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.

E-mail: E.A.Boiten@kent.ac.uk

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

E-mail: J.Derrick@dcs.shef.ac.uk

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

E-mail: schellhorn@informatik.uni-augsburg.de

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. Third, 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. In state-based notations, such as Z [Spi92], 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 [Hoa85, Ros98], 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 [Ros98]. In CCS, bisimulation is typically used [Mil89], whereas in LOTOS reduction, extension and conformance are defined [BoB88]. A survey of many of the most prominent refinement relations is given in [vGl01].

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 [dRE98]. Two different forms of simulations are needed to provide a complete method: downward and upward simulations. Theoretical background is given in [dRE98], and examples of their use in Z are given in [WoD96, DeB01].

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 threefold. 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 by-product 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 [Jos88], He [He89], Woodcock and Morgan [WoM90], Bolton and Davies [BoD02b, BoD06], Derrick and Boiten [BoD02a, DeB03] and Schneider (2006, unpublished). That due to Josephs [Jos88], He [He89], Woodcock and Morgan [WoM90] defines a basic correspondence between simulation rules and failures–divergences refinement. The more recent work of Bolton and Davies [BoD02b, BoD06], Derrick and Boiten [BoD02a, DeB03] and Schneider 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, and Bolton and Davies [BoD02b, BoD06] do the former, whilst Derrick and Boiten choose the latter [DeB03]. 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. 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 [DeB03] show 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 [DeB06] 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 Sect. 2 we provide the basic definitions and background, and we survey existing work in Sect. 3. The new relational datatype is defined in Sect. 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 Sect. 7 which also discusses the formalisation of the theory in KIV.

2. Background

In this section we present the standard refinement theory [DeB01] for ADTs in a relational setting. We describe the relational model for ADTs written as total relations over a global state in Sect. 2.1 and that for partial relations in Sect. 2.2. Section 2.3 introduces process algebraic variants of refinement that are relevant to our discussion. Section 2.4 shows how the relational theory can be applied to a specification language such as Z, and Sect. 2.5 studies the global state in more detail.

2.1. The standard relational model

Relational ADTs 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 relation from **State** to **G**. If all operations $\{\text{Op}_i\}$ and **Fin** 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., $\text{Init} = \text{G} \times \text{ran Init}$. $\square$

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 $\text{D} = (\text{State}, \text{Init}, \{\text{Op}_i\}_{i \in I}, \text{Fin})$ 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 $\text{Init} \circ \text{p} \circ \text{Fin}$. For example, if $\text{p} = \langle \text{p}_1, \dots, \text{p}_n \rangle$ then $\text{p}_\text{D} = \text{Init} \circ \text{Op}_{\text{p}_1} \circ \dots \circ \text{Op}_{\text{p}_n} \circ \text{Fin}$.

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

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

[HHS86, dRE98]. 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 $A = (AState, AInit, \{AOp_i\}_{i \in I}, AFin)$ and $C = (CState, CInit, \{COp_i\}_{i \in I}, CFin)$. A *downward* simulation is a relation R between $AState$ and $CState$ satisfying

$$\begin{aligned} CInit &\subseteq AInit \circ R \\ R \circ CFin &\subseteq AFin \\ \forall i : I \bullet R \circ COp_i &\subseteq AOp_i \circ R \end{aligned}$$

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

$$\begin{aligned} CInit \circ T &\subseteq AInit \\ CFin &\subseteq T \circ AFin \\ \forall i : I \bullet COp_i \circ T &\subseteq T \circ AOp_i \end{aligned}$$

$\square$

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 $State_\perp = State \cup \{\perp\}$ for some $\perp \notin State$. Two ways of defining a totalisation are the following.¹

Definition 5 (Blocking and non-blocking totalisation)

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

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

or in the blocking model by

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

$\square$

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 [DeB01]—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.

¹ There are others, which are discussed in [DeH06] and in terms of the terminology defined there the non-blocking totalisation we give is an *unstrict lifting*, and the blocking one is a *strict lifting*. Deutsch and Henson provide a careful examination of the possible ways to totalise and embed the element $\perp$ into a total data type.

Here we will consider semantic models of CSP, of which there are a number, 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, [Led91] and [VaT95] and the comparison given in [BoG05] 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 [BHR84, BrR85, Ros98]. 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

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

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.

$$\begin{aligned} (\langle \rangle, \emptyset) &\in \mathcal{F} & (F1) \\ (t_1 \frown t_2, \emptyset) &\in \mathcal{F} \Rightarrow (t_1, \emptyset) \in \mathcal{F} & (F2) \\ (t, X) \in \mathcal{F} \wedge Y \subseteq X &\Rightarrow (t, Y) \in \mathcal{F} & (F3) \\ (t, X) \in \mathcal{F} \wedge (\forall e \in Y \bullet (t \frown \langle e \rangle, \emptyset) \notin \mathcal{F}) &\Rightarrow (t, X \cup Y) \in \mathcal{F} & (F4) \end{aligned}$$

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

$$\begin{aligned} \mathcal{D} &\subseteq \text{dom } \mathcal{F} & (D1) \\ t_1 \in \mathcal{D} \wedge t_2 \in A^* &\Rightarrow t_1 \frown t_2 \in \mathcal{D} & (D2) \\ t \in \mathcal{D} \wedge X \subseteq A &\Rightarrow (t, X) \in \mathcal{F} & (D3) \end{aligned}$$

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 [BrR85]. 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 [Bol02] (and published in [BoD02b, BoD06]) 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).

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

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 (i.e., ones expressed in terms of $stop$, $\rightarrow$, $\sqcap$, $\square$ 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 [vGl01] and later in [BoL05] (which builds on [BoL03]).

2.4. Refinement in Z

In Sect. 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 [WoD96, DeB01] 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 (downward) retrieve relation between $AState$ and $CState$, this is modelled as

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

and analogously for upward simulation.

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

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

□

2.5.2. Refusals

We wish to observe traces and refusals.⁴ 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 (State , Init , $\{\text{Op}_i\}_{i \in I}$, Fin) with global state $\mathbf{G}$, a particular notion of events $\mathbf{E}$, and a refusal relation $\text{Ref} : \text{State} \leftrightarrow \mathbb{P} \mathbf{E}$. The ADT is said to be a *refusal embedding* (for Ref) iff it is program controlled, and a type $\mathbf{GB}$ exists such that

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

□

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 [vGI01] and [DuC05].

