , denoted p_D , is the relational composition $\text{Init} \circ p \circ \text{Fin}$. For example, if $p = \langle p_1, \dots, p_n \rangle$ then $p_D = \text{Init} \circ \text{Op}_{p_1} \circ \dots \circ \text{Op}_{p_n} \circ \text{Fin}$.

For total data types A and C , C *refines* A , denoted $A \sqsubseteq_{\text{data}} C$, iff for each finite sequence p over I , we have $p_C \subseteq p_A$. \square

As usual we assume that the data types are *conformal*, i.e., they use the same index set for the operations. Then *downward* and *upward* simulations form a sound and jointly complete proof method for verifying refinements [25, 16]. In a simulation a step-by-step comparison is made of each operation in the data types, and to do so the concrete and abstract states are related by a retrieve relation.

Definition 3 (Downward simulation).

Assume total data types $A = (\text{AState}, \text{AInit}, \{\text{AOp}_i\}_{i \in I}, \text{AFin})$ and $C = (\text{CState}, \text{CInit}, \{\text{COp}_i\}_{i \in I}, \text{CFin})$. A *downward simulation* is a relation R between AState and CState satisfying

$$\begin{aligned} \text{CInit} &\subseteq \text{AInit} \circ R \\ R \circ \text{CFin} &\subseteq \text{AFin} \\ \forall i : I \bullet R \circ \text{COp}_i &\subseteq \text{AOp}_i \circ R \end{aligned}$$

If such a simulation exists, we also say that C is a downward simulation of A and similarly for the corresponding operations of A and C . \square

Definition 4 (Upward simulation).

For total data types A and C as above, an *upward simulation* is a relation T between CState and AState such that

$$\begin{aligned} \text{CInit} \circ T &\subseteq \text{AInit} \\ \text{CFin} &\subseteq T \circ \text{AFin} \\ \forall i : I \bullet \text{COp}_i \circ T &\subseteq T \circ \text{AOp}_i \end{aligned}$$

\square

2.2. Partial Relations

Definition 2 and the simulations defined above provide a data refinement methodology for *total* ADTs only. However, in general, operations (for example in Z) may be *partial* relations. The domain of an operation is

the collection of before-states where it is required to deliver a well-defined result; from states outside the domain, the operation may either be *forbidden* or its result *unprescribed*. In the former case, we call this the *blocking* approach, and the domain acts as a *guard*; in the *non-blocking* approach, it acts as a *pre-condition*.

Consequently, there are two possibilities for a refinement theory for partial relations deriving from the total relations theory described above. They require an embedding of partial relations into total relations, so-called “totalisations”. This is done by adding a \perp value to the state space to represent “erroneous” behaviour, where we let $\text{State}_\perp = \text{State} \cup \{\perp\}$ for some $\perp \notin \text{State}$. One way of defining a totalisation is the following definition¹.

Definition 5 (Totalisation).

For a partial relation Op on State , its totalisation is a total relation on State_\perp , defined in the non-blocking model by

$$\widehat{\text{Op}}^{\text{nb}} == \text{Op} \cup \{x, y : \text{State}_\perp \mid x \notin \text{dom Op} \bullet (x, y)\}$$

or in the blocking model by

$$\widehat{\text{Op}}^{\text{b}} == \text{Op} \cup \{x : \text{State}_\perp \mid x \notin \text{dom Op} \bullet (x, \perp)\}$$

□

The simulation rules for partial operations are derived by applying the simulation rules to the totalised relations, and then eliminating \perp from the resulting conditions. For the detailed derivations and the resulting simulation rules on partial relations, see [18] – we will derive a single scheme generalising both the blocking and non-blocking model below.

2.3. Process refinement

A contrasting view of refinement is that offered by a process algebraic description of a system. There, instead of a relation over a global state being representative of a program, the traces of events (in essence, a record of all terminating programs) are recorded.

Here we will consider semantic models of CSP, of which there are a number and each of which induces its own refinement relation. These refinement relations are closely related to those in other process algebras. However, the CSP models take a particular approach with respect to divergence and it should be noted that other process algebraic models sometimes differ in this respect, see for example, [31] and [44] and the comparison given in [12] between CSP, CCS and LOTOS. We will assume this approach to divergence in our treatment of both process refinement and relational refinement, and this has particular consequences when we look at inclusion of internal events in our framework.

Failures-divergences semantics The standard semantics of CSP is the failures-divergences semantics developed in [13, 14, 36]. A process is modelled by the triple $(A, \mathcal{F}, \mathcal{D})$ where A is its alphabet, \mathcal{F} is its *failures* and \mathcal{D} is its *divergences*. The failures of a process are pairs (t, X) where t is a finite sequence of events that the process may undergo and X is a set of events the process may refuse to perform after undergoing t . That is, if the process after undergoing t is in an environment which only allows it to undergo events in X , it may deadlock. The divergences of a process are the sequences of events after which the process may undergo an infinite sequence of internal events, i.e. livelock. Unguarded recursion also leads to divergences.

Failures and divergences are defined in terms of the events in the alphabet of the process. The failures of a process with alphabet A are a set

$$\mathcal{F} \subseteq A^* \times \mathbb{P} A$$

¹ There are others, which are discussed in [21] and in terms of the terminology defined there the following totalisation is an *unstrict lifting*, Deutsch and Henson provide a careful examination of the possible ways to totalise and embed the element \perp into a total data type.

such that a number of properties hold. Properties $F1$ and $F2$ capture the requirement that the sequences of events a process can undergo form a non-empty, prefix-closed set. Property $F3$ states that if a process can refuse all events in a set X then it can refuse all events in any subset of X . Property $F4$ states that a process can refuse any event which cannot occur as the next event.

$$(\langle \rangle, \emptyset) \in \mathcal{F} \quad (F1)$$

$$(t_1 \wedge t_2, \emptyset) \in \mathcal{F} \Rightarrow (t_1, \emptyset) \in \mathcal{F} \quad (F2)$$

$$(t, X) \in \mathcal{F} \wedge Y \subseteq X \Rightarrow (t, Y) \in \mathcal{F} \quad (F3)$$

$$(t, X) \in \mathcal{F} \wedge (\forall e \in Y \bullet (t \wedge \langle e \rangle, \emptyset) \notin \mathcal{F}) \Rightarrow (t, X \cup Y) \in \mathcal{F} \quad (F4)$$

The divergences of a process with alphabet A and failures \mathcal{F} are a set $\mathcal{D} \subseteq A^*$ such that:

$$\mathcal{D} \subseteq \text{dom } \mathcal{F} \quad (D1)$$

$$t_1 \in \mathcal{D} \wedge t_2 \in A^* \Rightarrow t_1 \wedge t_2 \in \mathcal{D} \quad (D2)$$

$$t \in \mathcal{D} \wedge X \subseteq A \Rightarrow (t, X) \in \mathcal{F} \quad (D3)$$

The first property states that a divergence is a sequence of events. Properties $D2$ and $D3$ capture the idea that it is impossible to determine anything about a divergent process in a finite time. Therefore, the possibility that it might undergo further events cannot be ruled out. In other words, a divergent process behaves *chaotically*².

The failures-divergences semantics induces a refinement ordering defined in terms of failures and divergences [14]. A process Q is a refinement of a process P , denoted $P \sqsubseteq_{fd} Q$, iff

$$\mathcal{F}(Q) \subseteq \mathcal{F}(P) \text{ and } \mathcal{D}(Q) \subseteq \mathcal{D}(P)$$

There are two other semantics models for CSP relevant to this paper.

Traces-divergences semantics The traces-divergences semantics is just the failures-divergences semantics with the refusal information removed. A process P is now modelled by $(A, \mathcal{T}, \mathcal{D})$, where \mathcal{T} are the traces of P . It is obtained from the failures-divergences semantics by defining the traces as $\mathcal{T}(P) == \{tr \mid (tr, \emptyset) \in \mathcal{F}\}$.

The traces-divergences semantics induces a refinement ordering, where $P \sqsubseteq_{td} Q$ iff

$$\mathcal{T}(Q) \subseteq \mathcal{T}(P) \text{ and } \mathcal{D}(Q) \subseteq \mathcal{D}(P)$$

Singleton failures semantics The singleton failures semantics for CSP was used by Bolton [6] (and published in [7, 8]) in order to define an appropriate correspondence with blocking data refinement. Essentially the singleton failures semantics is a failures semantics where the refusal sets have cardinality at most one. Specifically, a process is now modelled by (A, \mathcal{S}) where $\mathcal{S} \subseteq A^* \times \mathbb{P}_1 A$ (and \mathbb{P}_1 forms subsets of cardinality at most one).

If P is a process expressed in terms of $stop$, \rightarrow , \sqcap , \square and \parallel , then its singleton failures are given as the obvious projection from its failures, that is:

$$\mathcal{S}(P) = \mathcal{F}(P) \cap (A^* \times \mathbb{P}_1 A)$$

The singleton failures semantics induces a refinement ordering, where $P \sqsubseteq_{sf} Q$ iff

$$\mathcal{S}(Q) \subseteq \mathcal{S}(P)$$

Clearly, failures-divergences refinement is stronger than traces-divergences refinement, that is, $P \sqsubseteq_{fd} Q \Rightarrow P \sqsubseteq_{td} Q$. For divergent-free processes we have $P \sqsubseteq_{sf} Q \Rightarrow P \sqsubseteq_{td} Q$, and for divergent-free basic processes

² The assumptions made here by the CSP models are not necessarily present in other process algebraic semantics.

(i.e., ones expressed in terms of $stop$, \rightarrow , \sqcap , \sqcup and \parallel) we have $P \sqsubseteq_{fd} Q \Rightarrow P \sqsubseteq_{sf} Q$. The relationship of singleton failures to other semantic models is discussed in [45] and later in [11] (which builds on [10]).

2.4. Refinement in Z

In Section 2.2 we described how partial relations are totalised in order that the relational theory of data refinement could be defined for a language such as Z. To complete the picture we provide a relational interpretation for Z specifications, that is, provide a semantics in terms of partial relations.

A Z specification can be thought of as a data type, defined as a tuple $(State, Init, \{Op_i\}_{i \in I})$. The operations Op_i are given in terms of (the variables of) $State$ (its before-state) and $State'$ (its after-state). The initialisation is expressed in terms of an after-state $State'$. In addition, operations consume inputs and produce outputs. The standard solution [46, 18] for embedding a Z specification into a relational model is as follows.

State, Initialisation and Finalisation The inputs and outputs are observable, so they are added to the global state, and because we require that every operation is embedded into a homogeneous relation, the local state space contains a representation of the Z state together with the sequence of inputs and outputs:

$$\begin{aligned} G &== \text{seq } Input \times \text{seq } Output \\ State &== \text{seq } Input \times \text{seq } Output \times State \end{aligned}$$

The initialisation transfers the sequence of inputs from the global state to the local state, and picks an initial local ADT state that satisfies the ADT's initialisation, and the finalisation makes visible the outputs produced by the program.

$$\begin{aligned} Init &== \{Init; is : \text{seq } Input; os : \text{seq } Output \bullet (is, os) \mapsto (is, \langle \rangle, \theta State')\} \\ Fin &== \{State; is : \text{seq } Input; os : \text{seq } Output \bullet (is, os, \theta State) \mapsto (\langle \rangle, os)\} \end{aligned}$$

Operations An operation Op is modelled as a relation which consumes inputs and produces outputs:

$$Op_i == \{Op_i; is : \text{seq } Input; os : \text{seq } Output \bullet (\langle \theta Input \rangle \frown is, os, \theta State) \mapsto (is, os \frown \langle \theta Output \rangle, \theta State')\}$$

Retrieve relations If R is a retrieve relation between $AState$ and $CState$, this is modelled as

$$R == \{R; is : \text{seq } Input; os : \text{seq } Output \bullet (is, os, \theta AState) \mapsto (is, os, \theta CState)\}$$

In a context where there is no input or output, the global state contains no information and is a one point domain, i.e., $G == \{*\}$, and the local state is $State == State$. In such a context the other components of the embedding collapse to the following:

$$\begin{aligned} Init &== \{Init \bullet * \mapsto \theta State'\} \\ Op &== \{Op \bullet \theta State \mapsto \theta State'\} \\ Fin &== \{(\theta State, *)\} \\ R &== \{R \bullet \theta AState \mapsto \theta CState\} \end{aligned}$$

Given these embeddings, we can translate the relational refinement conditions of downward and upward simulations into simulations conditions for Z ADTs. This is straightforward, the only point of note is that the conditions on the finalisation are always satisfied in this Z interpretation.

2.5. Observations: Inputs, Outputs and Refusals

A critical decision in developing a refinement theory is: what observations can be made on an ADT? The set of valid observations determines the global state – or, in principle, the aspect that is observed at finalisation. Two types of observations are of particular relevance: outputs and refusals, and we discuss these now.

2.5.1. Inputs and outputs

The standard construction described above embedded a sequences of inputs and outputs into the local and global states. Here, however, without any loss of generality, we can do away with the input sequence since any quantification over inputs comes hand in hand with the same quantification over the index set I . That is, it would make no difference for inputs in operations to appear as a shorthand for an increased alphabet of the ADT³ – as is commonly the interpretation of inputs in process algebras. Thus, we avoid mentioning inputs from this point. In Z versions of the rules, wherever it says $\forall i : I$ we may mentally insert $\forall \text{Input}$ as well. We do not rely on the conventional assumption that I is finite, and so are not requiring inputs to be from a finite domain either.

The outputs, however, do need an explicit representation in the relational embedding. The results of embeddings where the outputs are part of the global state, and initialisation and finalisation perform copying of them between global and local state are characterised by the following definition. The final two conditions state that every operation adds a single output to the output sequence and copies the rest, and that the previous outputs have no influence on the effect of this operation in general, including the new output value.

Definition 6 (Output embedding).

A basic ADT $(\text{State}, \text{Init}, \{\text{Op}_i\}_{i \in I}, \text{Fin})$ with global state G is said to be an *output embedding* iff types $GB, \text{Output}, \text{StateB}$ exist such that

$$\begin{aligned}
G &= GB \times \text{seq } \text{Output} \\
\text{State} &= \text{StateB} \times \text{seq } \text{Output} \\
\forall i, gs, ls, ls2, os, os2, \text{out} \bullet \\
&((gs, os), (ls, os2)) \in \text{Init} \Rightarrow os2 = \langle \rangle \\
&((ls, os), (gs, os2)) \in \text{Fin} \Rightarrow os2 = os \\
&((ls, os), (ls2, os2)) \in \text{Op}_i \Rightarrow \exists \text{out} \bullet os2 = os \hat{\ } \langle \text{out} \rangle \\
&((ls, os), (ls2, os \hat{\ } \langle \text{out} \rangle)) \in \text{Op}_i \Rightarrow ((ls, os2), (ls2, os2 \hat{\ } \langle \text{out} \rangle)) \in \text{Op}_i
\end{aligned}
\quad \square$$

2.5.2. Refusals

We wish to observe traces and refusals⁴. The traces are contained in the standard relational semantics, viz. through the notion of programs. The refusals are not present in the standard embedding given above. Thus to fully encode failures it is necessary to enhance the standard relational theory by adding the observation of refusals at the end of every program. This involves a change only to the global state and the finalisation, as embodied in the following definition.

Definition 7 (Refusal embedding).

Consider a basic ADT $(\text{State}, \text{Init}, \{\text{Op}_i\}_{i \in I}, \text{Fin})$ with global state G , a type of events E , and a refusal relation $\text{Ref} : \text{State} \leftrightarrow \mathbb{P}E$. The ADT is said to be an *refusal embedding* (for Ref) iff it is program controlled, and a type GB exists such that

$$\begin{aligned}
G &= \mathbb{P}E \times GB \\
\forall ls, r \bullet (\exists gs \bullet (ls, (r, gs)) \in \text{Fin}) \equiv (ls, r) \in \text{Ref}
\end{aligned}
\quad \square$$

As an example, in the context of a Z data type with no input and output, the refusals will be any subset E of the maximal refusals in a given state: $\{i : I \mid \neg \text{pre } \text{Op}_i\}$, where I is the set of operation names. Thus an

³ The standard embedding, taking into account ignored outputs at initialisation and ignored inputs at finalisation, effectively leads to a relation of type $\text{seq } \text{Input} \leftrightarrow \text{seq } \text{Output}$ for every program, i.e. a semantics of type $\text{seq } \text{Index} \rightarrow (\text{seq } \text{Input} \leftrightarrow \text{seq } \text{Output})$ overall. By not embedding inputs, we are replacing this by $\text{seq}(\text{Index} \times \text{Input}) \rightarrow \{*\} \leftrightarrow \text{seq } \text{Output}$, in functional programming terms: uncurrying and zipping.

⁴ For discussions about which aspects of a system should or could be observed see [45] and [23].

abstract data type $(State, Init, \{Op_i\}_{i \in I})$ in the refusals interpretation is embedded in the relational model as follows. The global state \mathbf{G} is $\mathbb{P}I$, finalisation is given by

$$Fin == \{State; E : \mathbb{P}I \mid (\forall i \in E \bullet \neg \text{pre } Op_i) \bullet \theta State \mapsto E\}$$

and initialisation is given by

$$Init == \{Init; E : \mathbb{P}I \bullet E \mapsto \theta State'\}$$

The local state and the embedding of operations are unchanged, see Section 2.4.

