

Slotted-Circus

A UTP-Family of Reactive Theories*

Andrew Butterfield¹, Adnan Sherif², and Jim Woodcock³

¹ Trinity College Dublin, Andrew.Butterfield@cs.tcd.ie

² Universidade Federal de Pernambuco, ams@cin.ufpe.br

³ University of York, Jim.Woodcock@cs.york.ac.uk

Abstract. We present a generic framework of UTP theories for describing systems whose behaviour is characterised by regular time-slots, compatible with the general structure of the *Circus* language [WC01a]. This “slotted-*Circus*” framework is parameterised by the particular way in which event histories are observable within a time-slot, and specifies what laws a desired parameterisation must obey in order for a satisfactory theory to emerge.

Two key results of this work are: the need to be very careful in formulating the healthiness conditions, particularly **R2**; and the demonstration that synchronous theories like SCSP [Bar93] do not fit well with the way reactive systems are currently formulated in UTP and *Circus*.

1 Introduction

1.1 *Circus* and slotted-*Circus*

The formal notation *Circus* [WC01a] is a unification of Z and CSP, and has been given a UTP semantics [WC02]. A *Circus* text describes behaviour as a collection of actions, which are a combination of processes with mutable state. However, apart from event sequencing, there is no notion of time in *Circus*.

A timed version of *Circus* (*Circus* Time Action or CTA) has been explored [SH02, She06] that introduces the notion of discrete time-slots in which sequences of events occur. In CTA, we have a two-level notion of history: the top-level views history as a sequence of time-slots; whilst the bottom-level records a history of events within a given slot. The key notion in this paper is that we can instantiate the bottom-level history in a variety of ways: as simple traces, or multisets of events, or as the more complex “micro-slot” structures used in the operational semantics of Handel-C [BW05].

This paper describes a generalisation of CTA called “slotted-*Circus*”, which is a collection of theories parameterised by different ways to instantiate the bottom-level event history within a time-slot. The motivation behind this work is the desire to re-cast existing semantics for Handel-C into the UTP framework so that *Circus* can be used as a specification language.

The Handel-C denotational [BW02] and operational semantics use this time-slot model, but with varying degrees of complexity in the slot structure, depending on which language constructs we wish to support. The slotted-*Circus* framework reported here is intended to be a foundation for formulating the common parts of these models, making it easier to explore the key differences.

1.2 UTP: General Principles

Theories in UTP are expressed as predicates over a pre-defined collection of free observation variables, referred to as the *alphabet* of the theory. The predicates are generally used to describe a

* research reported in this paper was partially supported by Qinetic

$$\begin{aligned}
P; Q &\hat{=} \exists obs_0 \bullet P[obs_0/obs'] \wedge Q[obs_0/obs] \\
P \triangleleft c \triangleright Q &\hat{=} c \wedge P \vee \neg c \wedge Q \\
P \sqcap Q &\hat{=} P \vee Q \\
\prod_{i \in I} P_i &\hat{=} \exists i : I \bullet P_i \\
S \sqsubseteq P &\hat{=} [P \Rightarrow S]
\end{aligned}$$

Fig. 1. Basic UTP Operators

relation between a before-state and an after-state, the latter typically characterised by dashed versions of the observation variables. A predicate whose free variables are all undashed is called a (*pre*-)condition. A given theory is characterised by its alphabet, and a series of *healthiness conditions* that constrain the valid assertions that predicates may make. In almost all cases there are some basic operators common to every theory (Figure 1). Sequential composition ($P; Q$) corresponds to relational composition, *i.e.*, the existence of an intermediate state (obs_0), such that P relates obs to obs_0 , whilst Q relates obs_0 to obs' . The conditional $P \triangleleft c \triangleright Q$ is generally used when c is a condition and asserts that P holds if c is true, otherwise it asserts Q . Nondeterminism between two predicates $P \sqcap Q$ is simply logical disjunction, which extends to an existential quantifier for a nondeterministic choice over an indexed set ($\prod_i P_i$). We capture the notion of refinement \sqsubseteq as logical entailment between the implementation and specification predicates, quantified over all free variables.

We note that UTP follows the key principle that “programs are predicates” [Hoa85b], and so does not distinguish between the syntax of some language and its semantics as alphabetised predicates. In other words, we view the language constructs as being “syntactic sugar” for their predicate semantics, rather than defining some form of semantic function mapping some abstract syntax type to some domain capturing its meaning.

1.3 Structure and Focus

The main technical emphasis of this paper is on the construction of the generic framework and the required healthiness conditions, with the semantics of the language constructs and a case study provided to give a feel for its utility. We first present the syntax §2, generic framework §3, healthiness conditions §4, and semantics §5. We then discuss instantiation §6 and describe a case-study §7, before mentioning related §8 and future §9 work, and concluding §10. Two appendices give supporting material.

2 Syntax

The syntax of Slotted-*Circus* is similar to that of *Circus*, and a subset is shown in Figure 2. The notation X^+ denotes a sequence of one or more X . We assume an appropriate syntax for describing expressions and their types, subject only to the proviso that booleans and non-negative integers are included.

The basic actions *Skip*, *Stop*, *Chaos* are similar to the corresponding CSP behaviours [Hoa85a, Sch00], respectively denoting actions that do nothing and terminate, do nothing and wait forever, or act unpredictably forever. We also introduce (multiple) assignment ($:=$) and event (communication) prefixes $\text{Comm} \rightarrow \text{Action}$ as basic actions. The communication prefixes range over communicating a value on a channel (Name.Expr), sending a value on a channel (Name!Expr), or receiving

$$\begin{aligned}
\text{Action} &::= \text{Skip} \mid \text{Stop} \mid \text{Chaos} \\
&\mid \text{Name}^+ := \text{Expr}^+ \mid \text{Comm} \rightarrow \text{Action} \mid \text{Action} \square \text{Action} \\
&\mid \text{Action} \llbracket \text{VS} \mid \text{CS} \mid \text{VS} \rrbracket \text{Action} \mid \text{Action} \setminus \text{CS} \\
&\mid \mu \text{Name} \bullet F(\text{Name}) \mid \text{Wait } t \mid \dots \\
\text{Comm} &::= \text{Name}.\text{Expr} \mid \text{Name}!\text{Expr} \mid \text{Name}?\text{Name} \\
\text{Expr} &::= \text{expression} \\
t &::= \text{positive integer valued expression} \\
\text{Name} &::= \text{channel or variable names} \\
\text{CS} &::= \text{channel name sets} \\
\text{VS} &::= \text{variable sets}
\end{aligned}$$

Fig. 2. Slotted-*Circus* Syntax

a value on a channel ($\text{Name}?\text{Name}$). The composite action operator \square denotes external choice, whilst parallel composition of actions ($\llbracket \text{VS} \mid \text{CS} \mid \text{VS} \rrbracket$) is parameterised by three sets, the first and third denoting the variables the corresponding action may modify, while the middle one specifies the synchronisation channels. We require that parallel processes modify disjoint parts of the state. We also have hiding ($\setminus \text{CS}$) and recursively defined actions ($\mu \text{Name} \bullet F(\text{Name})$).

The key construct related to time-slots, and hence not part of *Circus*, is $\text{Wait } t$ which denotes an action that simply waits for t time-slots to elapse, and then terminates.

3 Generic Slot-Theory

Both the semantics of Handel-C [BW05] and the timed extension to *Circus* called ‘‘Circus Timed Actions (CTA)’’ [SH02, She06] have in common the fact that the models involve a sequence of ‘‘slots’’ that capture the behaviour of the system between successive clock ticks. These slots contain information about the events that occurred during that time slot (‘‘history’’) as well as the events being refused at that point. A key feature of all these semantic models is that the progress of events during a time-slot is observable, rather than just the overall outcome for an entire slot. While the initial goal was to develop a synchronous variant of *Circus*, it rapidly became clear that it was worth investing time in a generic slot-based theory, which could then be specialised to cover synchronicity, CTA, and the various slot-models that could be used to characterise Handel-C and similar synchronous hardware languages at various levels of detail.

We begin our description of the generic slotted theory by noting that it is parametric in three inter-related aspects:

- A given set of events, E .
- A type constructor \mathcal{H} that builds a slot’s history-type from an event type.
- A collection of basic functions that work with $\mathcal{H} E$, which must satisfy certain laws.