As an example, in the context of a Z data type with no input and output, events in E are simply operation names from I , and refusals will be any subset E of the maximal refusals in a given state: $\{i : I \mid \neg \text{pre } Op_i\}$. Thus an ADT $(State, Init, \{Op_i\}_{i \in I})$ in the refusals interpretation is embedded in the relational model as follows. The global state 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 Sect. 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 [DeB03]).

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 [SmD02]).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 [HeI93] 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 ADT $(State, Init, \{Op_i\}_{i \in I})$ with outputs of type $Output$ is embedded in the relational model as follows. The global state is defined by

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

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

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

with in the angelic model $Fcond = FcondA(\theta State, E)$ where

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

and in the demonic model $Fcond = FcondD(\theta State, E)$ where:

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

□

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
Blocking data refinement with deterministic outputs	Singleton failures	Bolton and Davies
Blocking data refinement	Singleton failures of process and input process	Bolton and Davies
Blocking data refinement with strengthened applicability but no input/output	Failures	Josephs
Blocking data refinement with extended finalisations	Failures–divergences	Derrick and Boiten
Non-blocking data refinement with extended finalisations but no input/output	Failures–divergences	Derrick and Boiten

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 Sect. 3.3 below.

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

Inspired by the work by Bolton and Davies [BoD02b, BoD06] discussed below, Schneider (2006, unpublished) 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 Sect. 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 $process(A) \sqsubseteq_{id} process(C)$. $\square$

3.2. Blocking data refinement and the singleton failures semantics

Bolton in [Bol02] and Bolton and Davies in [BoD02b, BoD06] discuss the relationship between data refinement and the singleton failures semantics [vGl01] 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):

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

$$P_A(s) ==$$

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

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:

$$\text{inputProcess}_b(A) == \sqcap s \in \text{State}, (*, s) \in \text{Init} \bullet P_A(s)$$

$$P_A(s) == \sqcap i \in I, in \in \text{Input}, ((in), \langle \rangle, s) \in \text{dom } AOp_i \bullet \\ \sqcap s' \in \text{State} \mid (\exists out \in \text{Output} \mid ((in), \langle \rangle, s) \mapsto (\langle \rangle, \langle out \rangle, s') \in AOp_i) \bullet AOp_i.in \rightarrow P_A(s')$$

As can be seen, this is the same as *process_b* except that the outputs are unobservable. [BoD06] 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, $A \sqsubseteq_{\text{data}} C$ if and only if $\text{process}_b(A) \sqsubseteq_{\text{sf}} \text{process}_b(C)$ and $\text{inputProcess}_b(A) \sqsubseteq_{\text{sf}} \text{inputProcess}_b(C)$. For ADTs with deterministic outputs (or no input/output), this reduces to checking $\text{process}_b(A) \sqsubseteq_{\text{sf}} \text{process}_b(C)$. $\square$

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

$$\text{process}_b(A) \sqsubseteq_{\text{fd}} \text{process}_b(C) \Rightarrow A \sqsubseteq_{\text{data}} C$$

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

$$\text{process}_{nb}(A) == \sqcap s \in \text{State}, (*, s) \in \text{Init} \bullet Q_A(s)$$

$$Q_A(s) == \begin{array}{l} P_A(s) \\ \square \\ ((\sqcap i \in I, in \in \text{Input}, ((in), \langle \rangle, s) \notin \text{dom } AOp_i \bullet \\ \sqcap s' \in \text{State}, out \in \text{Output} \bullet AOp_i.in.out \rightarrow \text{Chaos}) \\ \sqcap \text{stop}) \end{array}$$

where *Chaos* == $(\sqcap i \in I, in \in \text{Input}, out \in \text{Output} \bullet AOp_i.in.out \rightarrow \text{Chaos}) \sqcap \text{stop}$ 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 [ReS06] 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.

Specifically, they discuss the difference between the embedded state $(\text{State} \times \text{seq } \text{Output})_{\perp}$ that we use (following [WoD96]), and $(\text{State}_{\perp} \times \text{seq } \text{Output})$ which is used by Bolton and Davies [BoD06]. 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 Fig. 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 $(\text{State}_{\perp} \times \text{seq } \text{Output})$ as the appropriate embedding.

Fig. 1. Singleton failures vs. data refinement [ReS06]

3.3. Defining a correspondence with failures–divergences refinement

Subsequent to the work of Bolton and Davies [BoD02a], Derrick and Boiten [DeB03] 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 Sect. 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 [Jos88] 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 [Jos88] are consistent with those in [DeB03].

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:

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

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

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

The standard upward simulation applicability condition requires that we consider pairs of abstract and concrete states for each operation. The finalisation condition, on the other hand, requires that for every abstract state we find a *single* concrete state such that the precondition of each abstract operation implies the precondition of its concrete counterpart.

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 outputs, the requirement that $CFin \subseteq T \mathbin{\text{\textcircled{S}}} AFin$, simplifies to:

$$\forall CState; E \bullet FcondD(\theta CState, E) \Rightarrow \exists AState \bullet T \wedge FcondD(\theta AState, E) \quad (2)$$

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 Sect. 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 is 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 $\text{Pre } Op == \exists State' \bullet Op$, and then using 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 $\forall CState' \bullet CInit \Rightarrow \exists AState' \bullet AInit \wedge R'$

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

DS.CorrNonBlock

$\forall i : I; Output; CState'; CState; AState \bullet \text{pre } AOp_i \wedge R \wedge COp_i \Rightarrow \exists AState' \bullet R' \wedge AOp_i$

DS.CorrBlock $\forall i : I; Output; CState'; CState; AState \bullet R \wedge COp_i \Rightarrow \exists AState' \bullet R' \wedge AOp_i$

DS.FinAng $\forall CState; AState; i : I; Output \bullet R \wedge \text{Pre } AOp_i \Rightarrow \text{Pre } COp_i$

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 $CState$ plus the following set:

US.Init $\forall CState'; AState' \bullet T' \wedge CInit \Rightarrow AInit$

US.AppBlock $\forall i : I; Output \bullet \forall CState \bullet \exists AState \bullet T \wedge \text{pre } AOp_i \Rightarrow \text{pre } COp_i$

US.CorrNonBlock

$\forall i : I; Output; AState'; CState'; CState \bullet T' \wedge COp_i \Rightarrow \exists AState \bullet T \wedge (\text{pre } AOp_i \Rightarrow AOp_i)$

US.CorrBlock $\forall i : I; Output; AState'; CState'; CState \bullet T' \wedge COp_i \Rightarrow \exists AState \bullet T \wedge AOp_i$

US.FinRef $\forall CState \bullet \exists AState \bullet T \wedge \forall i : I \bullet \text{pre } AOp_i \Rightarrow \text{pre } COp_i$

US.FinDem $\forall CState; E \bullet FcondD(\theta CState, E) \Rightarrow \exists AState \bullet T \wedge FcondD(\theta AState, E)$

US.FinAng $\forall CState \bullet \exists AState \bullet T \wedge \forall i : I; Output \bullet \text{Pre } AOp_i \Rightarrow \text{Pre } COp_i$

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

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 Fig. 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 e it is not the case that any abstract state has

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

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

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 $(\langle B \rangle, \{TF, EF\})$ which is not present in C , under a singleton failures model in A we just obtain singleton failures $(\langle B \rangle, \{TF\})$, $(\langle B \rangle, \{EF\})$, ... thus the difference is not observable. To recover failures-divergences refinement in the blocking model one needs 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 Fig. 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, b has as its refusals $\mathbb{P}\{E!true, T!true\}$, whereas d has as its refusals $\mathbb{P}\{E!true, T!false\}$ and e has $\mathbb{P}\{E!false, T!true\}$.

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 [DeB03], and Sect. 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 b , c , d and e has operations E and T 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.

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

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 $FcondD(A, E)$ holds. Now for this AS consider any operation Op , if $pre AOp$ holds in AS but $pre COp$ 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 [MBD00, BdR03] that the two approaches are not exclusive—and indeed, formalisms such as B [Abr96] 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 [DeB06] 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 [dRE98, Chaps. 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.

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.

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 $(\text{State}, G \times \text{Inits}, \{N_i\}_{i \in I}, \text{Fin})$ $\square$

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;
- 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 [BeZ86, Doo94, Fis97, HoH98]. 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 $(\text{State}_{\perp, \omega}, \text{Init}, \{\llbracket \text{Op}_i \rrbracket\}_{i \in I}, \llbracket \text{Fin} \rrbracket)$ where

$$\text{Init} == G_{\perp, \omega, \text{no}} \times \text{Inits}$$

$$\llbracket (N, B, D) \rrbracket == N \cup (B_{\perp, \omega} \times \{\perp\}) \cup (D_{\omega} \times \text{State}_{\omega})$$

$$\llbracket \text{Fin} \rrbracket == \text{Fin} \cup (\overline{\text{dom Fin}} \times \{\text{no}\}) \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_{\omega, \text{no}}$$

$\square$

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.

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

Refinement (i.e., downward and upward simulation) of process data types is derived by embedding them into total data types,⁷ 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.

4.1. Downward simulation for process data types

Consider the process data types $(A\text{State}, A\text{lnits}, \{AOp_i\}_{i \in I}, A\text{Fin})$ and $(C\text{State}, C\text{lnits}, \{COp_i\}_{i \in I}, C\text{Fin})$, both with global state G , and a candidate downward simulation relation R between $A\text{State}$ and $C\text{State}$. We also embed the simulation relation, in order to relate abstract and concrete blocking and divergence correctly:

$$\llbracket R \rrbracket == R \cup \{(\perp, \perp)\} \cup \{\omega\} \times C\text{State}_{\omega}$$

Initialisation

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

$$\begin{aligned} C\text{lnit} &\subseteq A\text{lnit} \circ \llbracket R \rrbracket \\ &\equiv \{ \text{definition of embeddings} \} \\ G_{\perp, \omega} \times C\text{lnits} &\subseteq (G_{\perp, \omega} \times A\text{lnits}) \circ \llbracket R \rrbracket \\ &\equiv \{ \text{definition of } \llbracket R \rrbracket \} \\ C\text{lnits} &\subseteq \text{ran}(A\text{lnits} \triangleleft R) \end{aligned}$$

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.

$$\begin{aligned} \llbracket R \rrbracket \circ \llbracket C\text{Fin} \rrbracket &\subseteq \llbracket A\text{Fin} \rrbracket \\ &\equiv \{ \text{embeddings} \} \\ (R \cup \{(\perp, \perp)\} \cup \{\omega\} \times C\text{State}_{\omega}) \circ (C\text{Fin} \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_{\omega} \cup (\overline{\text{dom } C\text{Fin}} \times \{\text{no}\})) \\ &\subseteq \\ A\text{Fin} \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_{\omega} \cup (\overline{\text{dom } A\text{Fin}} \times \{\text{no}\}) \\ &\equiv \{ \text{distribution of } \cup \text{ over } \circ \} \\ R \circ C\text{Fin} \cup R \circ (\overline{\text{dom } C\text{Fin}} \times \{\text{no}\}) \cup R \circ \{(\perp, \perp)\} \cup R \circ (\{\omega\} \times G_{\omega}) \cup \{(\perp, \perp)\} \circ C\text{Fin} \\ &\cup \{(\perp, \perp)\} \circ (\overline{\text{dom } C\text{Fin}} \times \{\text{no}\}) \cup \{(\perp, \perp)\} \circ \{(\perp, \perp)\} \cup \{(\perp, \perp)\} \circ (\{\omega\} \times G_{\omega}) \\ &\cup (\{\omega\} \times C\text{State}_{\omega}) \circ C\text{Fin} \cup (\{\omega\} \times C\text{State}_{\omega}) \circ (\overline{\text{dom } C\text{Fin}} \times \{\text{no}\}) \cup (\{\omega\} \times C\text{State}_{\omega}) \circ \{(\perp, \perp)\} \\ &\cup (\{\omega\} \times C\text{State}_{\omega}) \circ (\{\omega\} \times G_{\omega}) \\ &\subseteq \\ A\text{Fin} \cup (\overline{\text{dom } A\text{Fin}} \times \{\text{no}\}) \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_{\omega} \\ &\equiv \{ \perp, \omega, \text{no} \notin \text{domain or range of finalisation or } R \} \end{aligned}$$