Note that a basic ADT can be both an output embedding and a refusal embedding, for example having as the global state $\mathbb{P}E \times \text{seq } Output$. Indeed, the presence of outputs has a crucial interaction with the refusals as we discuss now (see also [19]).

In particular, when a system includes a non-deterministic choice of different output values, is the environment able to “select” its choice among these, or is the choice entirely inside the system? The refusals in each model differ slightly, and only the latter option causes refusals through choice of output. These options are called the *angelic* and *demonic* models of outputs, respectively (after [42]⁵). In the **Angelic** model the only refusals are the ones arising when an operation is not applicable, there are no refusals due to outputs. In the **Demonic** model [26] a process is, in addition, allowed to refuse all but one of the possible outputs. The difference between the two models is summarised in the following definition.

Definition 8 (Angelic and Demonic Embeddings).

An abstract data type $(State, Init, \{Op_i\}_{i \in I})$ with outputs of type *Output* is embedded in the relational model as follows. The global state is defined by

$$\begin{aligned} Event &== I \times Output \\ \mathbf{G} &== \text{seq } Output \times \mathbb{P}Event \end{aligned}$$

The embedding of operations is as in Section 2.4; initialisation is extended with an input set of events that is ignored. Finalisation is given by

$$Fin == \{State; os : \text{seq } Output; E : \mathbb{P}Event \mid Fcond \bullet (os, \theta State) \mapsto (os, E)\}$$

where *Fcond* is in the angelic model:

$$E \subseteq \{(i, out) \mid \neg \exists State'; Output \bullet Op_i \wedge \theta Output = out\}$$

and in the demonic model:

$$E \subseteq \{(i, out) \mid (\neg \exists State'; Output \bullet Op_i \wedge \theta Output = out) \vee (\exists State'; Output \bullet Op_i \wedge \theta Output \neq out \wedge (i, \theta Output) \notin E)\}$$

□

We will use this embedding when defining simulation rules that correspond to failures-divergences refinement below.

3. Relating data and process refinement

The previous section outlined the standard relational theory of refinement and relevant process based definitions. We now survey recent existing work relating relational refinement with process refinement. The correspondences are summarised in the following table.

⁵ A more natural naming would be “external” vs. “internal”, in analogy with CSP external and internal choice. However, in this paper the term “internal” is already used in a different context.

Relational refinement	Process model	Citations
Non-blocking data refinement	traces-divergences	Schneider [40]
Blocking data refinement with deterministic outputs	singleton failures	Bolton and Davies [7, 8]
Blocking data refinement	singleton failures of process and input process	Bolton and Davies [7, 8]
Blocking data refinement with strengthened applicability but no input/output	failures	Josephs [30]
Blocking data refinement with extended finalisations	failures-divergences	Derrick and Boiten [4, 19]
Non-blocking data refinement with extended finalisations but no input/output	failures-divergences	Derrick and Boiten [4, 19]

Both Bolton and Davies and Schneider consider the standard definition of data refinement. The 'extended finalisations' of Derrick and Boiten add conditions to the standard simulation rules, and these are detailed in Section 3.3 below.

3.1. Non-blocking data refinement and the traces-divergences semantics

Inspired by the work by Bolton and Davies [7, 8] discussed below, Schneider [40] shows that non-blocking data refinement corresponds to traces-divergences refinement in a process semantics. To show this he translates ADTs into CSP directly, and uses the traces-divergences semantics on the resulting CSP process.

As with all approaches the result is first proved for ADTs without input and output, and then extended to the general case. The extension to ADTs with inputs and outputs (which Schneider calls communicating data types following Bolton) involves the embedding of input and output sequences in the global states as defined above. The notation used is slightly different, but the construction is isomorphic to that given in Section 2.4 above. For such an ADT the translation of an ADT A into a CSP process $process(A)$ is given by

$$process(A) == \sqcap s \in State, (*, s) \in Init \bullet Proc_A(s)$$

$$Proc_A(s) ==$$

$$\begin{aligned} & \square i \in I, in \in Input, (\langle in \rangle, \langle \rangle, s) \in \text{dom } AOp_i \bullet \\ & \quad \sqcap s' \in State, out \in Output, (\langle in \rangle, \langle \rangle, s) \mapsto (\langle \rangle, \langle out \rangle, s') \in AOp_i \bullet AOp_i.in.out \rightarrow Proc_A(s') \\ & \square \\ & \square i \in I, in \in Input, (\langle in \rangle, \langle \rangle, s) \notin \text{dom } AOp_i \bullet \sqcap out \in Output \bullet AOp_i.in.out \rightarrow \text{div} \end{aligned}$$

Note that here div is the divergent CSP process, which ensures that all events are possible after an operation has been called outside its precondition. The following (in the notation used in this paper) is then proved.

Theorem 1. In the non-blocking model, $A \sqsubseteq_{data} C$ if and only if $process(A) \sqsubseteq_{td} process(C)$. \square

3.2. Blocking data refinement and the singleton failures semantics

Bolton in [6] and Bolton and Davies in [7, 8] discuss the relationship between data refinement and the singleton failures semantics [45] model. They consider both the blocking and non-blocking relational data type semantics, and, like Schneider, translate ADTs directly into CSP.

For the blocking model, the translation of an ADT A into a CSP process $process_b(A)$ is given by the following (using, for uniformity, the notation already introduced):

$$process_b(A) == \sqcap s \in State, (*, s) \in Init \bullet P_A(s)$$

$$\begin{aligned}
P_A(s) == & \\
& \square i \in I, in \in Input, (\langle in \rangle, \langle \rangle, s) \in \text{dom } AOp_i \bullet \\
& \quad \square s' \in State, out \in Output, (\langle in \rangle, \langle \rangle, s) \mapsto (\langle \rangle, \langle out \rangle, s') \in AOp_i \bullet AOp_i.in.out \rightarrow P_A(s')
\end{aligned}$$

As can be seen the enabling of this process is identical to that of Schneider, however, the effect of calling an operation outside its precondition is not now divergence but, since we are in the blocking model, simply inability to perform any event associated with that operation. This thus correctly reflects the intended meaning to the blocking model and with it data refinement corresponds to singleton failures refinement in the process model.

The inclusion of non-deterministic outputs complicates the process semantics needed, and an additional constraint is needed in order to characterise blocking data refinement. To do this a further partial translation is introduced, called *inputProcess* which provides a characterisation of when a particular input is in the domain. For the blocking model this is defined as:

$$\begin{aligned}
inputProcess_b(A) == & \square s \in State, (*, s) \in Init \bullet P_A(s) \\
P_A(s) == & \square i \in I, in \in Input, (\langle in \rangle, \langle \rangle, s) \in \text{dom } AOp_i \bullet \\
& \quad \square s' \in State \mid (\exists out \in Output \mid (\langle in \rangle, \langle \rangle, s) \mapsto (\langle \rangle, \langle out \rangle, s') \in AOp_i) \bullet AOp_i.in \rightarrow P_A(s')
\end{aligned}$$

As can be seen, this is the same as *process_b* except that the outputs are unobservable. [8] contains the following result: Blocking data refinement is equivalent to singleton failures refinement of both the process and the input process. That is:

Theorem 2. In the blocking model, $A \sqsubseteq_{data} C$ if and only if $process_b(A) \sqsubseteq_{sf} process_b(C)$ and $inputProcess_b(A) \sqsubseteq_{sf} inputProcess_b(C)$. For ADTs with deterministic outputs (or no input/output), this reduces to checking $process_b(A) \sqsubseteq_{sf} process_b(C)$. \square

Two corollaries are worth noting. First, that this characterisation is equivalent to checking the singleton failures of the input process and trace refinement of the process. Second, that

$$process_b(A) \sqsubseteq_{fd} process_b(C) \Rightarrow A \sqsubseteq_{data} C$$

in the blocking model.

Bolton and Davies also consider the non-blocking model. However, the corresponding process used is different to that of Schneider. Specifically, they define *process_{nb}* by

$$\begin{aligned}
process_{nb}(A) == & \square s \in State, (*, s) \in Init \bullet Q_A(s) \\
Q_A(s) == & \begin{array}{l} P_A(s) \\ \square \\ ((\square i \in I, in \in Input, (\langle in \rangle, \langle \rangle, s) \notin \text{dom } AOp_i \bullet \\ \quad \square s' \in State, out \in Output \bullet AOp_i.in.out \rightarrow Chaos) \\ \square stop) \end{array}
\end{aligned}$$

where *Chaos* == $(\square i \in I, in \in Input, out \in Output \bullet AOp_i.in.out \rightarrow Chaos) \square stop$ is the non-divergent process that can perform any event, yet also refuse any event. In the process Q_A the lower *stop* is being used to represent non-termination of the operation. With this corresponding process analogous results to the above are derived (i.e., data refinement corresponding to singleton as opposed to failures-divergences refinement). However, this non-blocking translation differs in key aspects from that defined by Schneider, and in particular, the use of *Chaos* and *stop* to model non-termination seems a less natural embedding of non-blocking than using explicit divergence.

However, Reeves and Streader have recently shown [34] that the equivalence given in Theorem 2 only holds if the relational semantics observes the exact point of deadlock, e.g., by producing trivial outputs, and preserving and counting those after deadlock.




Fig. 1. Singleton failures vs. data refinement [34].

Specifically, they discuss the difference between the embedded state $(State \times \text{seq } Output)_{\perp}$ that we use (following [46]) and $(State_{\perp} \times \text{seq } Output)$ which is used by Bolton and Davies [8]. By using the latter, the relational world makes finer distinctions: it is possible to observe how many outputs were produced before the state devolved to \perp . This allows to distinguish between, say, a trace abc having blocked on b or on c . The standard relational model does not allow this – in that model we can observe the outputs produced by any prefix of abc that does not yet block, but not the additional information of whether blocking immediately after that was possible, if this prefix already allowed blocking. They give the example depicted in Figure 1. The processes A and C in the figure are equivalent in a (standard) relational model with no outputs (in either the blocking or non-blocking model). However, C is not a singleton failures refinement of A , adding the failure $(ab, \{c\})$. For relational refinement to correspond to singleton failure refinement, it needs to include a way of observing the exact point of deadlock, for example, by using $(State_{\perp} \times \text{seq } Output)$ as the appropriate embedding.

3.3. Defining a correspondence with failures-divergences refinement

Subsequent to the work of Bolton and Davies, Derrick and Boiten in [4, 19] explored what additional conditions were needed on data refinement in order to achieve a correspondence with failures-divergences refinement in a process semantics. To define the appropriate correspondence Derrick and Boiten define a process semantics directly. With no inputs or outputs the traces, failures and divergences can be defined easily in each of the two models.

Blocking model Traces arise from sequences of operations which are defined within their guards. Refusals indicate the impossibility of applying an operation outside its precondition. Furthermore, there are no divergences since each operation is either blocked or gives a well-defined result.

Non-blocking model As no operation is blocked, every trace is possible: those that arise in the blocking model, and any other ones following divergence. There are no refusals beyond those after a divergence, since before the ADT diverges, no operation is blocked, it either gives a well-defined result or causes divergence. There are now, however, divergences, which arise from applying an operation outside its precondition.

Since refusals are not normally observed in data refinement, it is necessary to observe the refusals directly, and a refusal embedding (see Definition 7) is used. Thus in a context where the data type has no input or output the finalisation used is generalised from being $\{State \bullet \theta State \mapsto *\}$ to becoming $\{State \bullet \theta State \mapsto E\}$.

In defining the process semantics when inputs and outputs are included the only change is to the refusals since the traces and divergences remain the same. When extended to data types with outputs, the effect of adding outputs has consequences for the process semantics, and in particular the refusals of an ADT. Therefore, as detailed in Section 2.5.2, angelic and demonic refusal embeddings as given in Definition 8 are used for the finalisation. The following is proved (irrespective of output model chosen):

Theorem 3. In the non-blocking model, for ADTs with no input/output data refinement with extended finalisations corresponds to failures-divergences refinement.

In the blocking model, for ADTs with or without input/output, data refinement with extended finalisations corresponds to failures-divergences refinement. \square

Josephs' construction For the blocking model without input or output, Josephs in [30] showed that downward and upward simulations with the strengthened applicability condition (1) form a sound and jointly complete method for verifying failures refinement in CSP. Working entirely in the CSP semantics he uses the blocking model in the sense that his CSP processes are projections of ADTs under a semantics where there are no divergences and refusals arise outside an operation's precondition. When the link is made to the relational framework via a corresponding process, the results in [30] are consistent with those in [19].

Simulation conditions The downward and upward simulation rules as expressed for partial relations contain conditions on the finalisations. When we use a non-standard finalisation these conditions potentially impose additional constraints on the simulation conditions.

First in a model without input or output. Here, even with a refusal embedding a downward simulation places no further constraints than already present in the standard definition. For an upward simulation the finalisation conditions are:

$$\begin{aligned} \text{CFin} &\subseteq \text{T} \text{ ; } \text{AFin} \\ \forall c : \text{CState} \bullet \text{T}(\{c\}) &\subseteq \text{dom AFin} \Rightarrow c \in \text{dom CFin} \end{aligned}$$

With the refusals embedding the second is always satisfied, however, the first leads to a strengthening of the standard applicability condition from $\forall i : I \bullet \forall \text{CState} \bullet \exists \text{AState} \bullet T \wedge (\text{pre } AOp_i \Rightarrow \text{pre } COp_i)$ to

$$\forall \text{CState} \bullet \exists \text{AState} \bullet \forall i : I \bullet T \wedge (\text{pre } AOp_i \Rightarrow \text{pre } COp_i) \quad (1)$$

As noted in [19]: 'The standard upward simulation applicability condition requires that we have to consider pairs of abstract and concrete states for each operation. The finalisation condition, on the other hand, requires that for every abstract state we can find a *single* concrete state such that all the preconditions of the abstract operations imply the preconditions of their concrete counterparts.'

In a model with outputs we distinguish between the demonic and angelic models.

Demonic model of outputs. The consequences lie, as before, with the upward simulation conditions, since the downward simulation condition imposed by a refusal embedding is subsumed by the normal applicability and correctness rules.

In the case of upward simulations the finalisation condition leads to an extra condition, which is somewhat complicated involving the need to look at combinations of different operations, whilst considering possible output values individually. This complexity arises from the presence of non-deterministic outputs, and their interaction with the refusal sets. With only deterministic outputs the upward simulation condition is as above (i.e., just involving a strengthened applicability condition). In a model with inputs and outs, the requirement that $\text{CFin} \subseteq \text{T} \text{ ; } \text{AFin}$, simplifies to:

$$\forall \text{CState}; E \bullet \text{Fcond}_C \Rightarrow \exists \text{AState} \bullet T \wedge \text{Fcond}_A \quad (2)$$

This forces one to consider different linked abstract states for different maximal concrete refusal sets. In particular, even with just a single operation, it is, in general, necessary to look at different linked states for different output values. In fact, the above condition and its representation below, combine this aspect with the condition derived in the basic construction, which insisted on choosing *the same* linked abstract state for every set of enabled operations.

It is easy to prove that it is sufficient to consider only *maximal* refusal sets in each concrete state. This, and the fact that without outputs there is only one maximal refusal set, allowed to consider only a single linked abstract state. However, in the presence of non-deterministic outputs, multiple maximal refusal sets may exist in each concrete state, each of which may be verified by a different abstract state. See Section 5.6.2 for an example of why the condition is necessary. The condition is, of course, stronger than normal applicability, and additionally implies the totality of T .

Angelic model of outputs. The effect of the finalisation on the simulation rules are less drastic in the angelic model. Whereas in the demonic model we can reduce non-determinism in the outputs, in the angelic model one cannot. This difference is easily expressed in the simulation rules by using a different precondition operator, defined as $\text{Pre } Op == \exists \text{State}' \bullet Op$, and one then use Pre in place of pre in the strengthened applicability conditions.

A complete characterisation of the simulations rules can now be given. For downward simulation, we have the following potential collection of conditions:

- DS.Init** $\forall C\text{State}' \bullet C\text{Init} \Rightarrow \exists A\text{State}' \bullet A\text{Init} \wedge R'$
DS.App $\forall C\text{State}; A\text{State}; i : I \bullet R \wedge \text{pre } AOp_i \Rightarrow \text{pre } COp_i$
DS.CorrNonBlock $\forall i : I; \text{Output}; C\text{State}'; C\text{State}; A\text{State} \bullet \text{pre } AOp_i \wedge R \wedge COp_i \Rightarrow \exists A\text{State}' \bullet R' \wedge AOp_i$
DS.CorrBlock $\forall i : I; \text{Output}; C\text{State}'; C\text{State}; A\text{State} \bullet R \wedge COp_i \Rightarrow \exists A\text{State}' \bullet R' \wedge AOp_i$
DS.FinAng $\forall C\text{State}; A\text{State}; i : I; \text{Output} \bullet R \wedge \text{Pre } AOp_i \Rightarrow \text{Pre } COp_i$

Non-blocking data refinement (the standard model and the one corresponding to traces-divergences refinement) requires: **DS.Init**, **DS.App** and **DS.CorrNonBlock**.

Blocking data refinement requires: **DS.Init**, **DS.App** and **DS.CorrBlock**.