Given \mathcal{H} , we then define the notion of a slot (\mathcal{S}) as being a pair: a history and a set of events denoting a refusal:

$$\mathcal{S} E \hat{=} (\mathcal{H} E) \times (\mathbb{P} E) \tag{1}$$

In a sense a slot is similar to the notion of a failure in CSP [Ros97], except that it covers only the events within a single time-slot (*i.e.*, between two successive clock ticks). Given a notion of time-slot, we then introduce the top-level notion of event history as being a non-empty sequence

$$\begin{aligned}
Acc_{\mathcal{H}} &: \mathcal{H} E \rightarrow \mathbb{P} E \\
EqvTrc_{\mathcal{H}} &: E^* \leftrightarrow \mathcal{H} E \\
HNull_{\mathcal{H}} &: \mathcal{H} E \\
\preceq_{\mathcal{H}} &: \mathcal{H} E \leftrightarrow \mathcal{H} E \\
Hadd_{\mathcal{H}} &: \mathcal{H} E \times \mathcal{H} E \rightarrow \mathcal{H} E \\
Hsub_{\mathcal{H}} &: \mathcal{H} E \times \mathcal{H} E \rightarrow \mathcal{H} E \\
HHide_{\mathcal{H}} &: \mathbb{P} E \rightarrow \mathcal{H} E \rightarrow \mathcal{H} E \\
HSync_{\mathcal{H}} &: \mathbb{P} E \rightarrow \mathcal{H} E \times \mathcal{H} E \rightarrow \mathbb{P}(\mathcal{H} E)
\end{aligned}$$

Fig. 3. Generic Functions over $\mathcal{H} E$

of slots. The presence of clock-ticks in the history is denoted by the adjacency of two slots, so a slot-sequence of length $n + 1$ describes a situation in which the clock has ticked n times.

We can now describe the observational variables of our generic UTP theory:

$ok : \mathbb{B}$ —True if the process is stable, *i.e.*, not diverging.
 $wait : \mathbb{B}$ —True if the process is waiting, *i.e.*, not terminated.
 $state : Var \rightarrow Value$ —An environment giving the current values of slotted-*Circus* variables
 $slots : (\mathcal{S} E)^+$: —A non-empty sequence of slots recording the behaviour of the system.

The variables ok , $wait$ play the same role as the in the reactive systems theory in [HH98, Chp. 8], while $state$ follows the trend in [SH02] of grouping all the program variables under one observational variable, to simplify the presentation of the theory.

In order to give the generic semantics of the language, we need to provide six functions and two relations over $\mathcal{H} E$, listed in Figure 3. Function Acc returns the set of events mentioned (*Accepted*) in its history argument. The relation $EqvTrc$ relates a history to all event sequences (traces) compatible with it. $HNull$ is a constant denoting an empty history. Infix symbol \preceq captures the notion of one history being a prefix, of pre-history of another, and is required to be a pre-order. The functions $Hsub$ and $Hadd$ capture the notions of history subtraction and addition (extension). In particular we note that $Hsub$ is partial and is only defined when the second argument is a pre-history of the first. Function $HHide$ acts to remove a set of events from a history. Finally the $HSync$ function generates all the possible histories that can result from the synchronisation of two histories over a given event set.

In order to produce a coherent theory, the functions have to obey a number of laws, listed in Appendix A. Most of the properties concerned capture reasonable behaviours that one would expect of histories, *e.g.*, that history addition is associative, or that the null history acts as a unit. Most of these laws were determined by the needs of the general theory, in particular the definitions and proofs needed to establish the required healthiness conditions.

As an example, a variation of the CTA theory of [She06] can be captured by defining an event history ($\mathcal{H}_{CTA} E$) to be a sequence of events, and instantiating most of the functions and relations as the corresponding ones for sequences.

$$\mathcal{H}_{CTA} E \hat{=} E^* \tag{2}$$

3.1 Derived Types and Operators

Given the definition of \mathcal{H} , and the associated functions and relations, we need to use these to define the corresponding aspects for slots, and the slot-sequences that comprise our observational

$$\begin{aligned}
EqvTrace &: E^* \leftrightarrow (\mathcal{S} E)^* \\
Refs &: (\mathcal{S} E)^+ \rightarrow (\mathbb{P} E)^+ \\
EqvRef &: (\mathcal{S} E)^+ \rightarrow \mathbb{P} E \\
\preceq &: (\mathcal{S} E)^+ \leftrightarrow (\mathcal{S} E)^+ \\
\approx &: \mathcal{S} E \leftrightarrow \mathcal{S} E \\
\cong &: (\mathcal{S} E)^+ \leftrightarrow (\mathcal{S} E)^+ \\
Sadd_{\mathcal{S}} &: \mathcal{S} E \times \mathcal{S} E \rightarrow \mathcal{S} E \\
Ssub_{\mathcal{S}} &: \mathcal{S} E \times \mathcal{S} E \rightarrow \mathcal{S} E \\
\# &: ((\mathcal{S} E)^+ \times (\mathcal{S} E)^+) \rightarrow (\mathcal{S} E)^+ \\
\ll &: ((\mathcal{S} E)^+ \times (\mathcal{S} E)^+) \rightarrow (\mathcal{S} E)^+
\end{aligned}$$

Fig. 4. Derived Functions and Relations

variables (see Figure 4). *EqvTrace*, defined in terms of *EqvTrc*, relates traces to slot-sequences with which they are compatible. The functions *Refs* and *EqvRef* extract refusals from slot-sequences, with the former returning a refusal-set list, whilst the latter singles out the last refusal set. A slot-sequence s is a slot-prefix of a slot-sequence t , written $s \preceq t$ if the front of s is a prefix of t and the history component of the last slot of s is a history-prefix of the corresponding component of the first slot of $t - s$. The relation \preceq is a pre-order. Slot equivalence \approx and Slot-sequence equivalence (\cong) are the symmetric closure of \preceq and \approx respectively, giving equivalence relations. An important point to note here is that if $s \cong t$, then s and t are identical, except for the refusal values in the last slot in each.

The notions of adding (extending) and subtracting histories are lifted to the slot level, but here an issue immediately arises as to how the refusal components are handled. If we consider history addition, then $Hadd(h_1, h_2)$ is intended to capture the history resulting from the events of history h_1 , followed by those of h_2 . We now note that in most CSP-like theories, a failure consisting of a trace/history of events (h) coupled with a refusal set (r), is to be interpreted as stating that the process under consideration is refusing the events in r , *after* having performed the events in h . Given this interpretation, we are then required to specify slot addition and subtraction as follows:

$$\begin{aligned}
Sadd((h_1, -), (h_2, r_2)) &\hat{=} (Hadd(h_1, h_2), r_2) \\
Ssub((h_1, r_1), (h_2, -)) &\hat{=} (Hsub(h_1, h_2), r_1)
\end{aligned}$$

For history subtraction, the value $Hsub(h_1, h_2)$ is defined only if $h_2 \preceq h_1$, and denotes those events in h_1 that occurred after those in h_2 . The significance of this interpretation is important, as will be made clear when we consider an attempt to model Synchronous CSP (SCSP) [Bar93] later in this paper. A consequence of this interpretation is that one of the healthiness conditions discussed in the next section (**R2**) becomes more complex.

Given slot addition and subtraction, these can then be lifted to act on slot-sequences, as $\#$ and \ll respectively. The latter is only defined if its second argument is a \preceq -prefix of its first. Slot-sequence addition concatenates its two arguments, merging the last slot of the first with the first slot of the second:

$$slots_1 \# slots_2 \hat{=} front(slots_1) \wedge \langle Sadd(last(slots_1), head(slots_2)) \rangle \wedge tail(slots_2) \quad (3)$$

Slot-sequence subtraction $s \searrow t$ is defined when $t \preceq s$, in which case both s and t can be written as

$$\begin{aligned} s &= pfx \wedge \langle slot_s \rangle \wedge sfx \\ t &= pfx \wedge \langle slot_t \rangle \end{aligned}$$

In this case, the subtraction becomes:

$$s \searrow t \hat{=} \langle Ssub(slot_s, slot_t) \rangle \wedge sfx \quad (4)$$

4 Healthiness Conditions

Given that we are defining semantics as predicates over before- and after-observations, we need to ensure that what we write is feasible, in that we do not describe behaviour that is computationally or physically infeasible (*e.g.*, undoing past events). In UTP, the approach to handling feasibility is to define a number of so-called healthiness conditions that characterise the sort of predicates which make sense in the intended interpretation of the theory.