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.

$$\begin{aligned}
& R \circ \text{CFin} \cup (\text{dom}(R \triangleright (\text{dom CFin})) \times \{\text{no}\}) \cup \emptyset \cup \emptyset \cup \emptyset \cup \emptyset \cup \{(\perp, \perp)\} \cup \emptyset \cup (\{\omega\} \times \text{ran CFin}) \\
& \cup \{(\omega, \text{no})\} \cup \emptyset \cup \{\omega\} \times G_\omega \\
& \subseteq \\
& \text{AFin} \cup (\overline{\text{dom AFin}} \times \{\text{no}\}) \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_\omega \\
& \equiv \{ \text{calculus} \} \\
& R \circ \text{CFin} \subseteq \text{AFin} \quad \wedge \quad \text{dom}(R \triangleright (\text{dom CFin})) \subseteq \overline{\text{dom AFin}} \\
& \equiv \{ \text{calculus, taking domains of first conjunct to make second symmetric} \} \\
& R \circ \text{CFin} \subseteq \text{AFin} \quad \wedge \quad (\text{dom AFin}) \triangleleft R = R \triangleright (\text{dom CFin})
\end{aligned}$$

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 $\text{Op} = (N, B, D)$ 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 $\text{AOp} = (AN, AB, AD)$ and $\text{COp} = (CN, CB, CD)$, we have that

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

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 $\text{AOp}_i = (AN, AB, AD)$, $\text{COp}_i = (CN, CB, CD)$

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

□

The rules for initialisation and finalisation are identical to the usual rules.⁸ 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 $(B, D) = (\emptyset, \overline{\text{dom Op}_i})$ (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 $(A\text{State}, \text{Alnits}, \{\text{AOp}_i\}_{i \in I}, \text{AFin})$ and $(C\text{State}, \text{Clnits}, \{\text{COp}_i\}_{i \in I}, \text{CFin})$, both with global state G , and a candidate downward simulation relation T between $C\text{State}$ and $A\text{State}$, embedded as follows:

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

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:

$$\begin{aligned} \llbracket \text{CFin} \rrbracket &\subseteq \llbracket T \rrbracket \circ \llbracket \text{AFin} \rrbracket \\ &\equiv \{ \text{embeddings} \} \\ \text{CFin} \cup (\overline{\text{dom CFin}} \times \{\text{no}\}) \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_\omega \\ &\subseteq \\ (T \cup \{(\perp, \perp)\} \cup \{\omega\} \times A\text{State}_\omega) \circ (\text{AFin} \cup (\overline{\text{dom AFin}} \times \{\text{no}\}) \cup \{(\perp, \perp)\} \cup \{\omega\} \times G_\omega) \\ &\equiv \{ \text{domains; calculus (see also [DeB01, page 80])} \} \\ \text{CFin} &\subseteq T \circ \text{AFin} \wedge \overline{\text{dom CFin}} \subseteq \text{dom}(T \triangleright \text{dom AFin}) \end{aligned}$$

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 $C\text{State}$ 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

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

These derivations establish the following.

Theorem 5 (Upward simulation for process data types)

The relation T between $C\text{State}$ and $A\text{State}$ is an upward simulation between the process data types

⁸ The applicability rule for finalisation is less common, but is also included in [DeB01].

$(AState, Alnits, \{AOp_i\}_{i \in I}, AFin)$ and $(CState, Clnits, \{COp_i\}_{i \in I}, CFin)$, iff

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

and $\forall i : I$, for $AOp_i = (AN, AB, AD)$, $COp_i = (CN, CB, CD)$

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

□

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 [DeB03]) may still end up dominating some of the other conditions on operations. The property (see [DeB01]) 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 [DeB06] 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 $(\text{State}, \text{Init}, \{Op_i\}_{i \in I}, \tau, \text{Fin})$ such that $(\text{State}, \text{Init}, \{Op_i\}_{i \in I}, \text{Fin})$ is a basic data type, and the internal operation τ is a relation on State . □

5.2. Embedding into process data types

The addition of internal operations introduces a large number of, sometimes subtle, issues (see also [DBB98, DeB01]):

- (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 any 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).

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 $\lambda R. \bullet \text{id}_{\text{State}} \cup \tau \circ R$.

$\tau^{*|} == \tau^* \triangleright (\text{dom } \tau)$ 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.

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

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 **State** $\uparrow$, the ADT as a whole is divergent.

$\tau == (\text{State} \uparrow) \triangleleft \tau$ 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 $\tau^* = \text{id}_{\text{State} \uparrow} \cup (\text{State} \uparrow) \triangleleft \tau^*$.

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

$$\widetilde{\text{Init}} == \text{ran}(\text{Init} \circ \tau^*)$$

$$\widetilde{\text{Fin}} == \text{State} \uparrow \triangleleft \tau^{*|} \circ \text{Fin}$$

$$\widetilde{\text{Op}}_i == (\text{liv Op}_i \triangleleft \text{Op}_i, B_i, D_i)$$

where in the blocking approach

$$B_i == \text{dom } \text{Op}_i$$

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

and in the non-blocking approach

$$B_i == \emptyset$$

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

□

We argue informally that this embedding addresses each of the issues listed on page 84, with the exception of issue (c) concerning possible divergence of initial states:

- (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 [DBB98, DeB01] 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 only.¹⁰
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. □

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).

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 labelled 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 notation)

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

⁹ The corresponding encoding in [DeB06] 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 $\tau^* \mid \S \text{Fin} = \tau^* \triangleright (\text{dom } \tau) \S \text{Fin} = \tau^* \S (\text{dom } \tau) \triangleleft \text{Fin}$, and τ^* is always absorbed by preceding τ^* .