For a refusals embedding in the blocking model the following rules are required in the various situations. Each column represents a particular model for (inputs and) outputs; a missing entry indicates a condition dominated by the other conditions in the same column.

Outputs:	none	demonic	angelic
Init	DS.Init		
App	DS.App		-
Corr	DS.CorrBlock		
Fin	-	DS.FinAng	

For upward simulation, we have the totality of T on $C\text{State}$ plus the following set:

- US.Init** $\forall C\text{State}'; A\text{State}' \bullet T' \wedge C\text{Init} \Rightarrow A\text{Init}$
US.AppBlock $\forall i : I; \text{Output} \bullet \forall C\text{State} \bullet \exists A\text{State} \bullet T \wedge \text{pre } AOp_i \Rightarrow \text{pre } COp_i$
US.CorrNonBlock
 $\forall i : I; \text{Output}; A\text{State}'; C\text{State}'; C\text{State} \bullet T' \wedge COp_i \Rightarrow \exists A\text{State} \bullet T \wedge (\text{pre } AOp_i \Rightarrow AOp_i)$
US.CorrBlock $\forall i : I; \text{Output}; A\text{State}'; C\text{State}'; C\text{State} \bullet T' \wedge COp_i \Rightarrow \exists A\text{State} \bullet T \wedge AOp_i$
US.FinRef $\forall C\text{State} \bullet \exists A\text{State} \bullet T \wedge \forall i : I \bullet \text{pre } AOp_i \Rightarrow \text{pre } COp_i$
US.FinDem $\forall C\text{State}; E \bullet F\text{cond}_C \Rightarrow \exists A\text{State} \bullet T \wedge F\text{cond}_A$
US.FinAng $\forall C\text{State} \bullet \exists A\text{State} \bullet T \wedge \forall i : I; \text{Output} \bullet \text{Pre } AOp_i \Rightarrow \text{Pre } COp_i$

Non-blocking data refinement (the standard model and the one corresponding to traces-divergences refinement) requires: **US.Init**, **US.App** and **US.CorrNonBlock**.

Fig. 2. Non-blocking, no input/output = traces-divergences and failures-divergences

Blocking data refinement (corresponding to singleton-failures refinement) requires: **US.Init**, **US.App** and **US.CorrBlock**.

For a refusals embedding in the blocking model the following rules are required in the various situations.

Outputs:	none	demonic	angelic
Init	US.Init		
App	-		
Corr	US.CorrBlock		
Fin	US.FinRef	US.FinDem	US.FinAng

3.4. Discussion

We have seen that in both the non-blocking and blocking models it has been necessary to place additional restrictions (i.e., observations) on the standard definition of data refinement in order that failures-divergences refinement is achieved in a process semantics. Why this is, is perhaps best illustrated via a few examples.

Without input/output - non-blocking. We have seen that without input/output non-blocking data refinement is equivalent to traces-divergences refinement. However, it is worth noting that this does not mean that data refinement suffers from the weakness of the CSP traces model. Specifically, although traces refinement is normally considered too weak since the deadlocked behaviour *stop* refines all processes, such a behaviour is not a feasible translation of an ADT. That is, no ADT will have corresponding process *stop*, since the non-blocking model allows all traces due to no operation being refused. In addition, unlike in trace refinement there is no bottom of the refinement ordering since all ADTs with all operations deterministic and fully defined have no strict refinements in this framework.

Without input/output, non-blocking data refinement is, in fact, also equivalent to failures-divergences refinement. To see this, note that without input/output (specifically without output) the process semantics obtained identifies traces-divergences refinement and failures-divergences refinement, that is, $process(A) \sqsubseteq_{td} process(C)$ iff $process(A) \sqsubseteq_{fd} process(C)$. This is simply because there are no refusals (beyond those after a divergence) in the process semantics, since refusals only arise due to the presence of outputs.

Consider Figure 2, where in this and subsequent examples we define them via simple LTSs which represent the ADT's partial relations before totalisation.

These two specifications have the same traces and divergences, and are thus data refinement equivalent. There are no refusals, thus they are also failures-divergences equivalent. However, note that the stronger applicability condition needed for the blocking model does not hold here - for example for state *ensuite* it is not the case that any abstract state has

$$\forall i : I \bullet T \wedge (\text{pre } AOp_i \Rightarrow \text{pre } COp_i)$$

The difference (i.e., why this does not matter in a traces-divergences model) is that with failures the refusals are tied to the traces, whereas for divergences we simply require their inclusion.

Without input/output - blocking. However, when considered under a blocking totalisation A and C are *not* failures-divergences equivalent. In fact, in a blocking scenario these are singleton failures equivalent and hence a blocking data refinement. To see this note that in a blocking model there are no divergences, the traces are the same in each. Now, although A has failure $(\langle B \rangle, \{TVF, ESF\})$ which is not present in C , under a singleton failures model in A we just obtain singleton failures $(\langle B \rangle, \{TVF\})$, $(\langle B \rangle, \{ESF\})$, ... thus the difference is not observable. To recover failures-divergences refinement in the blocking model one needs

Fig. 3. Non-blocking with input/output = traces-divergences but not failures-divergences

to add the strengthened applicability condition. That is, it is precisely the condition

$$\forall CState \bullet \exists AState \bullet \forall i : I \bullet T \wedge (\text{pre } AOp_i \Rightarrow \text{pre } COp_i)$$

that fails in this example.

With input/output - non-blocking. When we extend to consider input/output, the presence and modelling of outputs forces a different condition (i.e., different, but related, to the strengthened applicability condition **US.FinRef**) required to recover failures-divergences equivalent.

To see why non-blocking data refinement is not failures-divergences refinement consider Figure 3, where different operations above have been replaced by an operation outputting a different value. Again, in the non-blocking model we have the same traces and same divergences in each specification. The presence of outputs, and non-determinism, causes failures in both specifications. For example, *simple* has as its refusals $\mathbb{P}\{HasES!true, HasTV!true\}$, whereas *tv* has as its refusals $\mathbb{P}\{HasES!true, HasTV!false\}$ and *ensuite* has $\mathbb{P}\{HasES!false, HasTV!true\}$.

This difference is not visible in the traces-divergences model, and consequently not in program observations made in the relational model. Moving some of the observable information (operation names here) from the first example into the outputs has kept that information observable but not under control of the environment, and it is this information that is captured in the failures.

We have discussed above the conditions that are required in a blocking model. The non-blocking model was not considered in [19], and Section 6.3 derives the extra condition that needs to be enforced in order to recover failures-divergences refinement. We will find that the condition is, of course, similar to the condition **US.FinDem** needed in the blocking model, but that it does not imply the strengthened applicability condition **US.FinRef**. For reasons we discuss later a condition such as **US.FinRef** is not needed in the non-blocking model.

With input/output - blocking. Considering the blocking model we already know we need an additional strengthening of applicability (even without input/output) to regain failures-divergences. This example shows that we need the condition **US.FinDem** on refusal sets due to outputs as well. Now the strengthened applicability condition holds - each state *simple*, *luxury*, *tv* and *ensuite* has operations *HasES* and *HasTV* enabled - so this does not pick up the different refusal information due to the outputs. Thus we need to impose **US.FinDem** to ensure this.

Note that in the case of the blocking model, **US.FinDem** implies **US.FinRef**. To see this, given a concrete state *CS* and maximal refusal set *E* in that state, then events are in *E* if the associated operation is blocked or there exists an alternate output for that operation. For this *E* we can find an abstract state *AS* such that $Fcond_A$ holds for *E*. Now for this *AS* consider any operation *Op*, if $\text{pre } AOp$ holds in *AS* but $\text{pre } COp$ does not hold in *CS*, then we have violated **US.FinDem**. Hence **US.FinRef** holds.

4. A Relational ADT with Divergence and Blocking

In the previous section, we sketched how partial relations would be interpreted in the model of total relation refinement, depending on whether the blocking or non-blocking approach was chosen. However, we have stated previously [32, 3] that the two approaches are not exclusive – and indeed, formalisms such as B [1] have both preconditions and guards. In general, specification formalisms may include both situations that lead to *divergence* and situations that lead to *deadlock*. For example, in our treatment of internal operations in [20] both occur – divergence appears through livelock, and the blocking interpretation is used for operations that are not enabled. The solution chosen there is to apply two totalisations in sequence: one to account for deadlock, and another to account for livelock. What follows here is a generalised reconstruction of that

solution, which allows us to derive refinement rules directly for any relational formalism which gives rise to both kinds of errors once the areas of divergence and blocking have been made explicit. In fact, it covers *any* relational formalism modelling at most two kinds of errors, one of which is chaotic, and whose combination satisfies the constraints discussed next.

Indeed, an important consideration is the relative ordering of the two kinds of erroneous behaviours. In particular, we need to decide what observations should be possible when the semantics leads to a non-deterministic choice⁶ between any combination of the three behaviours: “normal”, “divergent” and “blocking”.

First, “divergent” behaviour is usually viewed as “anything might happen”, which means that a choice between divergent and normal behaviour should appear as divergence. Consistent with the CSP chaotic interpretation of divergence (e.g., after divergence any refusal is possible), the choice between divergence and blocking should also result in divergence. All in all, this means that there is no observable difference between possible and certain divergence, and that divergence is a zero of non-deterministic choice. The remaining issue is the choice between normal and blocking behaviour. It would be possible, using a model of partial relations (see earlier discussion, and also [16, chapters 8-9]) to take deadlock as a unit of choice, and therefore to not observe possible blocking. Consistent with usual semantics for Z and for CSP, we will distinguish possible blocking in our model. These decisions are summarised in the following table, which also informally hints at the set-based model for this which we will introduce formally later. In particular, \perp represents blocking, and ω represents divergence. In the model, set union acts as the choice operator.

Choice	normal	divergence	deadlock	poss deadlock	model
normal	normal	divergence	poss deadlock	poss deadlock	sets not containing \perp, ω
divergence	divergence	divergence	divergence	divergence	$\text{State} \cup \{\omega\}$, possibly also \perp
deadlock	poss deadlock	divergence	deadlock	poss deadlock	$\{\perp\}$
poss deadlock	poss deadlock	divergence	poss deadlock	poss deadlock	sets containing \perp but not ω

With these considerations in mind, we now define a relational data type with partial operations allowing for both divergence and blocking. Its reduction is the basic data type obtained by removing all blocking and divergence information.

Definition 9 (Process data type; Reduction).

A *process data type* is a quadruple

$$(\text{State}, \text{Inits}, \{\text{Op}_i\}_{i \in I}, \text{Fin})$$

where Inits is a subset of State ; every operation $\{\text{Op}_i\}$ is a triple (N, B, D) such that $\text{dom } N, D$ and B form a partition of State ; Fin is a relation from State to G .

Its *reduction* is the basic data type $(\text{State}, G \times \text{Inits}, \{N_i\}_{i \in I}, \text{Fin})$ □

In an operation $\text{Op} = (N, B, D)$ the relation N represents the operation’s normal effect; the sets B and D represent states where the operation would lead to blocking and divergence, respectively. The three sets forming a partition excludes certain situations, such as miracles and possible (as opposed to certain) deadlock from a given state, and ensures that it can be represented by a *total* data type. Possible deadlock still occurs, however, whenever a program leads to multiple states, some but not all of which are deadlocked.

The blocking and non-blocking approaches to operations are, of course, special cases of process data types, in particular:

- the blocking operation Op is represented by $(\text{Op}, \overline{\text{dom Op}}, \emptyset)$, i.e., it never diverges, and blocks in the complement of the operation’s domain;

⁶ We do not explicitly identify choice operators in our formalism, although availability of multiple operations in a single state is close to external choice, and non-determinism in after-states and choice of outputs in the demonic model are closely related to internal choice.




Fig. 4. The original Op , and a divergent after-state; with $\mathbf{B}_\perp \times \{\perp\}$ added; finally also with $\mathbf{D}_\omega \times \text{State}_\omega \cup \{(\omega, \perp)\}$.

- the non-blocking operation Op is represented by $(\text{Op}, \emptyset, \overline{\text{dom Op}})$, i.e., it diverges in the complement of the operation's domain, but never blocks.

The definition of a process data type is not intended as a new and wonderful relational specification mechanism, to compete with similar approaches such as [2, 22, 24, 28]. It is solely an intermediate formalism that most partial relation frameworks can be embedded into. In turn, it is embedded into the total relations framework in order to define its simulation rules once and for all. The embedding is given below. In effect, it fixes the semantics of a process data type.

In the embeddings we will use state spaces enhanced with special values defined below. For simplicity, we assume in the rest of this paper that \perp (representing blocking), ω (representing divergence), and no are different values not already contained in any local or global state space of interest. The impossibility of making an observation in a final state is encoded in no , and this is added to the global state only. The embedding (and semantics) of a process data type is then defined as follows.

Definition 10 (Enhanced state; Embedding of a process data type).

For any set State , let

$$\text{State}_{\perp, \omega, \text{no}} == \text{State} \cup \{\perp, \omega, \text{no}\}$$

and similarly for sets subscripted with subsets of these special values.

A process data type $(\text{State}, \text{Inits}, \{(N_i, B_i, D_i)\}_{i \in I}, \text{Fin})$ with global state G is embedded into the total data type $(\text{State}_{\perp, \omega}, \text{Init}, \{[\text{Op}]_i\}_{i \in I}, [\text{Fin}])$ where

$$\begin{aligned} \text{Init} &== G_{\perp, \omega, \text{no}} \times \text{Inits} \\ [[N, B, D]] &== N \cup (\mathbf{B}_{\perp, \omega} \times \{\perp\}) \cup (\mathbf{D}_\omega \times \text{State}_\omega) \\ [[\text{Fin}]] &== \text{Fin} \cup (\overline{\text{dom Fin}} \times \{\text{no}\}) \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_{\omega, \text{no}} \end{aligned} \quad \square$$

As a process data type has a set of initial states rather than an initialisation relation, the embedding's initialisation relates every global state to all such states. The normal effect of an operation is part of the embedded one. In addition, every blocking state including \perp is related to \perp , every state where the operation diverges to every state including ω , and divergent state ω is linked to all states even including \perp . Finalisation makes both blocking and divergence visible globally. Figure 4 illustrates this operation embedding.

Refinement (i.e., downward and upward simulation) of process data types is derived by embedding them into total data types⁷, applying the simulation rules for total data types, and then eliminating \perp , ω , and no from the resulting rules – i.e., expressing them in terms of the process data types only.

⁷ The embedding of the initialisation is unusual, in the sense that it is non-strict in the erroneous behaviours \perp , ω , and no . This choice was made to ensure the ADT is program controlled; if we were to consider sequential composition of ADTs, strictness would be required.

4.1. Downward simulation for process data types

Consider the process data types $(AState, Alnits, \{AOp_i\}_{i \in I}, AFin)$ and $(CState, Clnits, \{COp_i\}_{i \in I}, CFin)$, both with global state G , and a candidate downward simulation relation R between $AState$ and $CState$. We also embed the simulation relation, in order to relate abstract and concrete blocking and divergence correctly:

$$\llbracket R \rrbracket == R \cup \{(\perp, \perp)\} \cup \{\omega\} \times CState_\omega$$

Initialisation

Applying the initialisation condition on the embedded data types with the embedded simulation relation $\llbracket R \rrbracket$ we calculate:

$$\begin{aligned} Clnit &\subseteq Alnit \circ \llbracket R \rrbracket \\ &\equiv \{ \text{definition of embeddings} \} \\ G_{\perp, \omega} \times Clnits &\subseteq (G_{\perp, \omega} \times Alnits) \circ \llbracket R \rrbracket \\ &\equiv \{ \text{definition of } \llbracket R \rrbracket \} \\ Clnits &\subseteq \text{ran}(Alnits \triangleleft R) \end{aligned}$$

Finalisation

This derivation is shown in detail; similar steps in later derivations (revolving around distribution of \cup over \circ , and subsequent simplifications) will be contracted.

$$\begin{aligned} \llbracket R \rrbracket \circ \llbracket CFin \rrbracket &\subseteq \llbracket AFin \rrbracket \\ &\equiv \{ \text{embeddings} \} \\ (R \cup \{(\perp, \perp)\} \cup \{\omega\} \times CState_\omega) \circ (CFin \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_\omega \cup \overline{\text{dom } CFin} \times \{\text{no}\}) \\ &\subseteq \\ AFin \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_\omega \cup \overline{\text{dom } AFin} \times \{\text{no}\} \\ &\equiv \{ \text{distribution of } \cup \text{ over } \circ \} \\ R \circ CFin \cup R \circ \overline{\text{dom } CFin} \times \{\text{no}\} \cup R \circ \{(\perp, \perp)\} \cup R \circ (\{\omega\} \times G_\omega) \cup \{(\perp, \perp)\} \circ CFin \\ &\cup \{(\perp, \perp)\} \circ \overline{\text{dom } CFin} \times \{\text{no}\} \cup \{(\perp, \perp)\} \circ \{(\perp, \perp)\} \cup \{(\perp, \perp)\} \circ (\{\omega\} \times G_\omega) \\ &\cup (\{\omega\} \times CState_\omega) \circ CFin \cup (\{\omega\} \times CState_\omega) \circ \overline{\text{dom } CFin} \times \{\text{no}\} \cup (\{\omega\} \times CState_\omega) \circ \{(\perp, \perp)\} \\ &\cup (\{\omega\} \times CState_\omega) \circ (\{\omega\} \times G_\omega) \\ &\subseteq \\ AFin \cup \overline{\text{dom } AFin} \times \{\text{no}\} \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_\omega \\ &\equiv \{ \perp, \omega, \text{no} \notin \text{domain or range of finalisation or } R \} \\ R \circ CFin \cup (\text{dom}(R \triangleright (\text{dom } CFin)) \times \{\text{no}\}) \cup \emptyset \cup \emptyset \cup \emptyset \cup \emptyset \cup \{(\perp, \perp)\} \cup \emptyset \cup (\{\omega\} \times \text{ran } CFin) \\ &\cup \{(\omega, \text{no})\} \cup \emptyset \cup \{\omega\} \times G_\omega \\ &\subseteq \\ AFin \cup \overline{\text{dom } AFin} \times \{\text{no}\} \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_\omega \\ &\equiv \{ \text{calculus} \} \\ R \circ CFin \subseteq AFin \quad \wedge \quad \text{dom}(R \triangleright (\text{dom } CFin)) \subseteq \overline{\text{dom } AFin} \\ &\equiv \{ \text{calculus, taking domains of first conjunct to make second symmetric} \} \\ R \circ CFin \subseteq AFin \quad \wedge \quad (\text{dom } AFin) \triangleleft R = R \triangleright (\text{dom } CFin) \end{aligned}$$

Note that the last step adds some information from the first conjunct onto the second, and thus adds some “finalisation correctness” onto the “finalisation applicability” condition; however, we will still refer to the second condition using the latter name.