While the notion of healthiness-conditions is well-understood in the UTP community, we are still going to take time for the presentation that follows, as we highlight a prevalent use of overloading that can have unexpected effects in inexperienced hands.

Given a healthiness condition called \mathbf{H} we introduce two functions, \mathbf{mkH} and \mathbf{isH} . In order to denote a healthiness condition, we require that the former is an idempotent monotonic predicate transformer, w.r.t. to the standard ordering used in UTP, namely that $S \sqsubseteq P$ iff $[P \Rightarrow S]$. The role of \mathbf{mkH} is to convert an un-healthy predicate into a healthy one, in some fashion, but also to leave already healthy predicates unchanged (hence the need for idempotency, so that a healthy predicate is a fixed-point of \mathbf{mkH}).

$$\begin{aligned} \mathbf{mkH} &: \text{Predicate} \rightarrow \text{Predicate} \\ \mathbf{mkH} &= \mathbf{mkH} \circ \mathbf{mkH} \end{aligned}$$

Function \mathbf{isH} asserts a healthiness condition, *i.e.*, is a higher order predicate that tests a given predicate to see if it is healthy:

$$\begin{aligned} \mathbf{isH} &: \text{Predicate} \rightarrow \mathbb{B} \\ \mathbf{isH}(P) &\hat{=} P \equiv \mathbf{mkH}(P) \end{aligned}$$

We can summarise by saying that a healthy predicate is a fixed-point of the corresponding healthiness predicate transformer. In most material on UTP, it is conventional to overload the notation \mathbf{H} to refer to both \mathbf{mkH} and \mathbf{isH} , with the use usually being clear from context. In either case it is also conventional to refer in general to \mathbf{H} as a healthiness condition, even in a context where it would actually be a predicate transformer. We shall adopt this convention in the sequel.

However a hazard can arise when alternative formulations of \mathbf{H} are available; note that different functions may have the same set of fixed-points. We illustrate this later when discussing $\mathbf{R2}$.

The healthiness conditions we introduce here for slotted-*Circus* parallel some of those in [HH98, Chp. 8] for general reactive systems, namely $\mathbf{R1}$, $\mathbf{R2}$, $\mathbf{R3}$ and $\mathbf{CSP1}$.

4.1 Reactive Healthiness

We shall discuss $\mathbf{R1}$ and $\mathbf{R3}$ first, as these are fairly straightforward, while $\mathbf{R2}$ deserves some discussion, as its adaption for slotted-*Circus* was decidedly non-trivial.

R1 simply states that a slotted-*Circus* process cannot undo the past, or in other words, that the $slots'$ observation must be an extension of $slots$, whilst **R3** deals with the situation when a process has not actually started to run, because a prior process has yet to terminate, characterised by $wait = \text{TRUE}$. In this case the action of a yet-to-be started process should simply be to do nothing, an action we call “reactive-skip” (\mathbb{I}). Reactive skip has two behavioural modes: if started in an unstable state (i.e the prior computation is diverging), then all it guarantees is that the slots may get extended somehow; otherwise it stays stable, and leaves all other observations unchanged.

$$\begin{aligned} \mathbf{R1}(P) &\hat{=} P \wedge slots \preceq slots' \\ \mathbf{R3}(P) &\hat{=} \mathbb{I} \triangleleft wait \triangleright P \\ \mathbb{I} &\hat{=} \neg ok \wedge slots \preceq slots' \vee ok' \wedge wait' = wait \wedge state' = state \wedge slots' = slots \end{aligned}$$

The purpose of the $slots$ observation variable in slotted-*Circus*, and its trace analogue (tr) in UTP reactive-process theory, is to facilitate the definition of operators such as sequential composition. What is not permitted however, is for a process to be able to base its actions on the history of past events as recorded by this variable—any such “memory” of the past must be captured by the $state$ observation. Healthiness condition **R2** is concerned with ensuring that a process can only specify how the history is extended, without reference to what has already happened. In [HH98, Chp. 8] this is captured by stating that P is **R2**-healthy if it is invariant under an arbitrary shift in the prehistory, or in other words, a non-deterministic choice over all possible values that tr might take:

$$\begin{aligned} \mathbf{R2-UTP}(P) &\hat{=} \sqcap_s P[s, s \wedge (tr' - tr)/tr, tr'] \\ &\equiv \exists s \bullet P[s, s \wedge (tr' - tr)/tr, tr'] \end{aligned}$$

It would seem reasonable to expect the slotted-*Circus* version to simply replace tr by $slots$ and use the slot-sequence analogues of sequence concatenation and subtraction. This would result in the following definition (here the \mathbf{a} indicates “almost”):

$$\mathbf{R2a}(P) \hat{=} \exists ss \bullet P[ss, ss \# (slots' \setminus\setminus slots)/slots, slots'] \quad (5)$$

Whilst this looks plausible, there is in fact a problem with it, which only becomes apparent when we attempt to apply the definition later on in the semantics and then prove certain key desirable properties. Consider the predicate $slots' = slots$ which asserts that no events occur. This predicate should be **R2**-healthy, as it describes a process that chooses to do nothing, regardless of the value of $slots$. However calculation shows that

$$\mathbf{R2a}(slots' = slots) \equiv slots' \cong slots.$$

The equality gets weakened to the slot-sequence equivalence introduced earlier. An immediate consequence of this is that \mathbb{I} is not healthy by this definition, as calculation shows that the slot-equality is weakened to slot-equivalence (underlined below).

$$\mathbf{R2a}(\mathbb{I}) \equiv \neg ok \wedge slots \preceq slots' \vee ok' \wedge wait' = wait \wedge state' = state \wedge \underline{slots' \cong slots}$$

Original work explored keeping **R2a** as is, and redefining \mathbb{I} to be that version shown above. However this then weakened a number of key properties of \mathbb{I} , most notably to do with its role as an identity for sequential composition under appropriate circumstances.

The underlying problem with **R2a** has to do with the fact that in slotted-*Circus*, unlike UTP, we have refusals interleaved with events in $slots$, and slot-sequence operators that treat refusals, particularly the last, in a non-uniform way. The problem is that **R2a** weakens the predicate a

little too much, so we need to find a way to strengthen its result appropriately. The appropriate way to handle this issue has turned out to be to modify the definition of **R2** to require that we only quantify over *ss* values that happen to agree with *slots* on the very last refusal. This has no impact on predicates like \preceq and \cong which are not concerned with the last refusals, but provides just enough extra information to allow slot-sequence equality be considered as **R2**-healthy. The slightly strengthened version now reads:

$$\mathbf{R2}(P) \triangleq \exists ss \bullet P[ss, ss \# (slots' \searrow slots)/slots, slots'] \wedge Ref(last(slots)) = Ref(last(ss))$$

The proof that **R2** is idempotent is somewhat more involved than those for **R1** and **R3**. Calculations show that predicates $slots \preceq slots'$, $slots' \cong slots$, $slots' = slots$ (se Appendix B) and \mathbb{I} , are all **R2**-healthy. It also distributes through disjunction, which is very important.

It is worth pointing out that two versions of **R2** are presented in [HH98]. The second, which we shall call **R2'** is shown in an appendix:

$$\mathbf{R2}'(P) \triangleq P[\langle \rangle, tr' - tr/tr, tr']$$

Both **R2** and **R2'** have the same set of fixed points, so can be used interchangeably as a test for healthiness. However, if used to make a predicate healthy, then **R2** is more forgiving than **R2'**:

$$\begin{aligned} \mathbf{R2}(tr = \langle a \rangle \wedge tr' = \langle a, b \rangle) &\equiv (tr' - tr) = \langle b \rangle \\ \mathbf{R2}'(tr = \langle a \rangle \wedge tr' = \langle a, b \rangle) &\equiv \mathbf{false} \end{aligned}$$

This is an example of where overloading the notation **H** to stand for both **mkH** and **isH** can be misleading. We note that the version of **R2** used in [She06] is the CTA equivalent of **R2'**.

Reactive Healthiness A reactive slotted-*Circus* process is one that satisfies all three of the above healthiness conditions, so we define an overall condition **R** as their composition:

$$\mathbf{R} \triangleq \mathbf{R3} \circ \mathbf{R2} \circ \mathbf{R1} \tag{6}$$

In fact all three conditions commute with each other, so we re-order the above composition to suit.