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

and $\Rightarrow$ 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

$$x \xrightarrow{() } x' == x = x' \qquad x \xrightarrow{p \widehat{ } q} x' == \exists x'' \bullet x \xrightarrow{p} x'' \wedge x'' \xrightarrow{q} x'$$

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 $T = (\text{State}, G \times \text{Inits}, \{\text{Op}_i\}_{i \in I}, \tau, \text{Fin})$ its divergences are

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

and its failures are

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

□

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 $\text{Init} \equiv G \times \text{ran Init}$;
- 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 (Definition 12), and $\llbracket x \rrbracket$ is the embedding of x from process data types into basic data types (Definition 10).¹¹

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

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

Proof. A simple proof shows that

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

Because T is not initially divergent, p is non-empty. We now prove

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

¹¹ This is different from the precursor paper [DeB06] where the various embeddings directly into basic data types were defined as $\hat{x}^t$ with varying t , and $\hat{x}$ was used for embedding simulations.

by induction over the construction of p .

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

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

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

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 $p \notin \text{div}(T)$, we have that

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

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 ran Init . The generalisation is proved by induction on p .

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

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

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

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

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 [DeB03] 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 $\text{div}(T)$, then

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

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 $p_{\llbracket \widetilde{T} \rrbracket}$ 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.

$$\text{CInit} \circ \tau_C^* \subseteq \text{AInit} \circ \tau_A^* \circ R \quad (3)$$

$$R \circ \text{CState} \uparrow \triangleleft \tau_C^* \circ \text{CFin} \subseteq \text{AState} \uparrow \triangleleft \tau_A^* \circ \text{AFin} \quad (4)$$

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

and for matching operations AOp and COp:

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

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

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

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 τ_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 $R \circ \tau_C^*$ is also a downward simulation between A and C .

Proof. By showing that (3) to (8) above imply the same conditions with $R \circ \tau_C^*$ 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 $R = R \circ \tau_C^*$. The initialisation condition (3) can then be simplified to $\text{CInit} \subseteq \text{AInit} \circ \tau_A^* \circ R$. The “correctness” condition (6) can be simplified using the divergence condition (7) to

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

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:

$$R \circ \text{CState} \uparrow \triangleleft \tau_C^* \circ \text{CFin} \subseteq \text{AState} \uparrow \triangleleft \tau_A^* \circ \text{AFin}$$

$$\equiv \{ \text{definition of } \subseteq \}$$

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

$$\Rightarrow a \notin \text{AState} \uparrow \wedge \exists a' : \text{AState} \bullet (a, a') \in \tau_A^* \wedge (a', E) \in \text{Ref}$$

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

$$\forall a : \text{AState}; E : \mathbb{P} I; c, c' : \text{CState} \bullet (a, c) \in R \wedge c \notin \text{CState} \uparrow \wedge (c, c') \in \tau_C^* \wedge (c', E) \in \text{Ref}$$

$$\Rightarrow a \notin \text{AState} \uparrow \wedge \exists a' : \text{AState} \bullet (a, a') \in \tau_A^* \wedge (a', E) \in \text{Ref}$$

We prove this by contradiction. Assume the implicand holds, but not the consequent. Thus, $(a', E) \notin \text{Ref}_A$ so for some $i \in E$ we have that $a' \in \text{dom } \overleftrightarrow{\text{AOp}}_i$. Theorem 7 implies that $(a, c') \in R$; from the blocking condition (8)

it then follows that $c' \in \text{dom } \overleftrightarrow{\text{COp}}_i$, and because c' is stable, also $c' \in \text{dom } \text{COp}_i$, contradicting $(c', E) \in \text{Ref}_C$. 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 (blocking))

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 $C = (\text{CState}, \text{CInit}, \{\text{COp}_i\}_{i \in I}, \tau_C, \text{CFin})$ iff the following conditions hold for $R = R' \circ \tau_C^*$:

$$\text{CInit} \subseteq \text{AInit} \circ \tau_A^* \circ R \quad (9)$$

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

and for matching operations AOp and COp:

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

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

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

□

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

$$\text{DS.InitBlock.internal } \forall CState' \bullet CInit \Rightarrow \exists AState' \bullet AInit \circ \tau_A^* \wedge R'$$

$$\text{DS.AppBlock.internal } \forall CState; AState; i : I \bullet R \wedge \text{pre } \overleftrightarrow{AOp}_i \Rightarrow \text{pre } \overleftrightarrow{COp}_i$$

$$\text{DS.CorrBlock.internal}$$

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

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

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

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 Sect. 3.3 above.

5.5.2. Upward simulation

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

$$CInit \circ \tau_C^* \circ T \subseteq AInit \circ \tau_A^* \quad (14)$$

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

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

and for matching operations AOp and COp:

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

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

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

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 $T \subseteq CState \times AState$ 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 $C = (CState, CInit, \{COp_i\}_{i \in I}, \tau_C, CFin)$ then $\tau_C^* \circ T \circ \tau_A^*$ is also an upward simulation between A and C.

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

The initialisation condition can be simplified using Theorem 9, clearly $\tau_C^* \circ T = T$. 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 (blocking))

The relation $T' \subseteq \text{CState} \times \text{AState}$ is an upward 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 $C = (\text{CState}, \text{CInit}, \{\text{COp}_i\}_{i \in I}, \tau_C, \text{CFin})$ iff the following conditions hold for $T \equiv \tau_C^* \circ T' \circ \tau_A^*$:

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

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

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

and for matching operations AOp and COp :

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

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

□

For specifications in Z , this leads to the following conditions. (See [BoD02a] for the essential steps in rephrasing the finalisation condition.)