Operations

We first simplify the compositions of an embedded operation $Op = (N, B, D)$ with an embedded simulation,

leading to:

$$\begin{aligned} \llbracket R \rrbracket \circ \llbracket \text{Op} \rrbracket &= R \circ N \cup R \circ (B \times \{\perp\}) \cup R \circ (D \times \text{State}_\omega) \cup \{(\perp, \perp)\} \cup \{\omega\} \times \text{State}_{\perp, \omega} \\ \llbracket \text{Op} \rrbracket \circ \llbracket R \rrbracket &= N \circ R \cup (B_{\perp, \omega} \times \{\perp\}) \cup D_\omega \times S_\omega \end{aligned}$$

For corresponding operations $\text{AOp} = (\text{AN}, \text{AB}, \text{AD})$ and $\text{COp} = (\text{CN}, \text{CB}, \text{CD})$, we have that

$$\begin{aligned} \llbracket R \rrbracket \circ \llbracket \text{COp} \rrbracket &\subseteq \llbracket \text{AOp} \rrbracket \circ \llbracket R \rrbracket \\ &\equiv \{ \text{above; simplifications} \} \\ R \circ \text{CN} \cup R \circ (\text{CB} \times \{\perp\}) \cup R \circ (\text{CD} \times \text{CState}_\omega) \\ &\subseteq \\ \text{AN} \circ R \cup (\text{AB} \times \{\perp\}) \cup \text{AD} \times \text{CState}_\omega \\ &\equiv \{ \text{inclusion of set union} \} \\ R \circ \text{CN} \subseteq \text{AN} \circ R \cup (\text{AB} \times \{\perp\}) \cup \text{AD} \times \text{CState}_\omega \\ \wedge R \circ (\text{CB} \times \{\perp\}) \subseteq \text{AN} \circ R \cup (\text{AB} \times \{\perp\}) \cup \text{AD} \times \text{CState}_\omega \\ \wedge R \circ (\text{CD} \times \text{CState}_\omega) \subseteq \text{AN} \circ R \cup (\text{AB} \times \{\perp\}) \cup \text{AD} \times \text{CState}_\omega \\ &\equiv \{ \text{simplification: domains} \} \\ R \circ \text{CN} \subseteq \text{AN} \circ R \cup \text{AD} \times \text{CState}_\omega \\ \wedge R \circ (\text{CB} \times \{\perp\}) \subseteq (\text{AB} \times \{\perp\}) \wedge R \circ (\text{CD} \times \text{CState}_\omega) \subseteq \text{AD} \times \text{CState}_\omega \\ &\equiv \{ \text{relational calculus} \} \\ \text{AD} \triangleleft R \circ \text{CN} \subseteq \text{AN} \circ R \quad \wedge \quad \text{dom}(R \triangleright \text{CB}) \subseteq \text{AB} \quad \wedge \quad \text{dom}(R \triangleright \text{CD}) \subseteq \text{AD} \end{aligned}$$

These derivations establish the following.

Theorem 4 (Downward simulation for process data types).

The relation R between AState and CState is a downward simulation between the process data types $(\text{AState}, \text{Alnits}, \{\text{AOp}_i\}_{i \in I}, \text{AFin})$ and $(\text{CState}, \text{Clnits}, \{\text{COp}_i\}_{i \in I}, \text{CFin})$, iff

$$\begin{aligned} \text{Clnits} &\subseteq \text{ran}(\text{Alnits} \triangleleft R) \\ R \circ \text{CFin} &\subseteq \text{AFin} \\ (\text{dom AFin}) \triangleleft R &= R \triangleright (\text{dom CFin}) \end{aligned}$$

and $\forall i : I$, for $\text{AOp}_i = (\text{AN}, \text{AB}, \text{AD})$, $\text{COp}_i = (\text{CN}, \text{CB}, \text{CD})$

$$\begin{aligned} \text{AD} \triangleleft R \circ \text{CN} &\subseteq \text{AN} \circ R \\ \text{dom}(R \triangleright \text{CB}) &\subseteq \text{AB} \\ \text{dom}(R \triangleright \text{CD}) &\subseteq \text{AD} \end{aligned}$$

□

The rules for initialisation and finalisation are identical to the usual rules⁸. The three rules for operations are the expected generalisations. The first is “correctness”, ensuring correct after-states, provided the abstract system does not diverge; in the blocking approach, this proviso is immaterial. The second and third rules both relate to what is normally known as “applicability”. When $(B, D) = (\emptyset, \text{dom Op}_i)$ (i.e., the non-blocking interpretation) the second is vacuously true and the third reduces to the usual $\text{ran}(\text{dom AOp}_i \triangleleft R) \subseteq \text{dom COp}_i$; in the opposite (blocking) case, the second reduces to the same familiar condition and the third is trivially true.

4.2. Upward simulation for process data types

Consider the process data types $(\text{AState}, \text{Alnits}, \{\text{AOp}_i\}_{i \in I}, \text{AFin})$ and $(\text{CState}, \text{Clnits}, \{\text{COp}_i\}_{i \in I}, \text{CFin})$, both with global state G , and a candidate downward simulation relation T between CState and AState , embedded

⁸ The applicability rule for finalisation is less common, but is also included in [18].

as follows:

$$\llbracket T \rrbracket == T \cup \{(\perp, \perp)\} \cup \{\omega\} \times \text{AState}_\omega$$

Initialisation A simple calculation shows

$$\text{Clnit} \circ \llbracket T \rrbracket \subseteq \text{Alnit} \equiv \text{ran}(\text{Clnits} \triangleleft T) \subseteq \text{Alnits}$$

Finalisation We derive:

$$\begin{aligned} \llbracket \text{CFin} \rrbracket &\subseteq \llbracket T \rrbracket \circ \llbracket \text{AFin} \rrbracket \\ &\equiv \{ \text{embeddings} \} \\ \text{CFin} \cup (\overline{\text{dom CFin}} \times \{\text{no}\}) &\cup \{(\perp, \perp)\} \cup \{\omega\} \times \mathbf{G}_\omega \\ &\subseteq \\ (T \cup \{(\perp, \perp)\} \cup \{\omega\} \times \text{AState}_\omega) &\circ (\text{AFin} \cup (\overline{\text{dom AFin}} \times \{\text{no}\}) \cup \{(\perp, \perp)\} \cup \{\omega\} \times \mathbf{G}_\omega) \\ &\equiv \{ \text{domains; calculus (see also [18, page 80])} \} \\ \text{CFin} \subseteq T \circ \text{AFin} \wedge \overline{\text{dom CFin}} &\subseteq \text{dom}(T \triangleright \text{dom AFin}) \end{aligned}$$

This cannot be further simplified: the first conjunct implies that every finalisable concrete state is linked to *some* finalisable abstract state, and the second similarly links non-finalisable states. However, neither excludes a concrete state being linked to both finalisable and non-finalisable abstract states. Together they do imply totality of T on CState as usual.

Operations

For corresponding operations $\text{AOp} = (\text{AN}, \text{AB}, \text{AD})$ and $\text{COp} = (\text{CN}, \text{CB}, \text{CD})$, we have that

$$\begin{aligned} \llbracket \text{COp} \rrbracket \circ \llbracket T \rrbracket &\subseteq \llbracket T \rrbracket \circ \llbracket \text{AOp} \rrbracket \\ &\equiv \{ \text{see downward simulation derivation} \} \\ \text{CN} \circ T \cup (\text{CB}_\perp \times \{\perp\}) \cup (\text{CD}_\omega \times \text{AState}_\omega) &\cup \{(\omega, \perp)\} \\ &\subseteq \\ T \circ \text{AN} \cup T \circ (\text{AB} \times \{\perp\}) \cup T \circ (\text{AD} \times \text{AState}_\omega) &\cup \{(\perp, \perp)\} \cup \{\omega\} \times \text{AState}_{\perp, \omega} \\ &\equiv \{ \text{inclusion of union; domain/range based simplifications} \} \\ \text{CN} \circ T \subseteq T \circ \text{AN} \cup T \circ (\text{AD} \times \text{AState}) \\ \wedge \text{CB} \times \{\perp\} \subseteq T \circ (\text{AB} \times \{\perp\}) \wedge \text{CD} \times \text{AState}_\omega &\subseteq T \circ (\text{AD} \times \text{AState}_\omega) \\ &\equiv \{ \text{calculus} \} \\ \text{dom}(T \triangleright \text{AD}) \triangleleft \text{CN} \circ T \subseteq T \circ \text{AN} \wedge \text{CB} \subseteq \text{dom}(T \triangleright \text{AB}) &\wedge \text{CD} \subseteq \text{dom}(T \triangleright \text{AD}) \end{aligned}$$

These derivations establish the following.

Theorem 5 (Upward simulation for process data types).

The relation T between CState and AState is an upward simulation between the process data types $(\text{AState}, \text{Alnits}, \{\text{AOp}_i\}_{i \in I}, \text{AFin})$ and $(\text{CState}, \text{Clnits}, \{\text{COp}_i\}_{i \in I}, \text{CFin})$, iff

$$\begin{aligned} \text{ran}(\text{Clnits} \triangleleft T) &\subseteq \text{Alnits} \\ \text{CFin} &\subseteq T \circ \text{AFin} \\ \overline{\text{dom CFin}} &\subseteq \text{dom}(T \triangleright \text{dom AFin}) \end{aligned}$$

and $\forall i : I$, for $\text{AOp}_i = (\text{AN}, \text{AB}, \text{AD})$, $\text{COp}_i = (\text{CN}, \text{CB}, \text{CD})$

$$\begin{aligned} \text{dom}(T \triangleright \text{AD}) \triangleleft \text{CN} \circ T &\subseteq T \circ \text{AN} \\ \text{CB} &\subseteq \text{dom}(T \triangleright \text{AB}) \\ \text{CD} &\subseteq \text{dom}(T \triangleright \text{AD}) \end{aligned}$$

□

4.3. Simulations on process data types and basic data types

In order to transfer results derived for basic data types, in particular those whose reductions contain an embedding of outputs or refusals, we compare the rules derived above for process data types with the usual definitions of simulations on total data types (Definitions 3 and 4).

As the essential information added in constructing a process data type is attached to the operations only, we would expect the conditions on initialisation and finalisation to be identical, and indeed they are (modulo the differences introduced by process data types using initialisation sets rather than relations, and possible partiality of finalisation). An important consequence of this is in adopting a refusal embedding: this only affects the correctness of finalisation, and so incurs the same extra simulation requirements on process data types as on total ones. Of course this does require the refusal notion to be expressed in terms of triples (N, B, D) , and the resulting requirement (as in [19]) may still end up dominating some of the other conditions on operations. The property (see [18]) that finalisation conditions resulting from an output embedding are trivially true also transfers across.

5. Failures-divergences with internal operations

In the previous section, we derived simulation rules for the generalised process data type by embedding it into a relational model with explicit error values \perp and ω . This section presents the major result of this paper: simulations for relational failures-divergences refinement of data types extended with an internal operation, defined through an embedding in process data types. There are two versions of this: the blocking and the non-blocking approaches, both deriving from the same process data type simulation rules. This reconstructs the approach in [20] which inspired the definition of process data types. The following section will then revisit these two approaches, adding outputs to the data types.

5.1. Basic data type with internal operation

We develop a relational verification method for failures-divergences refinement. Thus, we assume the basic data type is a refusal embedding – in the first instance in the blocking approach for the basic refusal relation characterised by

$$\forall s : \text{State}, E : \mathbb{P}I \bullet (s, E) \in \text{Ref}_{\text{State}} \equiv (\forall i : E \bullet s \notin \text{dom Op}_i)$$

Definition 11 (Basic data type with internal operation).

A basic data type with internal operation is a quintuple $(\text{State}, \text{Init}, \{\text{Op}_i\}_{i \in I}, \tau, \text{Fin})$ such that $(\text{State}, \text{Init}, \{\text{Op}_i\}_{i \in I}, \text{Fin})$ is a basic data type, and the internal operation τ is a relation on State . \square

5.2. Embedding into process data types

The addition of internal operations introduces a large number of, sometimes subtle, issues (see also [17, 18]):

- a) a finite number of internal operations may take place before and after every operation;
- b) in some states, unbounded internal evolution may be enabled (livelock), making those states and the operations leading into those states divergent;
- c) initial states may also be divergent; if only one initial state is, all traces of the ADT are divergent;

- d) the CSP failures-divergences semantics only observes refusals in stable states, i.e., where no internal evolution is possible. However, from the fact that traces involving unstable states are observable, enabledness of operations in such states may be observable; blocking in unstable states is more subtle, however. It also means that observing refusals at finalisation needs to be treated with care, in case the final state is not stable;
- e) although refusal of operations in unstable states is immaterial, the opposite, i.e., enabled operations in unstable states is *not*: we need to ensure that traces arising from operations being “momentarily” available in unstable states are included.

Notation We define the following auxiliary notations in order to deal with internal evolution. This includes notations for (maximal) finite internal evolution which are used to enhance the operations (etc).

$\text{State}\downarrow$ denotes the set of stable states, i.e. those from which it is not possible to do an internal evolution. That is, $v \in \text{State}\downarrow$ iff $v \notin \text{dom } \tau$.

τ^* denotes finite internal evolution, defined as the least fixed point of $\lambda R \bullet \text{id}_{\text{State}} \cup \tau \circledast R$.

$\tau^{*\downarrow} == \tau^* \triangleright (\text{dom } \tau)$ denotes maximal finite internal evolution, leading to a stable state. Note that all unstable states either are divergent, or are linked by finite internal evolution to a non-empty set of stable states.

$\overset{\leftrightarrow}{\text{Op}} == \tau^* \circledast \text{Op} \circledast \tau^*$ represents an operation with internal evolution beforehand and afterwards. Note that $\text{dom } \overset{\leftrightarrow}{\text{Op}} = \text{dom}(\tau^* \circledast \text{Op})$ as $\text{dom } \tau^* = \text{State}$.

$\text{State}\uparrow$ denotes the set of “divergent” states from which unbounded internal evolution is possible. It is defined as the largest fixed point of $\lambda S. \text{dom}(\tau \triangleright S)$. If any initial states are in $\text{State}\uparrow$, the ADT as a whole is divergent.

$\underline{\tau} == (\text{State}\uparrow) \triangleleft \tau$ denotes “relevant” internal transitions, excluding those from states which are already divergent. As everything is possible after divergence, the presence of finite internal behaviour from such states is semantically insignificant. Note that $\underline{\tau}^* = \text{id}_{\text{State}\uparrow} \cup (\text{State}\uparrow) \triangleleft \tau^*$.

liv Op characterises all the states where the application of Op might be followed by unbounded internal evolution, and is defined by

$$\text{liv Op} = \text{dom}(\text{Op} \triangleright \text{State}\uparrow)$$

Note that these are not necessarily states which are *themselves* divergent, but states from where Op might lead to divergent states.

Using these notations, we can construct embeddings of basic data types with an internal operation into process data types for the blocking and non-blocking approaches. This gives a semantics for these types, and we will later validate that, in combination with the right refusal embedding, this correctly encodes failures-divergences semantics.

Definition 12 (Embeddings of internal operations into process data type).