4.2 CSP Healthiness

In addition to the reactive-healthiness just introduced, shared by a range of concurrent theories including ACP and CSP, there are a number of aspects of healthiness specific to CSP-like theories. In [HH98, Chp. 8] there are five of these presented, but for our purposes it suffices to consider only the first one.

A process is **CSP1** healthy if *all* it asserts, when started in an unstable state (due to some serious earlier failure), is that the event history may be extended:

$$\mathbf{CSP1}(P) \triangleq P \vee \neg ok \wedge slots \preceq slots' \tag{7}$$

5 Slotted Semantics

We are now in a position to give the semantics of the slotted-*Circus* language which is presented for completeness in Figures 5 & 6.

We shall not give a detailed commentary to all the definitions shown but instead will focus on some key points.

$$\begin{aligned}
\text{Chaos} &\hat{=} \mathbf{R}(\text{true}) \\
\text{Stop} &\hat{=} \mathbf{CSP1}(\mathbf{R3}(ok' \wedge wait' \wedge \text{EqvTrace}(\langle \rangle, slots' \searrow slots))) \\
b \&A &\hat{=} A \triangleleft b \triangleright \text{Stop} \\
\text{Skip} &\hat{=} \mathbf{R}(\exists ref \bullet ref = \text{EqvRef}(slots) \wedge \mathbf{I}) \\
\text{Wait } t &\hat{=} \mathbf{CSP1}(\mathbf{R}(ok' \wedge \text{delay}(t) \wedge \text{EqvTrace}(\langle \rangle, slots' \searrow slots))) \\
&\text{delay}(t) = (\#slots' - \#slots < t) \triangleleft wait' \triangleright (\#slots' - \#slots = t \wedge state' = state) \\
x := e &\hat{=} \mathbf{CSP1} \left(\mathbf{R} \left(\begin{array}{l} ok = ok' \wedge wait = wait' \wedge slots = slots' \\ \wedge state' = state \oplus \{x \mapsto val(e, state)\} \end{array} \right) \right) \\
&val : \text{Expr} \times (\text{Name} \rightarrow \text{Value}) \mapsto \text{Value} \\
c.e \rightarrow \text{Skip} &\hat{=} \mathbf{CSP1}(ok' \wedge \mathbf{R}(\text{wait_com}(c) \vee \text{complete_com}(c.e))) \\
&wait_com(c) = wait' \wedge \text{possible}(c)(slots, slots') \wedge \text{EqvTrace}(\langle \rangle, slots' \searrow slots) \\
&\text{possible}(c)(slots, slots') = c \notin \bigcup \text{Refs}(slots' - \text{front}(slots)) \\
&\text{term_com}(c.e) = \neg wait' \wedge \#slots = \#slots' \wedge \text{EqvTrace}(\langle c \rangle, slots' \searrow slots) \\
&\text{complete_com}(c.e) = \text{term_com}(c.e) \vee \text{wait_com}(c); \text{term_com}(c.e) \\
c!e \rightarrow \text{Skip} &\hat{=} c.e \rightarrow \text{Skip} \\
c?x \rightarrow \text{Skip} &\hat{=} \exists e \bullet (c.e \rightarrow \text{Skip}[state_0/state] \wedge state' = state_0 \oplus \{x \mapsto e\}) \\
comm \rightarrow A &\hat{=} (comm \rightarrow \text{Skip}); A \\
A \square B &\hat{=} \mathbf{CSP2}(\text{ExtChoice1}(A, B) \vee \text{ExtChoice2}(A, B)) \\
&\text{ExtChoice1}(A, B) \hat{=} A \wedge B \wedge \text{Stop} \\
&\text{ExtChoice2}(A, B) \hat{=} (A \vee B) \wedge \text{DifDetected}(A, B) \\
&\text{DifDetected}(A, B) \hat{=} \neg ok' \vee \left(\left(\begin{array}{l} (ok \wedge \neg wait) \wedge \\ \left(\left(\begin{array}{l} A \wedge B \wedge ok' \wedge \\ wait' \wedge slots = slots' \end{array} \right) \vee \right) \\ \text{Skip} \end{array} \right) \right); \left(\begin{array}{l} (ok' \wedge \neg wait' \wedge slots' = slots) \vee \\ (ok' \wedge \text{ImmEvs}(slots, slots')) \end{array} \right) \right) \\
&\text{ImmEvs}(slots, slots') \hat{=} \neg \text{EqvTrc}(\langle \rangle, \text{head}(slots' \searrow slots))
\end{aligned}$$

Fig. 5. Slotted-Circus Semantics (part I)

$$\begin{aligned}
& A \llbracket s_A \mid \{ \} \mid cs \mid s_B \rrbracket B \\
& \cong \exists obs_A, obs_B \bullet \\
& \quad A[obs_A/obs'] \wedge B[obs_B/obs'] \wedge \\
& \quad \left(\begin{array}{l} \text{if } \left(\begin{array}{l} s_A \triangleleft state_A \neq s_A \triangleleft state \vee \\ s_B \triangleleft state_B \neq s_B \triangleleft state \vee \\ s_A \cap s_B \neq \emptyset \end{array} \right) \\ \text{then } \neg ok' \wedge slots \preceq slots' \\ \text{else } \left(\begin{array}{l} ok' = ok_A \wedge ok_B \wedge \\ wait' = (wait_A \vee 1.wait_B) \wedge \\ state' = (s_B \triangleleft state_A) \oplus (s_A \triangleleft state_B) \wedge \\ ValidMerge(cs)(slots, slots', slots_A, slots_B) \end{array} \right) \end{array} \right) \\
& \quad ValidMerge : \mathbb{P} E \rightarrow ((\mathcal{S} E)^+)^4 \rightarrow \mathbb{B} \\
& \quad ValidMerge(cs)(s, s', s_0, s_1) = dif(s', s) \in TSync(cs)(dif(s_0, s), dif(s_1, s)) \\
& \quad TSync : \mathbb{P} E \rightarrow (\mathcal{S} E)^* \times (\mathcal{S} E)^* \rightarrow \mathbb{P}((\mathcal{S} E)^+) \\
& \quad TSync(cs)(s_1, s_2) = TSync(cs)(s_2, s_1) \\
& \quad TSync(cs)(\langle \rangle, \langle \rangle) = \{ \} \\
& \quad TSync(cs)(\langle s \rangle, \langle \rangle) = \{ \langle s' \rangle \mid s' \in SSync(cs)(s, SNull(Ref(s))) \} \\
& \quad TSync(cs) \left(\begin{array}{l} s_1 \circledast S_1, \\ s_2 \circledast S_2 \end{array} \right) = \left\{ \begin{array}{l} s' \circledast S' \\ \mid s' \in SSync(cs)(s_1, s_2) \wedge \\ S' \in TSync(cs)(S_1, S_2) \end{array} \right\} \\
& \quad A \setminus hidn \cong \mathbf{R} \left(\exists s \bullet A[s/slots'] \wedge \right. \\
& \quad \left. slots' \searrow slots = map(SHide(hidn))(dif(s, slot)) \right); Skip \\
& \mu X \bullet F(X) \cong \sqcap \{ X \mid F(X) \sqsubseteq X \}
\end{aligned}$$

Fig. 6. Slotted-Circus Semantics (II)

The *STOP* action refuses all events, but does allow the clock to keep ticking. Assignment and channel-communication take less than a clock-cycle, so we can sequence arbitrarily many in a time-slot. This does raise the possibility of Zeno processes (infinite events within a time-slot), so some care will be required here (disallowing infinite histories). This is more power than that required for synchronous hardware, where we expect these actions to synchronise with the clock, but we can model that by postfixing a *Wait* 1 statement, as used in the case study shown later. An important point to note is the definition of channel input ($c?x \rightarrow P$), not only involves an event $c.e$ for some e , but also updates the state. This is exploited later to allow shared variables.

The definition of external choice is quite complex —see [She06, p69] for a discussion.

We define slotted-parallel in a direct fashion, similar to that used for *Circus*, avoiding the complexities of the UTP/CTA approaches, and also handling error cases in passing. An error occurs in $P \llbracket s_A \mid C \mid s_B \rrbracket Q$ if P (Q) modifies any variable in s_B (s_A).