US.InitBlock.internal $\forall \text{CState}' ; \text{AState}' \bullet T' \wedge \text{CInit} \Rightarrow \text{AInit} \circ \tau_A^*$

US.CorrBlock.internal $\forall i : I ; \text{CState} ; \text{CState}' ; \text{AState}' \bullet$

$$(\forall \text{AState} \bullet T \Rightarrow \neg \text{liv } \overleftrightarrow{\text{AOp}_i}) \wedge \text{COp}_i \wedge T' \Rightarrow \exists \text{AState} \bullet T \wedge \text{AOp}_i \circ \tau_A^*$$

US.FinRef.internal $\forall \text{CState} ; \text{CState}' \mid \neg \text{CState} \uparrow \wedge \tau_C^* \wedge \text{CState}' \downarrow \bullet$

$$\exists \text{AState} \bullet T \wedge \text{AState} \downarrow \wedge \forall i : I \bullet \text{pre } \text{AOp}_i \Rightarrow (\text{pre } \text{COp}_i)'$$

US.DivStates $\forall \text{CState} \bullet \text{CState} \uparrow \Rightarrow \exists \text{AState} \bullet T \wedge \text{AState} \uparrow$

US.DivOp $\forall i : I ; \text{CState} \bullet \text{liv } \overleftrightarrow{\text{COp}_i} \Rightarrow \exists \text{AState} \bullet T \wedge \text{liv } \overleftrightarrow{\text{AOp}_i}$

Again note that in the absence of internal operations, these conditions collapse to those in Sect. 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 (B, N, D) where

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

As discussed in Sect. 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 Sects. 5.5.1 and 5.5.2. Instantiating Theorem 4 with the embedding of Definition 12 initially gives:

$$\text{CInit} \circ \tau_C^* \subseteq \text{AInit} \circ \tau_A^* \circ R \quad (25)$$

$$R \circ CState \uparrow \triangleleft \tau_C^* \circ CFin \subseteq AState \uparrow \triangleleft \tau_A^* \circ AFin \quad (26)$$

$$(AState \uparrow \triangleleft R) = (R \triangleright CState \uparrow) \quad (27)$$

and for matching operations AOp and COp:

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

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

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

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

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

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:

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

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

□

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

DS.InitNonBlock.internal $\forall CState' \bullet CInit \circ \tau_C^* \Rightarrow \exists AState' \bullet AInit \circ \tau_A^* \wedge R'$

DS.AppNonBlock.internal

$$\forall CState; AState; i : I \bullet R \wedge \text{pre } \overleftrightarrow{AOp}_i \wedge \neg \text{liv } \overleftrightarrow{AOp}_i \Rightarrow \text{pre } \overleftrightarrow{COp}_i \wedge \neg \text{liv } \overleftrightarrow{COp}_i$$

DS.CorrNonBlock.internal

$$\forall i : I; AState; \Delta CState \bullet \neg \text{liv } \overleftrightarrow{AOp}_i \wedge \text{pre } \overleftrightarrow{AOp}_i \wedge R \wedge \overleftrightarrow{COp}_i \Rightarrow \exists AState' \bullet \overleftrightarrow{AOp}_i \wedge R'$$

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 $\text{div Op} == \text{liv } \overleftrightarrow{\text{Op}} \cup \text{dom } \overleftrightarrow{\text{Op}}$:

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

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

□

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

US.InitNonBlock.internal $\forall CState'; AState' \bullet CInit \circ \tau_C^* \wedge T' \Rightarrow AInit \circ \tau_A^*$

US.AppDivOp $\forall i : I; CState \bullet \text{div COp}_i \Rightarrow \exists AState \bullet T \wedge \text{div AOp}_i$

US.CorrNonBlock.internal $\forall i : I; \Delta CState; AState' \bullet$

$$(\forall AState \bullet T \Rightarrow \neg \text{div AOp}_i) \wedge \overleftrightarrow{\text{COp}}_i \wedge T' \Rightarrow \exists AState \bullet T \wedge \overleftrightarrow{\text{AOp}}_i$$

US.DivStates $\forall CState \bullet CState \uparrow \Rightarrow \exists AState \bullet T \wedge AState \uparrow$

US.NonDivStates $\forall CState \mid \neg CState \uparrow \bullet \exists AState \bullet T \wedge \neg AState \uparrow$

using the analogous definition of div Op , etc.

The downward and upward conditions derived in this section extend those discussed in Sect. 3.3 by adding internal evolution. As discussed in Sect. 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 Sect. 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 Sects. 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 [WoD96] and [DeB01, Sect. 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 [DeB01, 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 $(GB \times \text{seq } Output)_{\perp, \omega}$ when D_r has $GB \times \text{seq } Output$. 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 ($\mathbf{N}$, $\mathbf{B}$, $\mathbf{D}$) 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 sections.

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.

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 Sect. 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 Sect. 5.5.1, we end up with the proof obligation:

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

Unlike in Sect. 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).

$$Sim == I \rightarrow Output$$

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

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

DS.FinDemBlock.internal

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

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 Sect. 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 pres-

¹² Taking a CSP view, let $P_i := pli \rightarrow stop$ etc, and writing $(\tau \rightarrow A) \square (\tau \rightarrow B)$ as $A \sqcap B$, this example is $(P_1 \sqcap Q_1) \sqcap (P_2 \sqcap Q_2) \not\sqsubseteq (P_1 \sqcap P_2) \sqcap (Q_1 \sqcap Q_2)$.

ence of internal operations, a similar extra condition is needed for downward simulations, and this we have called **DS.FinDemBlock.internal**.

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.internal

$$\begin{aligned} \forall E : \text{Sim}; \Delta C\text{State} \mid \neg C\text{State} \uparrow \wedge \tau_C^* \wedge C\text{State}' \downarrow \wedge (\text{Maxsim}(E))' \bullet \\ \exists A\text{State}, F : \text{Sim} \bullet T \wedge A\text{State} \downarrow \wedge F \subseteq E \wedge \text{Maxsim}(F) \end{aligned}$$

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 $F\text{cond}D$ in Definition 8 as follows.

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

with $F\text{cond} = F\text{cond}ND(\theta\text{State}, E)$ where

$$F\text{cond}ND(s, E) == E \subseteq \left\{ (i, \text{out}) \mid \begin{aligned} &\exists \text{State}'; \text{Output} \bullet \theta\text{State} = s \wedge \\ &(\text{pre } Op_i \Rightarrow Op_i) \wedge \theta\text{Output} \neq \text{out} \wedge (i, \theta\text{Output}) \notin E \end{aligned} \right\} \quad (42)$$

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 Sect. 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

$$\begin{aligned} \forall a : A\text{State}; c, c' : C\text{State}; E \bullet (a, c) \in R \wedge (c, c') \in \tau_C^* \wedge c \notin C\text{State} \uparrow \wedge c' \in C\text{State} \downarrow \wedge F\text{cond}ND(c', E) \\ \Rightarrow \exists a' : A\text{State} \bullet (a, a') \in \tau_A^* \wedge a' \in A\text{State} \downarrow \wedge F\text{cond}ND(a', E) \end{aligned}$$

Compared to the condition in Sect. 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:

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

This leads to the following refinement condition.