A basic data type with internal operation $(\text{State}, \text{Init}, \{\text{Op}_i\}_{i \in I}, \tau, \text{Fin})$ is embedded into the process data type $(\text{State}, \widetilde{\text{Init}}, \{\widetilde{\text{Op}}_i\}_{i \in I}, \widetilde{\text{Fin}})$ with

$$\begin{aligned} \widetilde{\text{Init}} &== \text{ran}(\text{Init} \circledast \tau^*) \\ \widetilde{\text{Fin}} &== \text{State}\uparrow \triangleleft \tau^{*\downarrow} \circledast \text{Fin} \\ \widetilde{\text{Op}}_i &== (\text{liv Op}_i \triangleleft \overset{\leftrightarrow}{\text{Op}}_i, B_i, D_i) \end{aligned}$$

where in the blocking approach

$$B_i == \overline{\text{dom } \overset{\leftrightarrow}{\text{Op}}_i}$$

$$D_i == \text{liv Op}_i$$

and in the non-blocking approach

$$B_i == \emptyset$$

$$D_i == \text{liv Op}_i \cup \overline{\text{dom } \overset{\leftrightarrow}{\text{Op}}_i}$$

□

We argue informally that this embedding addresses each of the issues listed above, with the exception of a minor one:

- a) Using $\overset{\leftrightarrow}{\text{Op}}$ ensures finite internal evolution before and after every operation.
- b) By including $\text{liv } \overset{\leftrightarrow}{\text{Op}}$ in the divergence domain of an operation, we ensure that livelock becomes divergence.
- c) This encoding does *not* deal with divergent initial states. The process data type could be enhanced with a Boolean to represent initial divergence (and include ω in Inits whenever it is set), however, this pathological situation can more easily be singled out in the definition of refinement.
- d) In our previous work [17, 18] it was sufficient to include internal operations after each operation and initialisation only – the need to include internal operations *before* derives from observing refusals and blocking (in the blocking approach) and divergence outside the precondition (in the non-blocking approach) in stable states only.⁹ In general, in the context of *stable* failures, treating an operation that is not enabled as one that has an artificial result (\perp) does not work across the board – it depends on whether τ is enabled in the same state or not.
The finalisation, additionally, is preceded by maximal finite internal evolution in order to avoid observing refusals in unstable states, and restricted to non-divergent states. Alternatively, we could restrict finalisation to stable states only¹⁰.
Note that a finalisation which observes refusals in the non-blocking case is quite different from the blocking case, as there are no refusals in the non-blocking approach unless outputs are taken into account.
- e) This is addressed by including finite, but *not* necessarily maximal internal behaviour after every operation (and initialisation) – i.e., not considering only the stable after-states.

5.3. Refinement

The definition of refinement for basic data types with internal operations is given in terms of the embedding into process data types, with a small proviso to deal with divergent initial states.

Definition 13 (Refinement with internal operations).

A program controlled basic data type with internal operation $A = (\text{AState}, \text{AInit}, \{\text{AOp}_i\}_{i \in I}, \tau_A, \text{AFin})$ is refined by another such data type $C = (\text{CState}, \text{CInit}, \{\text{COp}_i\}_{i \in I}, \tau_C, \text{CFin})$ iff

- One of A 's initial states is divergent: $(\text{ran } \text{AInit}) \cap \text{AState} \uparrow \neq \emptyset$, or
- None of C 's initial states are divergent, i.e., $(\text{ran } \text{CInit}) \cap \text{CState} \uparrow = \emptyset$, and:
refinement holds between the embeddings of A and C in their embeddings into process data types. \square

By separating out divergent initial states in this way, we can restrict our discussion of simulations to those cases where the process data type embedding correctly reflects the original specification in all respects.

5.4. Correctness of the embedding: blocking approach

First, however, we prove that the embedding of internal operations into process data types, and through that into total data types, correctly reflects the failures-divergences semantics in the blocking approach (which is the more common one for a concurrency context).

⁹ The corresponding encoding in [20] is erroneous in this respect, and causes the relational semantics to make incorrect distinctions by including unstable refusals. This became evident in mechanising the proof.

¹⁰ This works because $\tau^* \upharpoonright \text{Fin} = \tau^* \triangleright (\text{dom } \tau) \text{ Fin} = \tau^* \text{ Fin} \triangleleft (\text{dom } \tau)$, and τ^* is always absorbed by preceding τ^* .

In order to do so, we first characterise formally the failures-divergences semantics of a data type with internal operations (ignoring its finalisation). We use some standard labeled transition notation to facilitate that, as follows. Note that (only) the definition of \xrightarrow{e} will become more complicated later when we also use outputs.

Definition 14 (Transition notations).

For a data type with internal operations $(\text{State}, \mathsf{G} \times \text{Inits}, \{\text{Op}_i\}_{i \in I}, \tau, \text{Fin})$ we define the following notations (ε denotes the empty trace):

$$\begin{aligned} x \xrightarrow{i} x' &== (x, x') \in \text{Op}_i & x \xrightarrow{i} &== \exists x' \bullet x \xrightarrow{i} x' \\ x \xrightarrow{\tau} x' &== (x, x') \in \tau & x \not\xrightarrow{i} &== \neg x \xrightarrow{i} \end{aligned}$$

and \Longrightarrow is the transitive reflexive closure of \longrightarrow , collecting all the labels in a sequence except for τ :

$$x \xrightarrow{s} x' == \exists p : \text{seq}(I \cup \{\tau\}) \bullet x \xrightarrow{p} x' \wedge s = p \upharpoonright I$$

where

$$x \xrightarrow{\langle \rangle} x' == x = x' \qquad x \xrightarrow{p \widehat{\ } q} x' == \exists x'' \bullet x \xrightarrow{p} x'' \wedge x'' \xrightarrow{q} x'$$

and an omitted first state indicates an initial one:

$$\xrightarrow{s} x == \exists x' : \text{Inits} \bullet x' \xrightarrow{s} x$$

□

Failures and divergences are then defined in the usual way. Note that they would be different in the non-blocking approach, producing no refusals but divergences instead.

Definition 15 (Failures-divergences semantics).

For a data type with internal operations $\mathsf{T} = (\text{State}, \mathsf{G} \times \text{Inits}, \{\text{Op}_i\}_{i \in I}, \tau, \text{Fin})$ its divergences are

$$\text{div}(\mathsf{T}) == \{s : \text{seq } I \mid \exists s' : \text{seq } I; x : \text{State} \uparrow \bullet s' \leq s \wedge \xrightarrow{s'} x\}$$

and its failures are

$$f(\mathsf{T}) == \{s : \text{seq } I; E : \mathbb{P} I \mid s \in \text{div}(\mathsf{T}) \vee (\exists x : \text{State} \downarrow \bullet \xrightarrow{s} x \wedge \forall e : E \bullet x \not\xrightarrow{e}) \bullet (s, E)\}$$

□

The correctness of the embedding will be proved using the following lemmas. Recall that we are working in a context where:

- T is defined as in Definition 15; and is program controlled, so $\text{Init} \equiv \mathsf{G} \times \text{ran } \text{Init}$;
- we use the blocking interpretation of T ;
- T observes refusals at finalisation;
- T is not initially divergent.

We have not defined failures and divergences at the level of process data types, and therefore we will have to prove the equivalence of refinement accounting for *two* levels of embedding. Recall that \tilde{x} is the embedding of x into process data types, and $\llbracket x \rrbracket$ is the embedding of x from process data types into basic data types¹¹.

Lemma 1. For all programs $p : \text{seq } I$, we have that

$$p \in \text{div}(\mathsf{T}) \equiv \omega \in \text{ran } p \llbracket \tilde{\ } \rrbracket$$

¹¹ This is different from the precursor paper [20] where the various embeddings directly into basic data types were defined as \widehat{x}^t with varying t , and \tilde{x} was used for embedding simulations.

Proof A simple proof shows that

$$\omega \in \text{ran } p_{\llbracket \widetilde{\mathbb{T}} \rrbracket} \equiv \omega \in \text{ran}(\llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p} \rrbracket)$$

Because \mathbb{T} is not initially divergent, p is non-empty. We now prove

$$\omega \in \text{ran}(\llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p} \rrbracket) \equiv p \in \text{div}(\mathbb{T})$$

by induction over the construction of p .

$$\begin{aligned} \omega \in \text{ran}(\llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p} \rrbracket) & \equiv \{ \text{Let } p = p' \wedge \langle i \rangle \} \\ \omega \in \text{ran}(\llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p}' \rrbracket \circ \llbracket \widetilde{\text{Op}}_i \rrbracket) & \equiv \{ \text{definition of } \widetilde{\text{Op}}_i \} \\ \omega \in \text{ran}(\llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p}' \rrbracket \circ \llbracket (\text{liv } \widetilde{\text{Op}}_i \triangleleft \overrightarrow{\text{Op}}_i, \overrightarrow{\text{dom } \text{Op}}_i, \text{liv } \overrightarrow{\text{Op}}_i) \rrbracket) & \equiv \{ \text{definition of } \llbracket (\mathbb{N}, \mathbb{B}, \mathbb{D}) \rrbracket: \omega \text{ only from } \mathbb{D} \text{ and } \omega \} \\ \omega \in \text{ran}(\llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p}' \rrbracket) \vee \text{ran}(\llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p}' \rrbracket) \cap \text{liv } \overrightarrow{\text{Op}}_i & \neq \emptyset \end{aligned}$$

The first disjunct can only hold when p' is nonempty (because \mathbb{T} is not initially divergent). In that case, by induction $p' \in \text{div}(\mathbb{T})$ and, as divergences are closed under extension, $p \in \text{div}(\mathbb{T})$ as required. For the second disjunct (including the base case), we have:

$$\begin{aligned} \text{ran}(\llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p}' \rrbracket) \cap \text{liv } \overrightarrow{\text{Op}}_i & \neq \emptyset \\ & \equiv \{ \text{definition of } \overrightarrow{\text{Op}}_i \text{ and liv; relations} \} \\ \llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p}' \rrbracket \circ (\tau^* \circ \text{Op}_i \circ \tau^*) \triangleright \text{State} \uparrow & \neq \emptyset \\ & \equiv \{ \llbracket \widetilde{p}' \rrbracket \text{ (in the base case: } \llbracket \widetilde{\text{Init}} \rrbracket) \text{ and } \text{State} \uparrow \text{ closed under composition with } \tau^* \} \\ \llbracket \widetilde{\text{Init}} \rrbracket \circ \llbracket \widetilde{p}' \rrbracket \circ \text{Op}_i \triangleright \text{State} \uparrow & \neq \emptyset \\ & \equiv \{ \text{definition of } \implies; \text{ no blocking in } p' \text{ as } \perp \notin \text{dom } \text{Op}_i \} \\ \exists x : \text{State} \uparrow \bullet \xrightarrow{p} x & \\ \Rightarrow \{ \text{definition of } \text{div} \} & \\ p \in \text{div}(\mathbb{T}) & \end{aligned}$$

We have now proved that $p = p' \wedge \langle i \rangle$ is divergent when p' is, and that p is divergent when p' is not and Op_i diverges in some final state of p' . As these are the only two cases in which p could be divergent, we have equivalence (rather than just implication) as required. \square

Lemma 2. For all programs $p : \text{seq } I$ and sets of events $E : \mathbb{P} I$ such that $p \notin \text{div}(\mathbb{T})$, we have that

$$(p, E) \in f(\mathbb{T}) \equiv E \in \text{ran } p_{\llbracket \widetilde{\mathbb{T}} \rrbracket}$$

Proof We generalise the definition of (non-divergent) refusals to

$$f'(S : \mathbb{P} \text{State}) == \{ s : \text{seq } I; E : \mathbb{P} I \mid \exists x : \text{State} \downarrow; y : S \bullet y \xrightarrow{s} x \wedge \forall e : E \bullet x \not\xrightarrow{e} \}$$

and the lemma to

$$\omega \notin \text{ran}(S \triangleleft \llbracket \widetilde{p} \rrbracket) \wedge S \cap \text{State} \uparrow = \emptyset \Rightarrow (p, E) \in f'(S) \equiv E \in \text{ran}(S \triangleleft \llbracket \widetilde{p} \rrbracket \circ \llbracket \widetilde{\text{Fin}} \rrbracket)$$

which proves the lemma when we instantiate S to $\text{ran } \text{Init}$. The generalisation is proved by induction on p .

Base case When $p = \langle \rangle$ we have:

$$(p, E) \in f'(S)$$

$$\begin{aligned}
&\equiv \{ \text{base case; definition of } f' \} \\
\exists x : \text{State} \downarrow ; y : \text{S} \bullet y &\xrightarrow{\varepsilon} x \wedge \forall e : E \bullet x \not\xrightarrow{e} \\
&\equiv \{ \text{definition of } \text{Fin}, \text{ definition of } \tau^{*!} \} \\
E \in \text{ran}(\text{S} \triangleleft (\text{State} \uparrow \triangleleft \tau^{*!} \wp \text{Fin})) \\
&\equiv \{ \text{embeddings of } \text{Fin}; \text{ insert empty program } \} \\
E \in \text{ran}(\text{S} \triangleleft \llbracket \tilde{p} \rrbracket \wp \llbracket \widetilde{\text{Fin}} \rrbracket)
\end{aligned}$$

Induction step Let $p = \langle i \rangle \hat{\ } p'$.

$$\begin{aligned}
(p, E) \in f'(\text{S}) \\
&\equiv \{ \text{definition of } f' \} \\
\exists x : \text{State} \downarrow ; y : \text{S} \bullet y &\xrightarrow{\langle i \rangle \hat{\ } p'} x \wedge \forall e : E \bullet x \not\xrightarrow{e} \\
&\equiv \{ \text{definition of } \implies ; \text{ relational calculus } \} \\
\exists x : \text{State} \downarrow ; y' : \text{ran}(\text{S} \triangleleft \overset{\leftrightarrow}{\text{Op}}_i) \bullet y' &\xrightarrow{p'} x \wedge \forall e : E \bullet x \not\xrightarrow{e} \\
&\equiv \{ \text{definition of } f' \} \\
(p', E) \in f'(\text{ran}(\text{S} \triangleleft \overset{\leftrightarrow}{\text{Op}}_i)) \\
&\equiv \{ \text{induction } \} \\
E \in \text{ran}(\text{ran}(\text{S} \triangleleft \overset{\leftrightarrow}{\text{Op}}_i) \triangleleft \llbracket \tilde{p}' \rrbracket \wp \llbracket \widetilde{\text{Fin}} \rrbracket) \\
&\equiv \{ \text{embeddings of operations } \} \\
E \in \text{ran}(\text{S} \triangleleft \llbracket \overset{\leftrightarrow}{\text{Op}}_i \rrbracket \wp \llbracket \tilde{p}' \rrbracket \wp \llbracket \widetilde{\text{Fin}} \rrbracket) \\
&\equiv \{ \text{program composition } \} \\
E \in \text{ran}(\text{S} \triangleleft \llbracket \tilde{p} \rrbracket \wp \llbracket \widetilde{\text{Fin}} \rrbracket)
\end{aligned}$$

This completes the proof of Lemma 2. \square

Note that, on first impression maybe surprisingly, the two lemmas capture all the relevant information (failures and divergences) without making any reference to \perp in the relational semantics. In fact, the information provided by \perp is redundant when refusals are observed at finalisation. Indeed, in [19] the simulation conditions derived from finalisations which observe refusals dominate the “applicability” conditions (derived by eliminating \perp), pointing at a similar redundancy in the situation without internal operations. The redundancy here is formalised in the following lemma (a healthiness condition).

Lemma 3. When no strict prefix p' of p is in $\text{div}(\text{T})$, then

$$\perp \in \text{ran } p \llbracket \widetilde{\text{T}} \rrbracket \equiv \exists p', p'', i \bullet p = p' \hat{\ } \langle i \rangle \hat{\ } p'' \wedge \{i\} \in \text{ran } p' \llbracket \widetilde{\text{T}} \rrbracket$$

i.e., when p is not a trace that already diverged previously and it may lead to blocking, then somewhere in p there has to be an index i which could be refused at that point. \square

Proof Similar to Lemma 2, generalising $\text{ran } \text{Init}$ in $p \llbracket \widetilde{\text{T}} \rrbracket$ to an arbitrary set and then by induction over p . \square

The main correctness theorem for the embedding is the following.

Theorem 6 (Equivalence of failures-divergences and data refinement).

For two data types A and C , failures-divergences refinement holds iff data refinement holds between their embeddings.

Proof By mutual implication. The proof that relational refinement implies failures-divergences refinement

uses Lemma 1 for divergences, and Lemma 2 for failures. The reverse direction has as its demonstrandum

$$x \in \text{ran } p_{\llbracket \widetilde{C} \rrbracket} \Rightarrow x \in \text{ran } p_{\llbracket \widetilde{A} \rrbracket}$$

which is proved by case distinction on x (\perp , ω , **no** or a refusal E), assuming failures-divergences refinement and using all three lemmas. \square

5.5. Simulations in the blocking approach

As failures-divergences refinement is correctly represented by relational refinement, we can instantiate simulation rules from Theorems 4 and 5 into simulation rules for data types with internal operation.

5.5.1. Downward simulation

Instantiating Theorem 4 with the embedding of Definition 12 initially gives the following conditions.

$$\text{CInit} \circledast \tau_C^* \subseteq \text{AInit} \circledast \tau_A^* \circledast R \tag{3}$$

$$R \circledast \text{CState} \uparrow \triangleleft \tau_C^{*|} \circledast \text{CFin} \subseteq \text{AState} \uparrow \triangleleft \tau_A^{*|} \circledast \text{AFin} \tag{4}$$

$$(\text{AState} \uparrow \triangleleft R) = (R \triangleright \text{CState} \uparrow) \tag{5}$$

and for matching operations **AOp** and **COp**:

$$(\text{liv } \overleftrightarrow{\text{AOp}} \triangleleft R \triangleright \text{liv } \overleftrightarrow{\text{COp}}) \circledast \overleftrightarrow{\text{COp}} \subseteq \overleftrightarrow{\text{AOp}} \circledast R \tag{6}$$

$$\text{dom}(R \triangleright \text{liv } \overleftrightarrow{\text{COp}}) \subseteq \text{liv } \overleftrightarrow{\text{AOp}} \tag{7}$$

$$\text{ran}(\text{dom } \overleftrightarrow{\text{AOp}} \triangleleft R) \subseteq \text{dom } \overleftrightarrow{\text{COp}} \tag{8}$$

In order to simplify these, in particular to take into account the particular nature of finalisation (observing refusals), we use the following theorem.