5.1 Laws

The language constructs displayed here obey a wide range of laws, many of which have been described elsewhere [HH98, WC01b, SH02, She06] for those constructs that slotted-*Circus* shares with other languages (e.g. non-deterministic choice, sequential composition, conditional, guards, *STOP*, *SKIP*). Here we simply indicate some of the laws regarding *Wait* that peculiar to slotted-*Circus* (Figure 7).

$$\begin{aligned}
& \textit{Wait } n; \textit{Wait } m = \textit{Wait } (m + n) \\
& \textit{Wait } n \square \textit{Wait } n + m = \textit{Wait } n \\
& (\textit{Wait } n; P) \square (\textit{Wait } n; Q) = \textit{Wait } n; (P \square Q) \\
& (\textit{Skip} \square (\textit{Wait } n; P)) = \textit{Skip}, \quad n > 0 \\
& (a \rightarrow P) \square (\textit{Wait } n; (a \rightarrow P)) = (a \rightarrow P)
\end{aligned}$$

Fig. 7. Laws of slotted-*Circus Wait*.

5.2 Links

In [HH98, §1.6, pp40–1], a general Galois connection between an abstract theory with observational variable a and a concrete theory over observation c is:

$$[(\exists c \bullet D(c) \wedge \ell(c, a)) \Rightarrow S(a)] \text{ iff } [D(c) \Rightarrow (\forall a \bullet \ell(c, a) \Rightarrow S(a))]$$

Here D and S are corresponding design (concrete) and specification (abstract) predicates respectively, while $\ell(c, a)$ is the linking predicate connecting observations at the two worlds. Of interest to us in the main are links between *Circus* (playing the role of the abstract theory with observations a) and various instantiations of slotted-*Circus* (concrete, with observations c). The difference between *Circus* and slotted-*Circus* is that the former has observations tr and ref , whilst the latter subsumes both into $slots$. However we can immediately exploit the method just presented by using the following relationship to define ℓ , which here relates the *Circus* observational variables to those of slotted-*Circus*:

$$\textit{EqvTrace}(tr, slots) \wedge ref = \textit{EqvRef}(slots) \tag{8}$$

So we get a Galois-link between *Circus* and any instantiation of slotted-*Circus* for free. Similarly, a given relationship between different \mathcal{H} types allows us to generate Galois-links between different slotted-*Circus* instantiations.

6 Instantiating Slotted-*Circus*

We now look at the issue of giving one or more concrete instantiations to the slotted-*Circus* framework just described. Originally, this work was aimed at producing a synchronous version of *Circus*, in which all events in a time-slot were to be considered as simultaneous. One motivation for this was to support the Handel-C language, which maps programs to synchronous hardware in which all variable updates are synchronised with a global clock edge marking the end of a computation cycle [Cel02]. However, there were two main difficulties with this approach.

The first was that the formal semantics developed for Handel-C outside of the UTP framework [BW02, BW05] actually modelled activity within a time-slot as a series of decision-making events spread out in time, all culminating in a set of simultaneous variable updates at the end of the slot. This approach, adopted in both the operational and denotational semantics, gives a very natural and intuitive description of what is taking place during Handel-C execution.

The second difficulty is more fundamental in nature, and exposed a key assumption underlying the UTP reactive theories, and those for CSP in general. Early work looked at the Ph.D thesis of Janet Barnes [Bar93] which introduced a synchronous version of CSP (SCSP). The key observation was a sequence of slots, each comprising two event sets, one denoting the events occurring in that slot (Acceptances) and the other describing the events refused (Refusals). A healthiness condition required that the acceptances and refusals in any slot be disjoint. However, implicit in this disjointedness condition is the notion that both the acceptances and refusals are truly simultaneous. However, in the failures of CSP, and the corresponding *tr* and *ref* observations of UTP, the key interpretation involved is that the refusals describe what is being refused given that the event history has just taken place. As a specific example, consider the process $a \rightarrow b \rightarrow P$. A possible (failure) observation of this process is $(\langle a \rangle, \{a\})$, *i.e.*, we have observed the occurrence of the a event and the fact that the process is now refusing to perform an a .

Consider trying to instantiate a slot where the history is simply an event-set, as per SCSP:

$$\begin{aligned} A \in \text{SCSP } E &\hat{=} \mathbb{P} E \\ \text{HNull}_{\text{SCSP}} &\hat{=} \emptyset \\ \text{Hadd}_{\text{SCSP}}(A_1, A_2) &\hat{=} A_1 \cup A_2 \\ \text{Hsub}_{\text{SCSP}}(A_1, A_2) &\hat{=} A_1 \setminus A_2 \\ &\dots \end{aligned}$$

We find that we cannot guarantee law [Sadd:unit] (Appendix A), even if the SCSP invariant is not required. This property is required to demonstrate that $\text{slots} \cong \text{slots}'$ is **R2**-healthy. The underlying problem is that the definition of **R2** relies on being able to deduce that slots is empty if subtracting slots' from slots leaves slots' unchanged. However at the history-as-set level, we cannot deduce $H = \emptyset$, given that $H' \setminus H = H'$.

6.1 Multiset History Instantiation

We can define an instantiation where the event history is a multiset or bag of events (\mathcal{H}_{MSA}), so event ordering is unimportant, but multiple event occurrences in a slot do matter (Figure 8). The bag notation used here is that of Z [Spi87]. The events accepted are simply the bag domain. A

$$\begin{aligned}
\mathcal{H}_{\mathcal{MSA}} E &\hat{=} E \leftrightarrow \mathbb{N}_1 \\
Acc(bag) &\hat{=} dom(bag) \\
EqvTrc(tr, bag) &\hat{=} items(tr) = bag \\
HNull &\hat{=} [] \\
bag_1 \preceq bag_2 &\hat{=} bag_1 \sqsubseteq bag_2 \\
Hadd(bag_1, bag_2) &\hat{=} bag_1 \oplus bag_2 \\
Hsub(bag_1, bag_2) &\hat{=} bag_1 \ominus bag_2 \\
HSync(cs)(bag_1, bag_2) &\hat{=} \{(cs \triangleleft (bag_1 \oplus bag_2)) \oplus (cs \triangleleft (bag_1 \cap bag_2))\} \\
\textbf{where } &\cap \text{ is bag interesection} \\
HHide(hdn)bag &\hat{=} hdn \triangleleft bag
\end{aligned}$$

Fig. 8. Multiset Action Instantiation (\mathcal{MSA})

trace corresponds to a bag if it contains the same number of events as that bag. A null history is simply an empty bag. A bag is a prefix if smaller than another bag. History addition and subtract are the bag equivalents. History synchronisation merges the parts of the two bags disjoint from the synchronisation set, with the intersection of all three. Hiding is modelled by bag restriction.

The proofs that the above instantiation satisfy the properties in Appendix A are all straightforward. The proof of law [ET:pf \times] for \mathcal{MSA} is shown in Appendix B.

7 Example Circus process

We illustrate slotted *Circus* using an example originally due to Hoare [Hoa85a]. The problem is to compute the weighted sums of consecutive pairs of inputs. Suppose that the input stream contains the following values: $x_0, x_1, x_2, x_3, x_4, \dots$; then the output stream will be

$$(a * x_0 + b * x_1), (a * x_1 + b * x_2), (a * x_2 + b * x_3), (a * x_3 + b * x_4), \dots$$

for weights a and b . We specify this problem with a synchronous process with two channels: *left*, used for input, and *right* used for output. Since each output requires two consecutive values from the input stream, the first output cannot occur before the third clock cycle.