Fig. 5. Refusals in the non-blocking model

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

DS.FinDemNonBlock.internal

$$\forall AState; E; CState; CState' \mid R \wedge \neg CState \uparrow \wedge \tau_C^* \wedge CState' \downarrow \wedge (Maxsimtot(E))' \bullet \\ \exists AState' \bullet \tau_A^* \wedge AState' \downarrow \wedge (Maxsimtot(E))'$$

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 Sect. 5.6.2 and, additionally, the conditions represented by the finalisation $CFin \subseteq T \S AFin$. This leads to the requirement that

$$\forall c, c' : CState; E \bullet c \notin CState \uparrow \wedge (c, c') \in \tau_C^* \wedge c' \in CState \downarrow \wedge FcondND(c', E) \\ \Rightarrow \exists a, a' : AState \bullet (c, a) \in T \wedge (a, a') \in \tau_A^* \wedge a' \in AState \downarrow \wedge FcondND(a', E)$$

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

$$US.FinDemNonBlock.internal \quad \forall \Delta CState; E : Sim \mid \neg CState \uparrow \wedge \tau_C^* \wedge CState' \downarrow \wedge (Maxsimtot(E))' \bullet \\ \exists \Delta AState \bullet T AState' \downarrow \wedge \tau_S^* \wedge (Maxsimtot(E))'$$

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 Fig. 5). The non-blocking totalisation of A (given as a CSP process) is

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

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 $FcondND$, 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 Fig. 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,

Fig. 7. Non-determinism and refusals inclusion

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 Fig. 7 the refusals initially in C are $\{Op.o_1\}$ and $\{Op.o_2\}$ but not $\{Op.o_1, Op.o_2\}$. The condition $\forall CState; E \bullet FcondD(\theta CState, E) \Rightarrow \exists AState \bullet T \wedge FcondD(\theta AState, E)$ 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 ADT specified in the usual Z “states-and-operations” style [Spi92] but with some additional syntax to make the ADT structure explicit [DeB01]. 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 [WoD96, DeB01];
- to derive refinement for Z with the “guard” interpretation, i.e., applying an undefined operation is impossible [BDW99, DeB01];
- to define failures refinement for Z, in both the “guard” and “precondition” interpretations, observing refusals and outputs [BoD02a, DeB03];
- to define readiness refinement for Z, using a version of relational refinement for partial relations [DeB03];
- to define failures–divergences refinement for Z ADTs, in both interpretations, allowing internal operations but not outputs [DeB06].

As noted in [MBD00, BdR03], the consideration of internal operations [DeB06] 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 (Sects. 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 Sect. 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 [RSS98]. Related theories verified previously include our earlier work on weak refinement ([DBB98], in [Sch05]) and Cooper, Stepney and Woodcock’s generalised Z refinement rules ([CSW02], in [SGH06]). 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 [Sch06].

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 Sect. 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 [Jos88] 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 Sects. 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 [ScT04] 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

- [Abr96] Abrial J-R (1996) The B-Book: assigning programs to meanings. Cambridge University Press, Cambridge