Theorem 7 (Downward simulation closed under \mathcal{T}_C^*).

If the relation $R \subseteq \text{AState} \times \text{CState}$ is a downward simulation between program controlled basic data types with internal operations $A = (\text{AState}, \text{AInit}, \{\text{AOp}_i\}_{i \in I}, \tau_A, \text{AFin})$ and $C = (\text{CState}, \text{CInit}, \{\text{COp}_i\}_{i \in I}, \tau_C, \text{CFin})$ then $R \circledast \mathcal{T}_C^*$ is also a downward simulation between A and C .

Proof By showing that (3) to (8) above imply the same conditions with $R \circledast \mathcal{T}_C^*$ substituted for R . \square

The relevance of the theorem is that we can, without loss of generality, restrict ourselves to retrieve relations R which satisfy $R = R \circledast \mathcal{T}_C^*$. The initialisation condition (3) can then be simplified to $\text{CInit} \subseteq \text{AInit} \circledast \tau_A^* \circledast R$. The ‘‘correctness’’ condition (6) can be simplified using the divergence condition (7) to

$$(\text{liv } \overleftrightarrow{\text{AOp}} \triangleleft R) \circledast \overleftrightarrow{\text{COp}} \subseteq \overleftrightarrow{\text{AOp}} \circledast R$$

The ‘‘finalisation’’ condition (4) is implied by the ‘‘blocking’’ condition (8) and finalisation applicability (5). Because the refusal finalisation refers explicitly to sets of refused operations, this proof is performed at the predicate calculus level:

$$\begin{aligned} R \circledast \text{CState} \uparrow \triangleleft \tau_C^{*|} \circledast \text{CFin} &\subseteq \text{AState} \uparrow \triangleleft \tau_A^{*|} \circledast \text{AFin} \\ &\equiv \{ \text{definition of } \subseteq \} \end{aligned}$$

$$\begin{aligned} \forall a : AState; E : \mathbb{P} I \bullet (\exists c : CState \bullet (a, c) \in R \wedge c \notin CState \uparrow \wedge \exists c' : CState \bullet (c, c') \in \tau_C^{*|} \wedge (c', E) \in Ref) \\ \Rightarrow a \notin AState \uparrow \wedge \exists a' : AState \bullet (a, a') \in \tau_A^{*|} \wedge (a', E) \in Ref \end{aligned}$$

$\equiv \{ \text{predicate calculus} \}$

$$\begin{aligned} \forall a : AState; E : \mathbb{P} I; c, c' : CState \bullet (a, c) \in R \wedge c \notin CState \uparrow \wedge (c, c') \in \tau_C^{*|} \wedge (c', E) \in Ref \\ \Rightarrow a \notin AState \uparrow \wedge \exists a' : AState \bullet (a, a') \in \tau_A^{*|} \wedge (a', E) \in Ref \end{aligned}$$

We prove this by contradiction. Assume the implicant holds, but not the consequent. Thus, $(a', E) \notin Ref_A$ so for some $i \in E$ we have that $a' \in \text{dom } AOp_i$. Theorem 7 implies that $(a, c') \in R$; from the blocking condition (8) it then follows that $c' \in \text{dom } \overleftrightarrow{COp}_i$, and because c' is stable, also $c' \in \text{dom } COp_i$, contradicting $(c', E) \in Ref_C$. This completes the proof that the finalisation correctness condition is dominated by the other conditions. The simplifications are summarised in the following theorem.

Theorem 8 (Downward simulation for data types with internal operations).

The relation $R' \subseteq AState \times CState$ is a downward simulation between program controlled basic data types with internal operations $A = (AState, AInit, \{AOp_i\}_{i \in I}, \tau_A, AFin)$ and $C = (CState, CInit, \{COp_i\}_{i \in I}, \tau_C, CFin)$ iff the following conditions hold for $R = R' \circ \tau_C^*$:

$$CInit \subseteq AInit \circ \tau_A^* \circ R \tag{9}$$

$$(AState \uparrow \triangleleft R) = (R \triangleright CState \uparrow) \tag{10}$$

and for matching operations AOp and COp :

$$(\text{liv } \overleftrightarrow{AOp} \triangleleft R) \circ \overleftrightarrow{COp} \subseteq \overleftrightarrow{AOp} \circ R \tag{11}$$

$$\text{dom}(R \triangleright \text{liv } \overleftrightarrow{COp}) \subseteq \text{liv } \overleftrightarrow{AOp} \tag{12}$$

$$\text{ran}(\text{dom } \overleftrightarrow{AOp} \triangleleft R) \subseteq \text{dom } \overleftrightarrow{COp} \tag{13}$$

□

For specifications in Z , embedded into relations in the usual way, this leads to the following conditions:

$$\mathbf{DS.Init.}\tau \quad \forall CState' \bullet CInit \Rightarrow \exists AState' \bullet AInit \circ \tau_A^* \wedge R'$$

$$\mathbf{DS.App.}\tau \quad \forall CState; AState; i : I \bullet R \wedge \text{pre } \overleftrightarrow{AOp}_i \Rightarrow \text{pre } \overleftrightarrow{COp}_i$$

$$\mathbf{DS.CorrBlock.}\tau \quad \forall i : I; AState; CState; CState' \bullet \neg \text{liv } \overleftrightarrow{AOp}_i \wedge R \wedge \overleftrightarrow{COp}_i \Rightarrow \exists AState' \bullet R' \wedge \overleftrightarrow{AOp}_i$$

$$\mathbf{DS.DivStates} \quad \forall R \bullet AState \uparrow \Leftrightarrow CState \uparrow$$

$$\mathbf{DS.DivOp} \quad \forall CState; AState; i : I \bullet R \wedge \text{liv } \overleftrightarrow{COp}_i \Rightarrow \text{liv } \overleftrightarrow{AOp}_i$$

using the obvious Z versions of notions such as \overleftrightarrow{Op} and $\text{liv } Op$, and assuming $R = R \circ \tau_C^*$. Note, that in the absence of internal operations the first three conditions collapse to the corresponding conditions given in Section 3.3 above.

5.5.2. Upward simulation

Instantiating Theorem 5 with the embedding of Definition 12 initially gives the following conditions.

$$CInit \circ \tau_C^* \circ T \subseteq AInit \circ \tau_A^* \tag{14}$$

$$\text{CState} \uparrow \triangleleft \tau_C^* \circ \text{CFin} \subseteq \text{T} \circ \text{AState} \uparrow \triangleleft \tau_A^* \circ \text{AFin} \quad (15)$$

$$\text{CState} \uparrow \subseteq \text{dom}(\text{T} \triangleright \text{AState} \uparrow) \quad (16)$$

and for matching operations AOp and COp :

$$\text{dom}(\text{T} \triangleright \text{liv} \overleftrightarrow{\text{AOp}}) \triangleleft (\text{liv} \overleftrightarrow{\text{COp}} \triangleleft \overleftrightarrow{\text{COp}}) \circ \text{T} \subseteq \text{T} \circ (\text{liv} \overleftrightarrow{\text{AOp}} \triangleleft \overleftrightarrow{\text{AOp}}) \quad (17)$$

$$\text{liv} \overleftrightarrow{\text{COp}} \subseteq \text{dom}(\text{T} \triangleright \text{liv} \overleftrightarrow{\text{AOp}}) \quad (18)$$

$$\overline{\text{dom} \overleftrightarrow{\text{COp}}} \subseteq \text{dom}(\text{T} \triangleright \text{dom} \overleftrightarrow{\text{AOp}}) \quad (19)$$

We can prove a stronger theorem about closure of simulation under internal evolution in this case:

Theorem 9 (Upward simulation closed under internal evolution).

If the relation $\text{T} \subseteq \text{CState} \times \text{AState}$ is an upward simulation between program controlled basic data types with internal operations $\text{A} = (\text{AState}, \text{AInit}, \{\text{AOp}_i\}_{i \in I}, \tau_A, \text{AFin})$ and $\text{C} = (\text{CState}, \text{CInit}, \{\text{COp}_i\}_{i \in I}, \tau_C, \text{CFin})$ then $\tau_C^* \circ \text{T} \circ \tau_A^*$ is also an upward simulation between A and C .

Proof By showing that (14) to (19) above imply the same conditions with $\tau_C^* \circ \text{T} \circ \tau_A^*$ substituted for T . \square

The initialisation condition can be simplified using Theorem 9, clearly $\tau_C^* \circ \text{T} = \text{T}$. Using totality of T and the blocking condition, the correctness condition can be simplified to remove two anti-restrictions.

In general, the conditions are not as clear-cut concerning divergent states as the downward simulation ones, where finalisation applicability ensures that the simulation links divergent states to divergent states only. Here, the finalisation conditions merely ensure that every abstract divergent state is linked to a concrete one, and vice versa. The blocking condition (19) also requires divergent concrete states to be linked to abstract states allowing the same operations. This is clearly a requirement inherited from the relational refinement theory that is redundant and irrelevant due to the chaotic interpretation of divergence represented in our semantics. Thus, using a general semantic theorem we restrict condition (19) to non-divergent states only; however, in those states it is dominated by the finalisation correctness condition (15). This is summarised, with some simplifications, in the following theorem.

Theorem 10 (Upward simulation for data types with internal operations).

The relation $\text{T}' \subseteq \text{CState} \times \text{AState}$ is an upward simulation between program controlled basic data types with internal operations $\text{A} = (\text{AState}, \text{AInit}, \{\text{AOp}_i\}_{i \in I}, \tau_A, \text{AFin})$ and $\text{C} = (\text{CState}, \text{CInit}, \{\text{COp}_i\}_{i \in I}, \tau_C, \text{CFin})$ iff the following conditions hold for $\text{T} = \tau_C^* \circ \text{T}' \circ \tau_A^*$:

$$\text{CInit} \circ \text{T} \subseteq \text{AInit} \circ \tau_A^* \quad (20)$$

$$\text{CState} \uparrow \triangleleft \tau_C^* \circ \text{Ref}_C \subseteq (\text{T} \triangleleft \text{AState} \downarrow) \circ \text{Ref}_A \quad (21)$$

$$\text{CState} \uparrow \subseteq \text{dom}(\text{T} \triangleright \text{AState} \uparrow) \quad (22)$$

and for matching operations AOp and COp :

$$\text{dom}(\text{T} \triangleright \text{liv} \overleftrightarrow{\text{AOp}}) \triangleleft \tau_C^* \circ \text{COp} \circ \text{T} \subseteq \text{T} \circ \text{AOp} \circ \tau_A^* \quad (23)$$

$$\text{liv} \overleftrightarrow{\text{COp}} \subseteq \text{dom}(\text{T} \triangleright \text{liv} \overleftrightarrow{\text{AOp}}) \quad (24)$$

\square

For specifications in \mathcal{Z} , this leads to the following conditions. (See [4] for the essential steps in rephrasing the finalisation condition.)

$$\begin{aligned}
 \mathbf{US.Init}.\tau & \forall CState'; AState' \bullet T' \wedge CInit \Rightarrow AInit \circlearrowleft \tau_A^* \\
 \mathbf{US.CorrBlock}.\tau & \forall i : I; CState; CState'; AState' \bullet \\
 & (\forall AState \bullet T \Rightarrow \neg \text{liv } \overleftrightarrow{AOp}_i) \wedge COp_i \wedge T' \Rightarrow \exists AState \bullet T \wedge AOp_i \circlearrowleft \tau_A^* \\
 \mathbf{US.FinRef}.\tau & \forall CState; CState' \mid \neg CState \uparrow \wedge \tau_C^* \wedge CState' \downarrow \bullet \\
 & \exists AState \bullet T \wedge AState \downarrow \wedge \forall i : I \bullet \text{pre } AOp_i \Rightarrow (\text{pre } COp_i)' \\
 \mathbf{US.DivStates} & \forall CState \bullet CState' \uparrow \Rightarrow \exists AState \bullet T \wedge AState \uparrow \\
 \mathbf{US.DivOp} & \forall i : I; CState \bullet \text{liv } \overleftrightarrow{COp}_i \Rightarrow \exists AState \bullet T \wedge \text{liv } AOp_i
 \end{aligned}$$

Again note that in the absence of internal operations, these conditions collapse to those in Section 3.3.

5.6. The non-blocking approach

Having considered refinement and simulation for data types with internal operations in the blocking approach in great detail, we now turn to the non-blocking approach, (as previously) not yet considering outputs. In that context, the semantics is considerably simpler: there are no refusals beyond those after divergence (the basic finalisation maps to a single global value representing a non-divergent run; its totalisation adds a representation of divergence). The embedding into process data types was already given above in Definition 12, the only difference being that an operation Op is embedded by $(\mathbf{B}, \mathbf{N}, \mathbf{D})$ where

$$\mathbf{B} == \emptyset \qquad \mathbf{D} == \text{liv } \overleftrightarrow{\text{Op}} \cup \text{dom } \overleftrightarrow{\text{Op}}$$

As discussed in Section 3.4, in this context failures-divergences refinement is equivalent to traces-divergences refinement. Moreover, the trace sets for all ADTs with the same alphabet are identical. We thus need to look at divergences only. For a definition of (non-blocking) failures-divergences semantics analogous to Definition 15, Lemma 1 holds and forms the essence of the correctness proof of this embedding.

5.6.1. Downward simulation

We now extract the simulation rules for the non-blocking approach in a way similar to that in Sections 5.5.1 and 5.5.2. Instantiating Theorem 4 with the embedding of Definition 12 initially gives:

$$CInit \circlearrowleft \tau_C^* \subseteq AInit \circlearrowleft \tau_A^* \circlearrowleft R \tag{25}$$

$$R \circlearrowleft CState \uparrow \triangleleft \tau_C^* \circlearrowleft CFin \subseteq AState \uparrow \triangleleft \tau_A^* \circlearrowleft AFin \tag{26}$$

$$(AState \uparrow \triangleleft R) = (R \triangleright CState \uparrow) \tag{27}$$

and for matching operations AOp and COp :

$$(\text{liv } \overleftrightarrow{\text{AOp}} \cup \text{dom } \overleftrightarrow{\text{AOp}}) \triangleleft R \circlearrowleft (\text{liv } \overleftrightarrow{\text{COp}}) \triangleleft \overleftrightarrow{\text{COp}} \subseteq (\text{liv } \overleftrightarrow{\text{AOp}} \triangleleft \overleftrightarrow{\text{AOp}}) \circlearrowleft R \tag{28}$$

$$\text{dom}(R \triangleright (\text{liv } \overleftrightarrow{\text{COp}} \cup \text{dom } \overleftrightarrow{\text{COp}})) \subseteq \text{liv } \overleftrightarrow{\text{AOp}} \cup \text{dom } \overleftrightarrow{\text{AOp}} \tag{29}$$

$$\text{dom}(R \triangleright \emptyset) \subseteq \emptyset \tag{30}$$

These can be simplified, although there is no theorem analogous to Theorem 7 (simulations closed under τ) – this is due to the fact that properties such as

$$s \in \text{dom } \overleftrightarrow{\text{Op}} \wedge (s, s') \in \tau^* \Rightarrow s' \in \text{dom } \overleftrightarrow{\text{Op}}$$

do not hold in general. The finalisation condition (26) is implied by finalisation applicability (27) for the trivial finalisation. The blocking condition (30) is obviously satisfied. Two domain restrictions in (28) can be removed, one due to preservation of divergence in (29). All in all, this establishes the following.

Theorem 11 (Downward simulation for data types with internal operations (non-blocking)).

The relation $R \subseteq AState \times CState$ is a downward simulation between program controlled basic data types with internal operations $A = (AState, AInit, \{AOp_i\}_{i \in I}, \tau_A, AFin)$ and $C = (CState, CInit, \{COp_i\}_{i \in I}, \tau_C, CFin)$ iff the following conditions hold:

$$CInit \circ \tau_C^* \subseteq AInit \circ \tau_A^* \circ R \quad (31)$$

$$(AState \uparrow \triangleleft R) = (R \triangleright CState \uparrow) \quad (32)$$

and for matching operations AOp and COp :

$$(\text{liv } AOp \cup \overline{\text{dom } AOp}) \triangleleft R \circ COp \subseteq AOp \circ R \quad (33)$$

$$\text{dom}(R \triangleright (\text{liv } COp \cup \overline{\text{dom } COp})) \subseteq \text{liv } AOp \cup \overline{\text{dom } AOp} \quad (34)$$

□

For specifications in Z this leads to the following conditions (using $\Delta CState$ for $CState$; $CState'$):

$$\mathbf{DS.Init.}\tau\tau \ \forall CState' \bullet CInit \circ \tau_C^* \Rightarrow \exists AState' \bullet AInit \circ \tau_A^* \wedge R'$$

$$\mathbf{DS.AppDivOp} \ \forall CState; AState; i : I \bullet R \wedge \text{pre } AOp_i \wedge \neg \text{liv } AOp_i \Rightarrow \text{pre } COp_i \wedge \neg \text{liv } COp_i$$

$$\mathbf{DS.CorrNonBlock.}\tau \ \forall i : I; AState; \Delta CState \bullet \neg \text{liv } AOp_i \wedge \text{pre } AOp_i \wedge R \wedge \overline{COp_i} \Rightarrow \exists AState' \bullet AOp_i \wedge R'$$

$$\mathbf{DS.DivStates} \ \forall AState; CState \mid R \bullet AState \uparrow \Leftrightarrow CState \uparrow$$

5.6.2. Upward simulation

Again instantiating the upward simulation theorem for process data types with the non-blocking embedding into process data types, we obtain the following, after some small simplifications.