<i>clock</i>	0	1	2	3	4	5	
<i>left</i>	x_0	x_1	x_2	x_3	x_4	x_5	\dots
<i>right</i>			$a * x_0 + b * x_1$	$a * x_1 + b * x_2$	$a * x_2 + b * x_3$	$a * x_3 + b * x_4$	\dots

Hoare's solution performs the two multiplications in parallel and then adds the results. Suppose the implementation technology is a single field-programmable gate array; the circuitry for the computation of the output would then be inherently parallel anyway. Let's assume instead that we want to implement the two multiplications on separate FPGAs. It's clear that the a -product is always ready one clock cycle before we need to perform the addition. Let's keep this intermediate result in the variable m :

<i>clock</i>	0	1	2	3	4	5	
<i>left</i>	x_0	x_1	x_2	x_3	x_4	x_5	\dots
m		$a * x_0$	$a * x_1$	$a * x_2$	$a * x_3$	$a * x_4$	\dots
<i>right</i>			$m + b * x_1$	$m + b * x_2$	$m + b * x_3$	$m + b * x_4$	\dots

First however, note we are going to target a Handel-C-like scenario where channel communication and assignment take one-clock cycle, and we have shared variables. We need to reason about interleavings of assignments, but rather than re-work the whole theory to have state-sequences, we simply convert assignments into channel communications. So for the following case study, we have the following shorthands:

shorthand	expansion
$c?_1x$	$c?x \rightarrow \text{Wait } 1.$
$c!_1x$	$c!x \rightarrow \text{Wait } 1.$
$x :=_1 e$	$(a!_1e \parallel [\emptyset \mid a \mid x] \parallel a?_1x)$ where a is fresh.
δP	variables modified by P i.e used in $x := \dots$ or $c?x$
$P \parallel Q$	$P \parallel [\delta P \mid \emptyset \mid \delta Q] Q$

In effect the clock-cycle wait is built into the communication and assignment notations, effectively avoid any Zeno hazards. Now we're ready to specify the problem as a slotted *Circus* process.

$$WS \hat{=} \text{var } x, m : \mathbb{N} \bullet (\text{left?}_1x ; (\text{left?}_1x \parallel m :=_1 a * x); \\ (\mu X \bullet (\text{left?}_1x \parallel m :=_1 a * x \parallel \text{right!}_1(m + b * x)); X))$$

The process WS is clearly deadlock and livelock free: it is a non-stopping process with no internal synchronisations; and it is hiding and chaos-free, with guarded recursion. Now we need to decompose WS into two parallel processes with encapsulated state. We can replace the use of m by a channel communication that passes the intermediate value. One process (WSL) will receive the input stream and compute the a -product; the other (WSR) will compute the b -product and the sum, and generate the output stream. But now we see a problem with WS . The value x_1 is received by WSL in the first clock cycle, and so it can be communicated to WSR in the second cycle. So it can't be used by WSR until the third clock cycle. So we need to delay the output on the *right* by another clock cycle. Our timing diagram shows this more clearly.

clock	0	1	2	3	4	5	
<i>left</i>	x_0	x_1	x_2	x_3	x_4	x_5	\dots
<i>w</i>		x_0	x_1	x_2	x_3	x_4	\dots
<i>m</i>			$a * x_0$	$a * x_1$	$a * x_2$	$a * x_3$	\dots
<i>right</i>				$m + b * x_1$	$m + b * x_2$	$m + b * x_3$	\dots

Here's another version of WS that does this.

$$WS' \hat{=} \text{var } w, x, m : \mathbb{N} \bullet \\ \text{left?}_1x ; (\text{left?}_1x \parallel w :=_1 x); \\ (\text{left?}_1x \parallel w :=_1 x \parallel m :=_1 a * w); \\ (\mu X \bullet (\text{left?}_1x \parallel w :=_1 x \parallel m :=_1 a * x \parallel \text{right!}_1(m + b * w)); X)$$

Our refinement strategy is to split into two processes. The variable x belongs in WSL , since it is used to store the current input. The variable m can be placed in WSR , since it is used directly in producing outputs, but its value must be computed in WSL , and so the value will have to be communicated from left to right. The variable w records the previous input, and this is used in both left and right processes; so we duplicate its value using a ghost variable v . The ghost variable can then be used in the right-hand process in the calculation of the output on the right. Our refinement starts with organising the variables. (To reduce clutter, we abbreviate left?_1x by $?_1x$

and $right!_1 e$ by $!_1 e$. We also separate the beginning and end of variable scopes.)

```

var w, x, m ;
  ?1x ; (?1x ||| w :=1 x) ; (?1x ||| w, m :=1 x, a * w) ;
  (μ X • (?1x ||| w, m :=1 x, a * w ||| !_1(m + b * w)) ; X) ;
end w, x, m
= { v ghosts w }
var w, x, m ;
  ?1x ; (?1x ||| w :=1 x) ;
  var v ;
    (?1x ||| v, w, m :=1 x, x, a * w) ;
    (μ X • (?1x ||| v, w, m :=1 x, x, a * w ||| !_1(m + b * v)) ; X) ;
  end v ;
end w, x, m
= { widen scope }
var v, w, x, m ;
  ?1x ; (?1x ||| w :=1 x) ; (?1x ||| v, w, m :=1 x, x, a * w) ;
  (μ X • (?1x ||| v, w, m :=1 x, x, a * w ||| !_1(m + b * v)) ; X) ;
end v, w, x, m

```

Our next step is to insert some hidden events to prepare for the communication of values between the two processes. We add two hidden channels: c communicates x 's value; and mid communicates m 's value. These events are not needed in the first two steps.

```

= { hiding }
( var v, w, x, m ;
  ?1x ; (?1x ||| w :=1 x) ; (?1x ||| c.x ||| mid.a * w ||| v, w, m :=1 x, x, a * w) ;
  (μ X • (?1x ||| c.x ||| mid.a * w ||| v, w, m :=1 x, x, a * w ||| !_1(m + b * v)) ; X) ;
end v, w, x, m ) \ {c, mid}

```

Now we can prepare for the parallel split by organising each step into parallel parts, examining each atomic action and assigning it to the left or right component. The right-hand process doesn't need to do anything during the first two steps, so we make it wait. In the third step, the the input belongs on the left. The pair of actions $(c.x ||| v :=_1 x)$ can be replaced by a communication: the left performs the output $c!_1 x$ and the right performs the input $c?_1 v$. Similarly, $(mid.a * w ||| m :=_1 a * w)$ can be replaced by $mid!_1 a * w$ and $mid?_1 m$. Finally, the assignment to w belongs on the left. The body of the recursion is split in exactly the same way, with the addition of the output being assigned to the right-hand process.

$$\begin{aligned}
&= \{ \textit{parallel} \} \\
&\quad (\textit{var } v, w, x, m ; \\
&\quad\quad (?_1 x \parallel \textit{Wait } 1) ; \\
&\quad\quad ((?_1 x \parallel w :=_1 x) \parallel \textit{Wait } 1) ; \\
&\quad\quad ((?_1 x \parallel c!_1 x \parallel \textit{mid}!_1 a * w \parallel w :=_1 x) \parallel (c?_1 v \parallel \textit{mid}?_1 m)) ; \\
&\quad\quad (\mu X \bullet ((?_1 x \parallel c!_1 x \parallel \textit{mid}!_1 a * w \parallel w :=_1 x) \\
&\quad\quad\quad \parallel (c?_1 v \parallel \textit{mid}?_1 m \parallel !_1(m + b * v)) ; X)) ; \\
&\quad \textit{end } v, w, x, m) \setminus \{c, \textit{mid}\}
\end{aligned}$$

We also need to split the recursion into two parallel parts. Since the body comprises two atomic steps in parallel, the fixed-point operator distributes cleanly through the parallel operator.

$$\begin{aligned}
&= \{ \textit{parallel} \} \\
&\quad (\textit{var } v, w, x, m ; \\
&\quad\quad (?_1 x \parallel \textit{Wait } 1) ; \\
&\quad\quad ((?_1 x \parallel w :=_1 x) \parallel \textit{Wait } 1) ; \\
&\quad\quad ((?_1 x \parallel c!_1 x \parallel \textit{mid}!_1 a * w \parallel w :=_1 x) \parallel (c?_1 v \parallel \textit{mid}?_1 m)) ; \\
&\quad\quad ((\mu X \bullet (?_1 x \parallel c!_1 x \parallel \textit{mid}!_1 a * w \parallel w :=_1 x) ; X) ; \\
&\quad\quad \parallel (\mu X \bullet (c?_1 v \parallel \textit{mid}?_1 m \parallel !_1(m + b * v)) ; X)) ; \\
&\quad \textit{end } v, w, x, m) \setminus \{c, \textit{mid}\}
\end{aligned}$$

Now we can perform the parallel split, using an interchange law for sequence and parallel that is similar to the spreadsheet rules in UTP. We create the left-hand process by encapsulating w and x , retaining the left-hand parts, and discarding the right-hand parts. We create the right-hand process similarly.