- [BdR03] Boiten EA, de Roever W-P (2003) Getting to the bottom of relational refinement: relations and correctness, partial and total. In: Berghammer R, Möller B (eds) 7th International seminar on relational methods in computer science (RelMiCS 7). University of Kiel, pp 82–88
- [BDW99] Bolton C, Davies J, Woodcock JCP (1999) On the refinement and simulation of data types and processes. In: Araki K, Galloway A, Taguchi K (eds) International conference on integrated formal methods 1999 (IFM'99). Springer, Berlin, pp 273–292
- [BeZ86] Berghammer R, Zierer H (1986) Relation algebraic semantics of deterministic and non-deterministic programs. *Theor Comp Sci* 43:123–147
- [BHR84] Brookes SD, Hoare CAR, Roscoe AW (1984) A theory of communicating sequential processes. *J ACM* 31(3):560–599
- [BoB88] Bolognesi T, Brinksma E (1988) Introduction to the ISO Specification Language LOTOS. *Comput Networks ISDN* 14(1):25–59
- [BoD02a] Boiten EA, Derrick J (2002) Unifying concurrent and relational refinement. In: Derrick J, Boiten EA, Woodcock JCP, von Wright J (eds) *Refine 2002*, ENTCS 70:94–131
- [BoD02b] Bolton C, Davies J (2002) Refinement in Object-Z and CSP. In: Butler M, Petre L, Sere K (eds) *Integrated formal methods (IFM 2002)*, Lecture notes in computer science, vol 2335. Springer, Berlin, pp 225–244
- [BoD06] Bolton C, Davies J (2006) A singleton failures semantics for communicating sequential processes. *Form Asp Comput* 18:181–210
- [BoG05] Bowman H, Gomez R (2005) *Concurrency theory: calculi and automata for modelling untimed and timed concurrent systems*. Springer, New York
- [Bol02] Bolton C (2002) On the refinement of state-based and event-based models. Ph.D. thesis, University of Oxford
- [BoL03] Bolton C, Lowe G (2003) A hierarchy of failures-based models. In: Corradini F, Nestmann U (eds) *Proceedings of express 2003: 10th international workshop on expressiveness in concurrency*. Elsevier Science, Amsterdam
- [BoL05] Bolton C, Lowe G (2005) A hierarchy of failures-based models: theory and application. *Theor Comp Sci* 330(3):407–438
- [BrR85] Brookes SD, Roscoe AW (1985) An improved failures model for communicating processes. In: Brookes SD, Roscoe AW, Winskel G (eds) *Seminar on concurrency, Lecture notes in computer science*, vol 197. Springer, Berlin, pp 281–305
- [CSW02] Cooper D, Stepney S, Woodcock J (2002) Derivation of Z refinement proof rules: forwards and backwards rules incorporating input/output refinement. Technical Report YCS-2002-347, University of York. URL: <http://www-users.cs.york.ac.uk/~susan/bib/ss/z/zrules.htm>
- [DBB98] Derrick J, Boiten EA, Bowman H, Steen MWAS (1998) Specifying and refining internal operations in Z. *Form Asp Comput* 10:125–159
- [DeB01] Derrick J, Boiten EA (2001) *Refinement in Z and object-Z: foundations and advanced applications*, FACIT series. Springer, London
- [DeB03] Derrick J, Boiten EA (2003) Relational concurrent refinement. *Form Asp Comput* 15(1):182–214
- [DeB06] Derrick J, Boiten EA (2006) Relational concurrent refinement with internal operations. In: Aichernig B, Boiten EA, Derrick J, Groves L (eds) *BCS-FACS Refinement Workshop*, ENTCS 187:35–53
- [DeH06] Deutsch M, Henson MC (2006) An analysis of refinement in an abortive paradigm. *Form Asp Comput* 18(3):329–363
- [Doo94] Doornbos H (1994) A relational model of programs without the restriction to Egli-Milner constructs. In: Olderog E-R (ed) *PROCOMET '94*, IFIP, pp 357–376
- [dRE98] De Roever W-P, Engelhardt K (1998) *Data refinement: model-oriented proof methods and their comparison*. Cambridge University Press, Cambridge
- [DuC05] Dunne S, Conroy S (2005) Process refinement in B. In: Treharne H, King S, Henson MC, Schneider S (eds) *ZB 2005: formal specification and development in Z and B*, 4th international conference of B and Z users, Lecture notes in computer science, vol 3455. Springer, Berlin, pp 45–64
- [Fis97] Fischer C (1997) CSP-OZ—A combination of CSP and Object-Z. In: Bowman H, Derrick J (eds) *Second IFIP international conference on formal methods for open object-based distributed systems*. Chapman & Hall, London, pp 423–438
- [He89] He J (1989) Process refinement. In: McDermid J (ed) *The theory and practice of refinement*. Butterworths, London
- [HeI93] Hennessy M, Ingólfsdóttir A (1993) A theory of communicating processes with value passing. *Inf Comput* 107(2):202–236
- [HHS86] He J, Hoare CAR, Sanders JW (1986) Data refinement refined. In: Robinet B, Wilhelm R (eds) *Proc. ESOP'86*. Lecture notes in computer science, vol 213. Springer, Berlin, pp 187–196
- [Hoa85] Hoare CAR (1985) *Communicating sequential processes*. Prentice-Hall, Englewood Cliffs
- [HoH98] Hoare CAR, He J (1998) *Unifying theories of programming*. Prentice-Hall, Englewood Cliffs
- [Jos88] Josephs MB (1988) A state-based approach to communicating processes. *Distrib Comput* 3:9–18
- [Led91] Leduc G (1991) On the role of implementation relations in the design of distributed systems using LOTOS. Ph.D. thesis, University of Liège
- [MBD00] Miarka R, Boiten EA, Derrick J (2000) Guards, preconditions and refinement in Z. In: Bowen JP, Dunne S, Galloway A, King S (eds) *ZB2000: Formal specification and development in Z and B*. Lecture notes in computer science, vol 1878. Springer, Berlin, pp 286–303
- [Mil89] Milner R (1989) *Communication and concurrency*. Prentice-Hall, Englewood Cliffs
- [ReS06] Reeves S, Streader D (2006) State- and event-based refinement. Technical report, Department of Computer Science, University of Waikato
- [Ros98] Roscoe AW (1998) *The theory and practice of concurrency*. Prentice-Hall, Englewood Cliffs
- [RSS98] Reif W, Schellhorn G, Stenzel K, Balser M (1998) Structured specifications and interactive proofs with KIV. In: Bibel W, Schmitt P (eds) *Automated deduction—a basis for applications*, vol II: systems and implementation techniques, Chap. 1: Interactive Theorem Proving. Kluwer, Dordrecht, pp 13–39
- [Sch05] Schellhorn G (2005) ASM refinement and generalizations of forward simulation in data refinement: a comparison. *Theor Comp Sci* 336(2–3):403–435
- [Sch06] Schellhorn G (2006) Web presentation of the KIV proofs of 'Relational Concurrent Refinement Part II: Internal Operations and Output'. URL: <http://www.informatik.uni-augsburg.de/swt/projects/Refinement/Web/CSPRef>
- [ScT04] Schneider S, Treharne H (2004) CSP theorems for communicating B machines. *Form Asp Comput* 17(4):390–422

- [SGH06] Schellhorn G, Grandy H, Haneberg D, Reif W (2006) The Mondex challenge: machine checked proofs for an electronic purse. In: Misra J, Nipkow T, Sekerinski E (eds) Formal methods 2006, Proceedings. Lecture notes in computer science, vol 4085. Springer, Berlin, pp 16–31
- [SmD02] Smith G, Derrick J (2002) Abstract specification in Object-Z and CSP. In: George C, Miao H (eds) Formal methods and software engineering. Lecture notes in computer science, vol 2495. Springer, Berlin, pp 108–119
- [Spi92] Spivey JM (1992) The Z notation: a reference manual, 2nd edn. Prentice-Hall, Englewood Cliffs
- [VaT95] Valmari A, Tienari M (1995) Compositional failure-based semantics models for basic LOTOS. *Form Asp Comput* 7(4):440–468
- [vGl01] van Glabbeek, RJ (2001) The linear time—branching time spectrum I. The semantics of concrete sequential processes. In: Bergstra JA, Ponse A, Smolka SA (eds) Handbook of process algebra. North-Holland, Amsterdam pp 3–99
- [WoD96] Woodcock JCP, Davies J (1996) Using Z: specification, refinement, and proof. Prentice-Hall, Englewood Cliffs
- [WoM90] Woodcock JCP, Morgan CC (1990) Refinement of state-based concurrent systems. In: Bjørner D, Hoare CAR, Langmaack H (eds) VDM'90: VDM and Z!—formal methods in software development, Lecture notes in computer science, vol 428. Springer, Berlin

Received 16 January 2007

Accepted in revised form 26 November 2007 by B. K. Aichernig, M. J. Butler and L. Groves

Published online 4 January 2008