Theorem 12 (Upward simulation for data types with internal operations (non-blocking)).

The relation $T \subseteq CState \times AState$ is an upward simulation between program controlled basic data types with internal operations $A = (AState, AInit, \{AOp_i\}_{i \in I}, \tau_A, AFin)$ and $C = (CState, CInit, \{COp_i\}_{i \in I}, \tau_C, CFin)$ iff the following conditions hold:

$$CInit \circ \tau_C^* \circ T \subseteq AInit \circ \tau_A^* \quad (35)$$

$$\overline{CState \uparrow} \subseteq \text{dom}(T \triangleright AState \uparrow) \quad (36)$$

$$CState \uparrow \subseteq \text{dom}(T \triangleright AState \uparrow) \quad (37)$$

and for matching operations AOp and COp , using $\text{div } Op == \text{liv } Op \cup \overline{\text{dom } Op}$:

$$\text{dom}(T \triangleright \text{div } AOp) \triangleleft COp \circ T \subseteq T \circ AOp \quad (38)$$

$$\text{div } COp \subseteq \text{dom}(T \triangleright \text{div } AOp) \quad (39)$$

□

For specifications in Z , this leads to the following conditions:

$$\begin{aligned}
\mathbf{US.Init.}\tau\tau & \forall CState'; AState' \bullet CInit \text{ } \S \tau_C^* \wedge T' \Rightarrow AInit \text{ } \S \tau_A^* \\
\mathbf{US.AppDivOp} & \forall i : I; CState \bullet \text{div } COp_i \Rightarrow \exists AState \bullet T \wedge \text{div } AOp_i \\
\mathbf{US.CorrNonBlock.}\tau & \forall i : I; \Delta CState; AState' \bullet \\
& (\forall AState \bullet T \Rightarrow \neg \text{div } AOp_i) \wedge \overleftrightarrow{COp}_i \wedge T' \Rightarrow \exists AState \bullet T \wedge \overleftrightarrow{AOp}_i \\
\mathbf{US.DivStates} & \forall CState \mid CState \uparrow \bullet \exists AState \bullet T \wedge AState \uparrow \\
\mathbf{US.NonDivStates} & \forall CState \mid \neg CState \uparrow \bullet \exists AState \bullet T \wedge \neg AState \uparrow
\end{aligned}$$

using the analogous definition of $\text{div } Op$, etc.

The downward and upward conditions derived in this section extend those discussed in Section 3.3 by adding internal evolution. As discussed in Section 3.4 no additional conditions are necessary in order to achieve equivalence with failures-divergences refinement – that only becomes necessary when we add outputs to the model, which we do now for both blocking and non-blocking models.

6. Outputs

In this section we enhance the results of the previous section by also considering data types with outputs. The course of our argument will be:

- The addition of outputs as characterised in data types with output embeddings has a very localised effect on the refinement conditions. This is explored in Section 6.1.
- The addition of outputs leads to refusals, even in the non-blocking approach, and even when outputs are deterministic; as a consequence more stringent simulation conditions derive from the finalisation observing refusals. This is explored in Sections 6.2 and 6.3.

6.1. Refinement conditions for output embeddings

Embedding, for example, a Z state-and-operations specification with state *State* and output type *Output* into a relational framework normally involves the creation of what we have called an output embedding, where the global and local state contain an output sequence as well as the “real” state, see Definition 6.

The standard derivation of Z refinement rules from such embeddings is given in [46] and [18, Section 4.5]. From the latter, it is clear that the properties contained in Definition 6 suffice to derive refinement conditions which vary only in small details from those without outputs:

- the initialisation and finalisation applicability conditions are unaffected;
- all quantifications over after-states in the operation conditions, including in the definition of the precondition *pre*, are extended with the same quantification over the operation’s output *Output*.

This result is obtained using a retrieve relation that is the identity on the output sequence – this is enforced by the particular finalisation which prevents change of output (type) in refinement (IO-refinement [18, Chapter 10] removes this restriction).

Consider a process data type D , its reduction D_r and its embedding in total data types D_e , where D_r is an output embedding. First, clearly D_e is *not* an output embedding: its global state will be $(\mathbf{GB} \times \text{seq } Output)_{\perp, \omega}$ when D_r has $\mathbf{GB} \times \text{seq } Output$. However, the construction of D_e still guarantees (in both the blocking and

non-blocking case) the crucial property: an operation's result is independent of previously produced outputs. A state (ls, os) being in the domain of an operation only depends on ls , due to the last condition of Definition 6, and thus ls alone determines whether an operation will block or diverge for that reason.

Some obvious restrictions on the internal operation τ are also required to ensure this. First, τ cannot produce outputs (as this would make its occurrence visible), and second, it must be independent of previous outputs just like normal operations (cf. the last condition of Definition 6). So in the context of data types with outputs, we will also require:

$$((ls, os), (ls2, os2)) \in \tau \Rightarrow os2 = os \quad (40)$$

$$((ls, os), (ls2, os)) \in \tau \Rightarrow ((ls, os2), (ls2, os2)) \in \tau \quad (41)$$

This guarantees that the occurrence of livelock is independent of previous outputs, and thus the previous output sequence does not affect outcomes of operations in *any* of the three parts (N, B, D) of an operation. Consequently, the derivations in the previous section can be adapted with minor modifications as previously.

From this point on, we will concentrate on deriving the Z rules (which are explicit about outputs) rather than the relational ones, taking the rules derived in the previous section as our basis. Note that the embedding of Z operations into relational ones ensures that conditions in Definition 6 and properties (40) to (41) are satisfied. Additional consequences from output refusals in finalisations will be outlined in the next subsections.

6.2. Outputs in the blocking approach

Recall that the inclusion of outputs changes the notion of an event – rather than just an index $i : I$, it now also includes an output value. Traces (including divergences) and refusal sets will have these events as their elements. Definition 8 provides two possible refusal finalisations in the blocking approach; here we only consider the *demonic* view of output, where the system is in charge of selecting an output value and consequently output values may be refused if alternative output values are possible. This is characterised by the predicate

$$Fcond(s, E) == \forall(i, out) : E \bullet (\neg \exists Op_i \bullet \theta State = s \wedge \theta Output = out) \vee (\exists Op_i \bullet \theta State = s \wedge \theta Output \neq out \wedge (i, \theta Output) \notin E)$$

The proof that the relational embedding using this finalisation correctly represent the failures-divergences semantics characterised by demonic output refusals proceeds similarly to the proof of Theorem 6, with the following modifications:

- Traces (divergences) and refusals now have (operation, output) pairs as their elements.
- Lemma 1 (characterisation of divergences) needs to be strengthened, as the output generated up to the point of divergence needs to be recovered before the state degenerates to ω .
- Lemma 2 (correctness of observation of a successful run) has no such issues, as the output sequence forms part of the global state value generated.
- Lemma 3 (consistency of \perp with refusals) is still essentially correct (generalising the refused operation to refusing all possible outputs for it), and allows recovery of the output sequence at the point of blocking.
- The proof of the main theorem is then mostly identical, using modified lemmas as suggested above.

6.2.1. Downward simulation

As indicated above, the initialisation conditions and conditions for operations remain unchanged from Section 5.5.1. Finalisation applicability is unchanged: we still finalise only in non-divergent states. What is also

unchanged is that in any such state, we observe all possible refusals in all stable states reachable from it by (necessarily finite) internal evolution. Moreover, we can assume that the retrieve relation is closed under composition with τ_C^* (Theorem 7). Using similar reasoning as for the corresponding condition in Section 5.5.1, we end up with the proof obligation:

$$\begin{aligned} & \forall a : AState; c' : CState; E \bullet (a, c') \in R \wedge c' \in CState \downarrow \wedge Fcond(c', E) \\ & \Rightarrow \exists a' : AState \bullet (a, a') \in \tau_A^* \wedge Fcond(a', E) \end{aligned}$$

Unlike in Section 5.5.1, this is *not* implied by the other conditions. This is clear from the following counterexample.

$\frac{AState}{a : a0 \mid a1 \mid a2}$	$\frac{CState}{c : \{c0\}}$
$\frac{AInit}{AState'}$	$\frac{CInit}{CState'}$
$a' = a0$	
$\frac{\tau_A}{\Delta AState}$	
$a = a0 \wedge a' \in \{a1, a2\}$	
$\frac{P_A}{\Delta AState; x! : \{1, 2\}}$	$\frac{P_C}{\Delta CState; x! : \{1, 2\}}$
$(a = a1 \wedge x! = 1) \vee (a = a2 \wedge x! = 2)$	
$a' = a0$	
$Q_A == P_A$	$Q_C == P_C$

The concrete data type can refuse $\{P!1, Q!2\}$ after the empty trace; this is not possible in the abstract data type. Thus, it is not a refinement, however, it satisfies all downward simulation conditions except the finalisation condition.¹²

The remaining finalisation condition looks cumbersome to check (quantifying over all sets of events that a concrete state might refuse); however two further simplifications are possible:

- only sets E that are maximal in the concrete state need to be considered; downward closedness of refusals in the linked abstract state then ensures that the property is also satisfied for subsets;
- only events refused because of the availability of alternative outputs need to be considered; operations whose precondition does not hold in the concrete state must be refused in any linked abstract state, due to the blocking condition.

Maximal refusals are most easily characterised by their complements, which select a single output value for each enabled operation (and refuse all other output values, whether possible or impossible), i.e., they are partial functions from I to $Output$. *Maxsim* is a schema operation, in this case effectively a predicate on states, parameterised by such a partial function (and implicitly by the state schema).

$$\begin{aligned} Sim & == I \leftrightarrow Output \\ Maxsim(E) & == \forall i : I \setminus \text{dom } E \bullet \neg \text{pre } Op_i \wedge \forall i : \text{dom } E \bullet \exists State'; Output \bullet Op_i \wedge \theta Output = E(i) \end{aligned}$$

Thus, in addition to the conditions from Section 5.5.1, we require here

¹² Taking a CSP view, and writing $(\tau \rightarrow p) \sqcap (\tau \rightarrow q)$ as $p \sqcap q$, this example is $(p!1 \sqcap q!1) \sqcap (p!2 \sqcap q!2) \not\sqsubseteq (p!1 \sqcap p!2) \sqcap (q!1 \sqcap q!2)$.

DS.FinDemBlock. τ

$$\begin{aligned} & \forall AState; CState; E \mid R \wedge CState \downarrow \wedge Maxsim(E) \bullet \\ & \quad \exists AState' \bullet \tau_A^* \wedge AState' \downarrow \wedge \forall i : \text{dom } E \bullet (\exists (AOp_i)' \bullet \theta Output' = E(i)) \vee \neg \exists (AOp_i)' \end{aligned}$$

where the consequent ensures that for each operation that is enabled in the concrete state, some stable abstract state connected by $R \circ \tau_A^*$ either selects the same output, or disables the operation (leading to a superset of refusals in either case).

In the counterexample above, this condition fails on the maximal refusal set $\{P!1, Q!2\}$ represented by its complement $\{P!2, Q!1\}$. There is no abstract state which, for each of these events, either allows it or completely disallows the operation (P or Q).

In Section 3.4 we discussed how, in the blocking model with input/output, we needed the condition **US.FinDem** to correctly link up refusal sets due to outputs. The example above has shown that, in the presence of internal operations, a similar extra condition is needed for downward simulations, and this we have called **DS.FinDemBlock. τ** .

6.2.2. Upward simulation

The reasoning here is similar to the corresponding case without outputs. Again, for a concrete non-divergent state c , we need to find a linked abstract state for each stable state c' reachable from c by internal evolution. However, as in this model such a state c' does not necessarily have a single maximal refusal set, we need to find a (possibly different) linked abstract state for each maximal refusal set in each such c' .

US.FinDemBlock. τ

$$\begin{aligned} & \forall E : Sim; \Delta CState \mid \neg CState \uparrow \wedge \tau_C^* \wedge CState' \downarrow \wedge (Maxsim(E))' \bullet \\ & \quad \exists AState, F : Sim \bullet T \wedge AState \downarrow \wedge F \subseteq E \wedge Maxsim(F) \end{aligned}$$

In the absence of internal operations this, of course, collapses to **US.FinDem**.

6.3. Outputs in the non-blocking approach

In the non-blocking approach with demonic outputs, it may at first seem surprising that there are refusals (other than those after divergence), even in states where operations are not enabled (and not blocked either!) First, we still expect output values to be determined “by the system”, and thus output values will still be refused if other values are possible. Moreover, consider an operation that is not enabled (in a particular state), and thus leads to divergence. If none of the possible output values of that operation is refused, it also means that defining the operation to be enabled and have a particular output (or any choice of outputs) will not be a valid refinement, as it introduces new refusals. Thus, where an operation is not enabled in a particular state, we will allow any set of output values to be refused, except for the full set. In particular, this means that an operation whose output is from a singleton set will produce no refusals. This is consistent with the view that having no output, or having an output from a singleton set are isomorphic.

The resulting definition of refusals is then given by modifying the predicate $Fcond$ in Definition 8 as follows.

$$Fin == \{State; os : \text{seq } Output; E : \mathbb{P}(I \times Output) \mid Fcond \bullet (os, \theta State) \mapsto (os, E)\}$$

where $Fcond$ is

$$E \subseteq \{(i, out) \mid (\exists State'; Output \bullet (\text{pre } Op_i \Rightarrow Op_i) \wedge \theta Output \neq out \wedge (i, \theta Output) \notin E)\} \quad (42)$$

Thus, the first disjunct from Definition 8 defining refusals when the operation is not applicable is dropped. The operation Op_i is replaced by (effectively a totalisation) $\text{pre } Op_i \Rightarrow Op_i$ which allows an arbitrary result outside the precondition.

6.3.1. Downward simulation

The example in Section 6.2.1 also shows why the rules from Theorem 11 do not suffice in the case of demonic output refusals. (It is not a refinement, but it satisfies all the conditions of Theorem 11. However, the finalisation condition fails.)

Assuming the other conditions, the finalisation condition in this case reduces to

$$\begin{aligned} & \forall a : AState; c, c' : CState; E \bullet (a, c) \in R \wedge (c, c') \text{ in } \tau_C^* \wedge c \notin CState \uparrow \wedge c' \in CState \downarrow \wedge Fcond_{c'} \\ & \Rightarrow \exists a' : AState \bullet (a, a') \in \tau_A^* \wedge a' \in AState \downarrow \wedge Fcond_{a'} \end{aligned}$$

Compared to the condition in Section 6.2.1, we require consideration of τ_C^* as well, as we cannot assume that R is closed under composition with τ_C^* . As usual, it suffices to consider only maximal refusal sets in particular states (here: c'), and these can be represented by their complements. These now select an output for *every* operation (enabled or not), and thus they are *total* functions. This is characterised by:

$$Maxsimtot(E) == \text{dom } E = I \wedge \forall i : I \bullet \text{pre } Op_i \Rightarrow \exists State'; Output \bullet Op_i \wedge \theta Output = E(i)$$

This leads to the following refinement condition.

DS.FinDemNonBlock. τ

$$\begin{aligned} & \forall AState; E; CState; CState' \mid R \wedge \neg CState \uparrow \wedge \tau_C^* \wedge CState' \downarrow \wedge (Maxsimtot(E))' \bullet \\ & \quad \exists AState' \bullet \tau_A^* \wedge AState' \downarrow \wedge (Maxsimtot(E))' \end{aligned}$$

As E is a total function from I to $Output$, it necessarily also represents a maximal set of simultaneously enabled events in A .

6.3.2. Upward simulation

For upward simulations we require the conditions given in Section 5.6.2 and, additionally, the conditions represented by the finalisation $\text{CFin} \subseteq \text{T} \text{ ; } \text{AFin}$. This leads to the requirement that

$$\begin{aligned} & \forall c, c' : CState; E \bullet c \notin CState \uparrow \wedge (c, c') \in \tau_C^* \wedge c' \in CState \downarrow \wedge Fcond_{c'} \\ & \Rightarrow \exists a, a' : AState \bullet (c, a) \in T \wedge (a, a') \in \tau_A^* \wedge a' \in AState \downarrow \wedge Fcond_{a'} \end{aligned}$$

where $Fcond$ is the predicate in (42) above. Concentrating on maximal refusal sets in c' and rephrasing that in terms of their complements gives:

$$\text{US.FinDemNonBlock.}\tau \quad \forall \Delta CState; E : Sim \mid \neg CState \uparrow \wedge \tau_C^* \wedge CState' \downarrow \wedge (Maxsimtot(E))' \bullet \exists \Delta AState \bullet TAState' \downarrow \wedge \tau_S^* \wedge (Maxsimtot(E))'$$

In this non-blocking model, refusals arise solely from the presence of outputs, since the internal choice of an observable aspect means refusals can occur if the environment was only prepared to accept a different value. To understand the requirement, consider the following scenario, in a simpler context without any internal operations.