$$\begin{aligned}
&= \{ \textit{parallel split} \} \\
&\quad ((\textit{var } w, x ; \\
&\quad\quad ?_1 x ; \\
&\quad\quad (?_1 x \parallel w :=_1 x) ; \\
&\quad\quad (?_1 x \parallel c!_1 x \parallel \textit{mid}!_1 a * w \parallel w :=_1 x) ; \\
&\quad\quad (\mu X \bullet (?_1 x \parallel c!_1 x \parallel \textit{mid}!_1 a * w \parallel w :=_1 x) ; X) ; \\
&\quad \textit{end } w, x) \\
&\quad \parallel \\
&\quad (\textit{var } v, m ; \\
&\quad\quad \textit{Wait } 1 ; \textit{Wait } 1 ; \\
&\quad\quad (c?_1 v \parallel \textit{mid}?_1 m) ; \\
&\quad\quad (\mu X \bullet (c?_1 v \parallel \textit{mid}?_1 m \parallel !_1(m + b * v)) ; X) ; \\
&\quad \textit{end } v, m) \\
&\quad) \setminus \{c, \textit{mid}\}
\end{aligned}$$

Now we can tidy up the processes for our final result.

$$\begin{aligned}
&(\textit{var } w, x : \mathbb{N} \bullet \\
&\quad \textit{left}?_1 x ; (\textit{left}?_1 x \parallel w :=_1 x) ; (\textit{left}?_1 x \parallel c!_1 x \parallel \textit{mid}!_1 a * w \parallel w :=_1 x) ; \\
&\quad (\mu X \bullet (\textit{left}?_1 x \parallel c!_1 x \parallel \textit{mid}!_1 a * w \parallel w :=_1 x) ; X) ; \\
&\parallel \\
&\quad \textit{var } v, m : \mathbb{N} \bullet (\textit{Wait } 2 ; (c?_1 v \parallel \textit{mid}?_1 m) ; \\
&\quad\quad (\mu X \bullet (c?_1 v \parallel \textit{mid}?_1 m \parallel \textit{right}!_1(m + b * v)) ; X) ;) \\
&\quad) \setminus \{c, \textit{mid}\}
\end{aligned}$$

Of course, since this is equal to WS , it is deadlock and livelock-free and computes the right results.

A key point of the above case-study is that it works in any of the instantiations mentioned so far for slotted-*Circus*, namely CTA or MSA .

8 Related Work

During the development of Handel-C at Oxford, a lot of the principles and theory was developed and published [PL91, HIJ93]. Here the emphasis was very much on the verified compilation into hardware of an *occam*-like language. However with the commercialisation of this as the language Handel-C the formal aspects and hardware compilation parted company, and the Handel-C language acquired new constructs like “prialt” that were not treated in the literature.

Modern Handel-C [Cel02] also has the idea of connecting hardware with different clocks together using tightly controlled asynchronous interfaces. Modelling this kind of behaviour requires a theory that mixes time and asynchronicity, such as timed-CSP [Sch00].

There has been work done on hardware semantics, ranging from the “reFLect” language used by Intel for hardware verification [GJ06], to the language Esterel used mainly for the development of flight avionics [BG92]. The Intel approach is a suite of hardware description languages, model-checkers and theorem provers all written and/or integrated together using the reFLect language, aimed mainly at the verification of computer datapath hardware. The Esterel language is a hardware description language with a formal semantics, and so is quite low-level in character, and so in the context of this research could be considered a potential replacement of Handel-C as an implementation technology. However, it is unclear how well it would link to the kind of specification and refinement style of work that we are proposing to support.

9 Future Work

We have described a generic framework for instantiating a wide range of slotted-theories, capturing their common features. An important aspect that has yet to be covered is what distinguishes the various instantiations from one another, i.e. how do the laws of CTA differ from those of MSA , for instance. We know for example that the following is a law of MSA , but not of CTA , or slotted-*Circus* in general:

$$a \rightarrow b \rightarrow P = b \rightarrow a \rightarrow P$$

Also worthy of exploration are the details of the behaviour of the Galois links inbetween different instances of slotted-*Circus*, and between those and standard *Circus*. These details will provide a framework for a comprehensive refinement calculus linking all these reactive theories together.

In order to deal with the asynchronously interfaced multiple-clock hardware now supported by Handel-C we will need to exploit the link from the slotted theories to the generally asynchronous *Circus* theory itself.

Also of interest will be to consider to what extent the work on “generic composition” [Che02, Che06] can contribute to a clear and or tractable presentation of this theory.

10 Conclusions

A framework for giving UTP semantics to a class of reactive systems whose execution is demarcated by regular clock ticks has been presented. The general nature of the observational variables and the key operations on same have been discussed, showing how they are used build to both the healthiness conditions and the language semantics. A key result of this work has been the care needed to get a satisfactory definition of **R2**, and exposing the fact that certain synchronous theories like SCSP do not fit this particular UTP pattern for describing reactive systems.

10.1 Acknowledgement

We would like to thank the Dean of Research at TCD and QinetiQ for their support of this work, and the comments of the anonymous reviewers, which helped improve key material in this paper.

References

- [Bar93] Janet E. Barnes. A Mathematical Theory of Synchronous Communication. Technical Monograph PRG-112, Oxford University Computing Laboratory Programming Research Group, Hilary Term 1993.
- [BG92] G. Berry and G. Gonthier. The ESTEREL synchronous programming language: design, semantics, implementation. *Science of Computer Programming*, 19:87–152, 1992.
- [BW02] Andrew Butterfield and Jim Woodcock. Semantic domains for handel-C. *Electr. Notes Theor. Comput. Sci.*, 74, 2002.
- [BW05] Andrew Butterfield and Jim Woodcock. `primalt` in Handel-C: an operational semantics. *International Journal on Software Tools for Technology Transfer (STTT)*, 7(3):248–267, June 2005.
- [Cel02] Celoxica Ltd. *Handel-C Language Reference Manual, v3.0*, 2002. URL: www.celoxica.com.
- [Che02] Yifeng Chen. Generic composition. *Formal Asp. Comput.*, 14(2):108–122, 2002.
- [Che06] Yifeng Chen. Hierarchical organisation of predicate-semantic models. In Steve Dunne and Bill Stoddart, editors, *UTP*, volume 4010 of *Lecture Notes in Computer Science*, pages 155–172. Springer, 2006.
- [GJ06] Melham T. Grundy, J. and O’Leary J. A reflective functional language for hardware design and theorem proving. *Journal of Functional Programming*, 16(2):157–196, 2006.
- [HH98] C. A. R. Hoare and Jifeng He. *Unifying Theories of Programming*. Series in Computer Science. Prentice Hall, 1998.
- [HIJ93] H. Jifeng, I. Page, and J. Bowen. Towards a provably correct hardware implementation of Occam. In G.J. Milne and L. Pierre, editors, *Correct Hardware Design and Verification Methods*, volume 683 of *Lecture Notes in Computer Science*, pages 214–225, Arles, France, May 1993. IFIP WG10.2, Springer-Verlag.
- [Hoa85a] C. A. R. Hoare. *Communicating Sequential Processes*. Intl. Series in Computer Science. Prentice Hall, 1985.
- [Hoa85b] C. A. R. Hoare. Programs are predicates. In *Proc. of a discussion meeting of the Royal Society of London on Mathematical logic and programming languages*, pages 141–155, Upper Saddle River, NJ, USA, 1985. Prentice-Hall, Inc.
- [PL91] I. Page and W. Luk. Compiling Occam into field-programmable gate arrays. In W. Moore and W. Luk, editors, *FPGAs, Oxford Workshop on Field Programmable Logic and Applications*, pages 271–283, 15 Harcourt Way, Abingdon OX14 1NV, UK, 1991. Abingdon EE&CS Books.
- [Ros97] A. W. Roscoe. *The Theory and Practice of Concurrency*. international series in computer science. Prentice Hall, 1997.
- [Sch00] Steve Schneider. *Concurrent and Real-time Systems — The CSP Approach*. Wiley, 2000.
- [SH02] Adnan Sherif and Jifeng He. Towards a time model for circus. In Chris George and Huaikou Miao, editors, *ICFEM*, volume 2495 of *Lecture Notes in Computer Science*, pages 613–624. Springer, 2002.
- [She06] Adnan Sherif. *A Framework for Specification and Validation of Real Time Systems using Circus Action*. Ph.d. thesis, Universidade Federal de Pernambuco, Recife, Brazil, Jan 2006.
- [Spi87] Spivey. *The Z Notation: A Reference Manual*. Prentice Hall, 1987.
- [WC01a] J. C. P. Woodcock and A. L. C. Cavalcanti. A Concurrent Language for Refinement. In A. Butterfield and C. Pahl, editors, *IWFM’01: 5th Irish Workshop in Formal Methods*, BCS Electronic Workshops in Computing, Dublin, Ireland, July 2001.
- [WC01b] Jim Woodcock and Ana Cavalcanti. *Circus: a concurrent refinement language*. Technical report, University of Kent at Canterbury, October 2001.
- [WC02] Jim Woodcock and Ana Cavalcanti. The semantics of circus. In *ZB*, pages 184–203, 2002.

A Generic Laws

The functions and relations over $\mathcal{H} E$ required to define a slotted-*Circus* theory, need to satisfy the following laws:

- [ET:elems] $EqvTrc(tr, hist) \Rightarrow elems(tr) = Acc(hist)$
- [HIST:eq] $(h_1 = h_2) \equiv \forall tr \bullet EqvTrc(tr, h_1) \equiv EqvTrc(tr, h_2)$
- [HN:null] $Acc(HNull) = \{\}$
- [pfx:refl] $hist \preceq hist = \text{TRUE}$
- [pfx:trans] $hist_1 \preceq hist_2 \wedge hist_2 \preceq hist_3 \Rightarrow hist_1 \preceq hist_3$
- [pfx:anti-sym] $hist_1 \preceq hist_2 \wedge hist_2 \preceq hist_1 \Rightarrow hist_1 = hist_2$
- [SN:pfx] $HNull \preceq hist$
- [ET:pfx] $hist_1 \preceq hist_2 \Rightarrow \exists tr_1, tr_2 \bullet EqvTrc(tr_1, hist_1) \wedge EqvTrc(tr_2, hist_2) \wedge tr_1 \leq tr_2$
- [Sadd:events] $Acc(Sadd(h_1, h_2)) = Acc(h_1) \cup Acc(h_2)$
- [Sadd:unit] $Sadd(h_1, h_2) = h_1 \equiv h_2 = HNull$
- [Sadd:assoc] $Sadd(h_1, Sadd(h_2, h_3)) = Sadd(Sadd(h_1, h_2), h_3)$
- [Sadd:prefix] $h \preceq Sadd(h, h')$
- [Sub:pre] $pre Ssub(h_1, h_2) = h_2 \preceq h_1$
- [Sub:events] $h_2 \preceq h_1 \wedge h' = Ssub(h_1, h_2) \Rightarrow$
 $Acc(h_1) \setminus Acc(h_2) \subseteq Acc(h') \subseteq Acc(h_1)$
- [SSub:self] $Ssub(h, h) = HNull$
- [SSub:nil] $Ssub(h, HNull) = h$
- [SSub:same] $hist \preceq hist'_a \wedge hist \preceq hist'_b \Rightarrow$
 $Ssub(hist'_a, hist) = Ssub(hist'_b, hist) \equiv hist'_a = hist'_b$
- [SSub:subsub] $hist_c \preceq hist_a \wedge hist_c \preceq hist_b \wedge hist_b \preceq hist_a$
 $\Rightarrow Ssub(Ssub(hist_a, hist_c), Ssub(hist_b, hist_c)) = Ssub(hist_a, hist_b)$
- [Sadd:Ssub] $hist \preceq hist' \Rightarrow Sadd(hist, Ssub(hist', hist)) = hist'$
- [Ssub:Sadd] $Ssub(Sadd(h_1, h_2), h_1) = h_2$
- [SHid:evts] $Acc(SHide(hid)(h)) = Acc(h) \setminus hid$
- [SNC:sym] $SSync(cs)(h_1, h_2) = SSync(cs)(h_2, h_1)$
- [SNC:one] $\forall h' \in SSync(cs)(h_1, HNull) \bullet Acc(h') \subseteq Acc(h_1) \setminus cs$
- [SNC:only] $h' \in Acc(SSync(cs)(h_1, h_2)) \Rightarrow Acc(h') \subseteq Acc(h_1) \cup Acc(h_2)$
- [SNC:sync] $h' \in Acc(SSync(cs)(h_1, h_2)) \Rightarrow cs \cap Acc(h') \subseteq cs \cap (Acc(h_1) \cap Acc(h_2))$
- [SNC:assoc] $SyncSet(cs)(h_1)(SSync(cs)(h_2, h_3)) = SyncSet(cs)(h_3)(SSync(cs)(h_1, h_2))$

B Proofs for R2-ness of = and \mathcal{MSA} prefix

$$\begin{aligned}
& \mathbf{R2}(slots' = slots) \\
\equiv & \text{ “ defn. } \mathbf{R2}, \text{ apply substitution, shorthand } RL(s) = Ref(last(s)) \text{ ”} \\
& \exists ss \bullet ss \# (slots' \searrow slots) = ss \wedge RL(slots) = RL(ss) \\
\equiv & \text{ “ Property 1 (below) ”} \\
& \exists ss \bullet slots' \searrow slots = \langle SNull(RL(ss)) \rangle \wedge RL(slots) = RL(ss) \\
\equiv & \text{ “ Property 2 (below) ”} \\
& \exists ss \bullet front(slots') = front(slots) \wedge tail(slots').1 = tail(slots).1 \\
& \quad \wedge RL(slots') = RL(ss) \wedge RL(slots) = RL(ss) \\
\equiv & \text{ “ Liebniz, restrict quantification scope ”} \\
& front(slots') = front(slots) \wedge tail(slots').1 = tail(slots).1 \\
& \quad \wedge RL(slots') = RL(slots) \wedge \exists ss \bullet RL(slots) = RL(ss) \\
\equiv & \text{ “ defn. of equality, witness } ss = slots \text{ ”} \\
& slots = slots' \\
\text{Property 1: } & (ss \# tt = ss) \equiv tt = \langle SNull(RL(ss)) \rangle \\
\text{Property 2: } & (tt' \searrow tt) = \langle SNull(r) \rangle \equiv front(tt) = front(tt') \wedge last(tt).1 = last(tt').1 \wedge RL(tt') = r
\end{aligned}$$

$$\begin{aligned}
& bag_1 \preceq bag_2 \\
\equiv & \text{ “ defn. of prefix ”} \\
& bag_1 \sqsubseteq bag_2 \\
\equiv & \text{ “ bag property ”} \\
& \exists bag_\Delta \bullet bag_2 = bag_1 \oplus bag_\Delta \\
\equiv & \text{ “ bag property: } \forall bag \bullet \exists tr \bullet items(tr) = bag \text{ ”} \\
& \exists bag_\Delta, tr_\Delta, tr_1, \bullet bag_2 = bag_1 \oplus bag_\Delta \wedge items(tr_\Delta) = bag_\Delta \wedge items(tr_1) = bag_1 \\
\equiv & \text{ “ One-point rule backwards } tr_2 = tr_1 \hat{\wedge} tr_\Delta \text{ ”} \\
& \exists bag_\Delta, tr_\Delta, tr_1, tr_2 \bullet bag_2 = bag_1 \oplus bag_\Delta \wedge items(tr_\Delta) = bag_\Delta \wedge items(tr_1) = bag_1 \wedge tr_2 = tr_1 \hat{\wedge} tr_\Delta \\
\equiv & \text{ “ One-point rule } bag_\Delta, \text{ Liebniz } bag_1 \text{ ”} \\
& \exists tr_\Delta, tr_1, tr_2 \bullet bag_2 = items(tr_1) \oplus items(tr_\Delta) \wedge items(tr_1) = bag_1 \wedge tr_2 = tr_1 \hat{\wedge} tr_\Delta \\
\equiv & \text{ “ items is a sequence homomorphism ”} \\
& \exists tr_\Delta, tr_1, tr_2 \bullet bag_2 = items(tr_2) \wedge bag_1 = items(tr_1) \wedge tr_2 = tr_1 \hat{\wedge} tr_\Delta \\
\equiv & \text{ “ sequence property ”} \\
& \exists tr_\Delta, tr_1, tr_2 \bullet bag_2 = items(tr_2) \wedge bag_1 = items(tr_1) \wedge tr_\Delta = tr_2 - tr_1 \\
\equiv & \text{ “ One point rule: } tr_\Delta, \text{ requires definedness of } tr_2 - tr_1 \text{ ”} \\
& \exists tr_1, tr_2 \bullet bag_2 = items(tr_2) \wedge bag_1 = items(tr_1) \wedge tr_1 \leq tr_2 \\
\equiv & \text{ “ def. of } EquTrc, \text{ backwards ”} \\
& \exists tr_1, tr_2 \bullet EquTrc(tr_2, bag_2) \wedge EquTrc(tr_1, bag_1) \wedge tr_1 \leq tr_2
\end{aligned}$$