In the first there is just one operation in C without input or output, and none in A (see Figure 5.). The non-blocking totalisation of A (given as a CSP process) is

$$Op \rightarrow \text{div}$$

Fig. 5. Refusals in the non-blocking model**Fig. 6.** Refusals in the non-blocking model - with outputs

In the initial state, the refusals of A and C are thus the same (i.e., none). This is consistent with the construction of E and $Fcond$, since in both A and C we cannot find different output values occurring, thus $E = \emptyset$. The finalisation condition thus does not impose any further constraint on upward simulations in a model without input or output, since there are no refusals.

However, now consider a model with input and output. Consider A and C with one operation that can potentially output o_1 or o_2 . Adapting the above example, we derive that in Figure 6. The non-blocking totalisation of A now has an internal choice of possible output values

$$(Op.o_1 \rightarrow \text{div}) \sqcap (Op.o_2 \rightarrow \text{div})$$

Thus refusals of A after the empty trace are $\{Op.o_1\}$ and $\{Op.o_2\}$, whereas those of C are just $\{Op.o_2\}$.

The presence of non-determinism in the specification, together with refusals due to outputs complicates the finalisation requirement. In particular, if an operation has a non-deterministic output at a given state, this can, in a refinement, be transferred to non-determinism prior to the operation invocation. This means the subsetting of refusals can be split across different states. For example, in Figure 7 the refusals initially in C are $\{Op.o_1\}$ and $\{Op.o_2\}$ but not $\{Op.o_1, Op.o_2\}$. The condition $\forall CState; E \bullet Fcond_C \Rightarrow \exists AState \bullet T \wedge Fcond_A$ allows the chosen $AState$ for $E = \{Op.o_1\}$ to be $s = 2$ but that for $E = \{Op.o_2\}$ to be $s = 1$.

7. Conclusions

In this paper we have formed a general framework for reasoning about relational and process refinement. To do so we started out with an abstract data type specified in the usual Z “states-and-operations” style [43] but with some additional syntax to make the ADT structure explicit [18]. We then defined a particular process semantics for ADTs, and embedded an ADT in a relational data type in one of several ways, depending on the notion of observation. This enabled us to show that relational refinement of the embedded ADT was equivalent to a refinement of the process semantics. As a consequence, simulation rules on that particular relational embedding can then be used to verify the process refinement relation.

As we mentioned in the introduction similar approaches have been applied previously:

- to derive the usual Z refinement where the “process” semantics is not a concurrent one at all, observing only output values, with the “precondition” interpretation of operations – i.e., applying an undefined operation leads to divergence [46, 18];
- to derive refinement for Z with the “guard” interpretation, i.e., applying an undefined operation is impossible [9, 18];
- to define failures refinement for Z, in both the “guard” and “precondition” interpretations, observing refusals and outputs [4, 19, 40];
- to define readiness refinement for Z, using a version of relational refinement for partial relations [19];
- to define failures-divergences refinement for Z ADTs, in both interpretations, allowing internal operations but not outputs [20].

Fig. 7. Non-determinism and refusals inclusion

As noted in [32, 3], the consideration of internal operations [20] once again brings home that the two different “erroneous” behaviours, often appearing as different interpretations of partiality in operations, can – and often need to – co-exist. The “guards” or “blocking” interpretation leads to *deadlock*, i.e., applying an operation outside its domain leads to a special state which represents that a deadlock has occurred. If that deadlock is the only possible outcome for a particular trace (or series of choice points), this means that the trace is not allowed to occur, and thus this deadlock may not be removed in refinement. The “pre-condition” (or “non-blocking”, or “contract”) interpretation introduces *divergence*: an operation applied outside its domain is not contracted to produce any particular result, and may therefore produce any result including deadlock, or “worse”. Typically one would take this “any result” view as underspecification, and allow for it to be removed in refinement. However, internal operations (at least in the CSP interpretation) bring in the possibility of divergence through livelock, independently of how partial operations are interpreted. In addition, we saw that outputs led to refusals even when we assume the non-blocking approach.

These considerations led to the definition of a relational datatype (which we called a *process data type*) with explicit characterisations of “normal” behaviour, divergence, and blocking, and observation of refusals at finalisation. Simulation rules (Sections 4.1 and 4.2) were derived for this generalised scheme “once and for all”. We then subsequently showed how process data types could be given a number of different instantiations. The key result is the instantiation given in Section 5, which showed how a failures-divergences interpretation could be given for process data types. The correctness of the embedding was proved and simulation conditions derived which highlighted the role of divergence.

The main line of development that leads to Theorems 6, 8 and 10 has been formally verified using the KIV theorem prover [35]. Related theories verified previously include our earlier work on weak refinement ([17], in [38]) and Cooper, Stepney and Woodcock’s generalised Z refinement rules ([15], in [39]). In contrast to these, formal theories and proofs were done during the development, and not a posteriori given the finished work. Mechanisation required three weeks of work: one for setting up specifications, one for proving theorems and one for adapting to changes. It uncovered a few missing assumptions (such as the stronger induction hypothesis for Lemma 2 using f' rather than f), and helped to shorten proofs, such as that the finalisation condition is implied by the blocking condition for Theorem 8. A Web presentation automatically generated from KIV specifications and proofs is available [37].

These results extend previous work in a number of directions, specifically divergence due to internal evolution and the inclusions of outputs. In the survey given in Section 3 we gave the simulation conditions necessary for the standard non-blocking model (which corresponds to traces-divergences refinement), the standard blocking model (which corresponds to singleton failures refinement), and for a refusals embedding in a blocking model which corresponds to failures-divergences refinement. Section 5.5 extended this latter model by adding in the potential of internal evolution (but without outputs). This enabled the systematic derivation of the extension of Josephs’ rules [30] to include arbitrary internal behaviour, which we believe has not been published previously. Section 5.6 gave the simulation conditions for a refusals embedding in the non-blocking approach (again without outputs). The full generality is reached in Sections 6.2 and 6.3 when the additional conditions are given when both outputs and internal evolution are included in the models. The inclusion of outputs is, of course, what makes the models subtle in terms of refusal information, and one direction of further research would be to attempt to further simplify condition **US.FinDem** and the various other conditions in terms of the complements of maximal refusal sets.

We have concentrated on the demonic interpretation of outputs, as the model that is, perhaps, more common, and presents us with more subtleties than the angelic. The angelic has, however, been advocated in the context of integrating formal methods. For example, Treharne and Schneider argue in [41] that the angelic model is more natural when CSP controllers are combined with B components in integrations of CSP and B. In particular, the CSP acts as a non-discriminating controller which cannot block on the values that the B component provides.

It is also worth noting that the angelic model can be more discriminating than the demonic model, e.g.,

$$(a \rightarrow b!1 \rightarrow stop) \sqcap (a \rightarrow (b!2 \rightarrow stop \sqcap b!3 \rightarrow stop))$$

is demonically equivalent to

$$(a \rightarrow b!1 \rightarrow stop) \sqcap (a \rightarrow b!2 \rightarrow stop) \sqcap (a \rightarrow b!3 \rightarrow stop)$$

However, these processes are not equivalent angelically. It is left to future work to explore the consequences of angelic outputs and how they might be used, for example, in integrations of formal methods in ways which were not possible for the demonic model.

Acknowledgements

We would like to thank Christie Bolton, Jim Davies, Steve Schneider, Helen Treharne and Heike Wehrheim for discussions on these issues, and the anonymous reviewers for their suggestions and comments. John Derrick was supported by the Leverhulme Trust via a Research Fellowship for this work.

References

- [1] J.-R. Abrial. *The B-Book: Assigning Programs to Meanings*. CUP, 1996.
- [2] R. Berghammer and H. Zierer. Relation algebraic semantics of deterministic and non-deterministic programs. *Theoretical Computer Science*, 43:123–147, 1986.
- [3] E. A. Boiten and W.-P. de Roever. Getting to the bottom of relational refinement: Relations and correctness, partial and total. In R. Berghammer and B. Möller, editors, *7th International Seminar on Relational Methods in Computer Science (RelMiCS 7)*, pages 82–88. University of Kiel, May 2003.
- [4] E. A. Boiten and J. Derrick. Unifying concurrent and relational refinement. *ENTCS*, 70, 2002. Proceedings REFINe’02, Editors: J. Derrick, E. A. Boiten, J. von Wright and J. C. P. Woodcock.
- [5] T. Bolognesi and E. Brinksma. Introduction to the ISO Specification Language LOTOS. *Computer Networks and ISDN Systems*, 14(1):25–59, 1988.
- [6] C. Bolton. *On the refinement of state-based and event-based models*. PhD thesis, University of Oxford, 2002.
- [7] C. Bolton and J. Davies. Refinement in Object-Z and CSP. In M. Butler, L. Petre, and K. Sere, editors, *Integrated Formal Methods (IFM 2002)*, volume 2335 of *Lecture Notes in Computer Science*, pages 225–244. Springer-Verlag, 2002.
- [8] C. Bolton and J. Davies. A singleton failures semantics for Communicating Sequential Processes. *Formal Aspects of Computing*, 18:181–210, 2006.
- [9] C. Bolton, J. Davies, and J. C. P. Woodcock. On the refinement and simulation of data types and processes. In K. Araki, A. Galloway, and K. Taguchi, editors, *International Conference on Integrated Formal Methods 1999 (IFM’99)*, pages 273–292, York, July 1999. Springer.
- [10] C. Bolton and G. Lowe. A hierarchy of failures-based models. In F. Corradini and U. Westmann, editors, *Proceedings of Express 2003: 10th International Workshop on Expressiveness in Concurrency*. Elsevier Science Publishers, 2003.
- [11] C. Bolton and G. Lowe. A hierarchy of failures-based models: Theory and application. *Theoretical Computer Science*, 330(3):407–438, 2005.
- [12] Howard Bowman and Rodolfo Gomez. *Concurrency Theory: Calculi and Automata for Modelling Untimed and Timed Concurrent Systems*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.
- [13] S. D. Brookes, C. A. R. Hoare, and A. W. Roscoe. A theory of communicating sequential processes. *Journal of the ACM*, 31(3):560–599, 1984.
- [14] S.D. Brookes and A.W. Roscoe. An improved failures model for communicating processes. In S.D. Brookes, A.W. Roscoe, and G. Winskel, editors, *Seminar on Concurrency*, volume 197 of *Lecture Notes in Computer Science*, pages 281–305. Springer, 1985.
- [15] D. Cooper, S. Stepney, and J. Woodcock. Derivation of Z refinement proof rules: forwards and backwards rules incorporating input/output refinement. Technical Report YCS-2002-347, University of York, December 2002. URL: <http://www-users.cs.york.ac.uk/~susan/bib/ss/z/zrules.htm>.
- [16] W.-P. de Roever and K. Engelhardt. *Data Refinement: Model-Oriented Proof Methods and their Comparison*. CUP, 1998.
- [17] J. Derrick, E. A. Boiten, H. Bowman, and M. W. A. Steen. Specifying and refining internal operations in Z. *Formal Aspects of Computing*, 10:125–159, December 1998.
- [18] J. Derrick and E.A. Boiten. *Refinement in Z and Object-Z: Foundations and Advanced Applications*. FACIT. Springer Verlag, May 2001.
- [19] J. Derrick and E.A. Boiten. Relational concurrent refinement. *Formal Aspects of Computing*, 15(1):182–214, November 2003.
- [20] J. Derrick and E.A. Boiten. Relational concurrent refinement with internal operations. In B. Aichernig, E.A.

- Boiten, J. Derrick, and L. Groves, editors, *BCS-FACS Refinement Workshop*, volume 187 of *ENTCS*, pages 35–53, 2006.
- [21] M. Deutsch and M.C. Henson. An analysis of refinement in an abortive paradigm. *Formal Aspects of Computing*, 18(3):329–363, 2006.
- [22] H. Doornbos. A relational model of programs without the restriction to Egli-Milner constructs. In E.-R. Olderog, editor, *PROCOMET '94*, pages 357–376. IFIP, 1994.
- [23] Steve Dunne and Stacey Conroy. Process refinement in B. In Helen Treharne, Steve King, Martin C. Henson, and Steve Schneider, editors, *ZB 2005: Formal Specification and Development in Z and B, 4th International Conference of B and Z Users*, volume 3455 of *Lecture Notes in Computer Science*, pages 45–64. Springer, 2005.
- [24] C. Fischer. CSP-OZ - A combination of CSP and Object-Z. In H. Bowman and J. Derrick, editors, *Second IFIP International Conference on Formal Methods for Open Object-Based Distributed Systems*, pages 423–438. Chapman & Hall, July 1997.
- [25] He Jifeng, C. A. R. Hoare, and J. W. Sanders. Data refinement refined. In B. Robinet and R. Wilhelm, editors, *Proc. ESOP'86*, volume 213 of *Lecture Notes in Computer Science*, pages 187–196. Springer-Verlag, 1986.
- [26] Matthew Hennessy and Anna Ingólfssdóttir. A theory of communicating processes with value passing. *Inf. Comput.*, 107(2):202–236, 1993.
- [27] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
- [28] C.A.R. Hoare and He Jifeng. *Unifying Theories of Programming*. Prentice-Hall, 1998.
- [29] He Jifeng. Process refinement. In J. McDermid, editor, *The Theory and Practice of Refinement*. Butterworths, 1989.
- [30] M. B. Josephs. A state-based approach to communicating processes. *Distributed Computing*, 3:9–18, 1988.
- [31] G. Leduc. *On the Role of Implementation Relations in the Design of Distributed Systems using LOTOS*. PhD thesis, University of Liège, Liège, Belgium, June 1991.
- [32] R. Miarka, E. A. Boiten, and J. Derrick. Guards, preconditions and refinement in Z. In J. P. Bowen, S. Dunne, A. Galloway, and S. King, editors, *ZB2000: Formal Specification and Development in Z and B*, volume 1878 of *Lecture Notes in Computer Science*, pages 286–303. Springer-Verlag, September 2000.
- [33] R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
- [34] S. Reeves and D. Streader. State- and event-based refinement. Technical report, Department of Computer Science, University of Waikato, September 2006.
- [35] W. Reif, G. Schellhorn, K. Stenzel, and M. Balsler. Structured specifications and interactive proofs with KIV. In W. Bibel and P. Schmitt, editors, *Automated Deduction—A Basis for Applications*, volume II: Systems and Implementation Techniques, chapter 1: Interactive Theorem Proving, pages 13 – 39. Kluwer Academic Publishers, Dordrecht, 1998.
- [36] A.W. Roscoe. *The Theory and Practice of Concurrency*. International Series in Computer Science. Prentice Hall, 1998.
- [37] G. Schellhorn. Web presentation of the KIV proofs of ‘Relational Concurrent Refinement Part II: Internal Operations and Output’, 2006. URL: <http://www.informatik.uni-augsburg.de/swt/projects/Refinement/Web/CSPRef/>.
- [38] G. Schellhorn. ASM refinement and generalizations of forward simulation in data refinement: A comparison. *Theoretical Computer Science*, 336(2-3):403–435, May 2005.
- [39] G. Schellhorn, H. Grandy, D. Haneberg, and W. Reif. The Mondex challenge: Machine checked proofs for an electronic purse. In J. Misra, T. Nipkow, and E. Sekerinski, editors, *Formal Methods 2006, Proceedings*, volume 4085 of *LNCS*, pages 16–31. Springer, 2006.
- [40] S. Schneider. Non-blocking data refinement and traces-divergences semantics, 2006. Personal communication.
- [41] S. Schneider and H. Treharne. CSP Theorems for Communicating B Machines. *Formal Aspects of Computing*, 17(4):390 – 422, December 2005.
- [42] G. Smith and J. Derrick. Abstract specification in Object-Z and CSP. In C. George and H. Miao, editors, *Formal Methods and Software Engineering*, volume 2495 of *Lecture Notes in Computer Science*, pages 108–119. Springer, November 2002.
- [43] J. M. Spivey. *The Z Notation: A Reference Manual*. International Series in Computer Science. Prentice Hall, 2nd edition, 1992.
- [44] Antti Valmari and Martti Tienari. Compositional failure-based semantics models for basic LOTOS. *Formal Aspects of Computing*, 7(4):440–468, 1995.
- [45] R. J. van Glabbeek. The linear time - branching time spectrum I. The semantics of concrete sequential processes. In J.A. Bergstra, A. Ponse, and S.A. Smolka, editors, *Handbook of Process Algebra*, pages 3–99. North-Holland, 2001.
- [46] J. C. P. Woodcock and J. Davies. *Using Z: Specification, Refinement, and Proof*. Prentice Hall, 1996.
- [47] J. C. P. Woodcock and C. C. Morgan. Refinement of state-based concurrent systems. In D. Bjorner, C. A. R. Hoare, and H. Langmaack, editors, *VDM'90: VDM and Z!- Formal Methods in Software Development*, volume 428 of *Lecture Notes in Computer Science*. Springer-Verlag, 1990.