

Isabelle/*Circus* : a Process Specification and Verification Environment

Abderrahmane Feliachi, Marie-Claude Gaudel and Burkhart Wolff

¹ Univ. Paris-Sud, Laboratoire LRI, UMR8623, Orsay, F-91405, France

² CNRS, Orsay, F-91405, France

{Abderrahmane.Feliachi, Marie-Claude.Gaudel, Burkhart.Wolff}@lri.fr

Abstract. The *Circus* specification language combines elements for complex data and behavior specifications, using an integration of Z and CSP with a refinement calculus. Its semantics is based on Hoare and He’s unifying theories of programming (UTP).

We develop a machine-checked, formal semantics based on a “shallow embedding” of *Circus* in Isabelle/UTP (our semantic theory of UTP based on Isabelle/HOL). We derive proof rules from this semantics and implement tactic support that finally allows for proofs of refinement for *Circus* processes (involving both data and behavioral aspects).

This proof environment supports a syntax for the semantic definitions which is close to textbook presentations of *Circus*.

Keywords: *Circus*, denotational semantics, Isabelle/HOL, Process Algebras, Refinement

1 Introduction

Many systems involve both complex (sometimes infinite) data structures and interactions between concurrent processes. Refinement of abstract specifications of such systems into more concrete ones, requires an appropriate formalisation of refinement and appropriate proof support.

There are several combinations of process-oriented modeling languages with data-oriented specification formalisms such as Z or B or CASL; examples are discussed in [3, 10, 17, 14]. In this paper, we consider *Circus* [18], a language for refinement, that supports modeling of high-level specifications, designs, and concrete programs. It is representative of a class of languages that provide facilities to model data types, using a predicate-based notation, and patterns of interactions, without imposing architectural restrictions. It is this feature that makes it suitable for reasoning about both abstract and low-level designs.

We present a “shallow embedding” of the *Circus* semantics enabling state variables and channels in *Circus* to have arbitrary HOL types. Therefore, the entire handling of typing can be completely shifted to the (efficiently implemented) Isabelle type-checker and is therefore implicit in proofs. This drastically simplifies definitions and proofs, and makes the reuse of standardized proof procedures

possible. Compared to implementations based on a “deep embedding” such as [19] this significantly improves the usability of the resulting proof environment.

Our representation brings particular technical challenges and contributions concerning some important notions about variables. The main challenge was to represent alphabets and bindings in a typed way that preserves the semantics and improves deduction. We provide a representation of bindings without an explicit management of alphabets. However, the representation of some core concepts in the unifying theories of programming (UTP) and *Circus* constructs (variable scopes and renaming) became challenging. Thus, we propose a (stack-based) solution that allows the coding of state variables scoping with no need for renaming. This solution is even a contribution to the UTP theory that does not allow nested variable scoping. Some challenging and tricky definitions (e.g. channels and name sets) are explained in this paper.

This paper is organized as follows. The next section gives an introduction to the basics of our work: Isabelle/HOL, UTP and *Circus* with a short example of a *Circus* process. In Section 3, we present our embedding of the basic concepts of *Circus* (alphabet, variables ...). We introduce the representation of some *Circus* actions and process, with an overview of the Isabelle/*Circus* syntax. In Section 4, we show on an example, how Isabelle/*Circus* can be used to write specifications. We give some details on what is happening “behind the scenes” when the system parses each part of the specification. In the last part of this section, we show how to write proofs based on specifications, and give a refinement proof example. A more developed version of this paper can be found in [9].

2 Background

2.1 Isabelle, HOL and Isabelle/HOL

Isabelle [12] is a generic theorem prover implemented in SML. It is based on the so-called “LCF-style architecture”, which makes it possible to extend a small trusted logical kernel by user-programmed procedures in a logically safe way. New object logics can be introduced to Isabelle by specifying their syntax and semantics, by deriving its inference rules from there and program specific tactic support for the object logic. Isabelle is based on a typed λ -calculus including a Haskell-style type-system with type-classes (e.g. in $\alpha :: \text{order}$, the type-variable ranges over all types that posses a partial ordering.)

Higher-order logic (HOL) [7, 1] is a classical logic based on a simple type system. It provides the usual logical connectives like $_ \wedge _$, $_ \Rightarrow _$, $\neg _$ as well as the object-logical quantifiers $\forall x \bullet P x$ and $\exists x \bullet P x$; in contrast to first-order logic, quantifiers may range over arbitrary types, including total functions $f : \alpha \Rightarrow \beta$. HOL is centered around extensional equality $_ = _ : \alpha \Rightarrow \alpha \Rightarrow \text{bool}$. HOL is more expressive than first-order logic, since, e. g., induction schemes can be expressed inside the logic. Being based on some polymorphically typed λ -calculus, HOL can be viewed as a combination of a programming language like SML or

Haskell and a specification language providing powerful logical quantifiers ranging over elementary and function types.

Isabelle/HOL is an instance of Isabelle with higher-order logic. It provides a rich collection of library theories like sets, pairs, relations, partial functions lists, multi-sets, orderings, and various arithmetic theories which only contain rules derived from conservative, i. e. logically safe definitions. Setups for the automated proof procedures like `simp`, `auto`, and arithmetic types such as `int` are provided.

2.2 Advanced Specification Constructs in Isabelle/HOL

Constant definitions. In its easiest form, constant definitions are definitional logical axioms of the form $c \equiv E$ where c is a fresh constant symbol not occurring in E which is closed (both wrt. variables and type variables). For example:

```
definition upd :: ( $\alpha \Rightarrow \beta$ )  $\Rightarrow$   $\alpha \Rightarrow \beta \Rightarrow (\alpha \Rightarrow \beta)$     ("_(|_ := _)")
where      upd f x v  $\equiv$   $\lambda z.$  if x=z then v else f z
```

The pragma ("`_(|_ := _)`") for the Isabelle syntax engine introduces the notation `f(x:=y)` for `upd f x y`. Moreover, some elaborate preprocessing allows for recursive definitions, provided that a termination ordering can be established. Such recursive definitions are thus internally reduced to definitional axioms.

Type definitions. Types can be introduced in Isabelle/HOL in different ways. The most general way to safely introduce new types is using the `typedef` construct. This allows introducing a type as a non-empty subset of an existing type. More precisely, the new type is specified to be isomorphic to this non-empty subset. For instance:

```
typedef mytype = "{x::nat. x < 10}"
```

This definition requires that the set is non-empty: $\exists x. x \in \{x::\text{nat}. x < 10\}$, which is easy to prove in this case:

```
by (rule_tac x = 1 in exI, simp)
```

where `rule_tac` is a tactic that applies an introduction rule, and `exI` corresponds to the introduction of the existential quantification.

Similarly, the `datatype` command allows the definition of inductive datatypes. It introduces a datatype using a list of *constructors*. For instance, a logical compiler is invoked for the following introduction of the type `option`:

```
datatype  $\alpha$  option = None | Some  $\alpha$ 
```

which generates the underlying type definition and derives distinctness rules and induction principles. Besides the *constructors* `None` and `Some`, the following match-operator and his rules are also generated:

```
case x of None  $\Rightarrow$  ... | Some a  $\Rightarrow$  ...
```

Extensible records. Isabelle/HOL’s support for *extensible records* is of particular importance for our work. Record types are denoted, for example, by:

```
record T = a::T1
          b::T2
```

which implicitly introduces the record constructor $(\mathbf{a}:=\mathbf{e}_1, \mathbf{b}:=\mathbf{e}_2)$ and the update of record \mathbf{r} in field \mathbf{a} , written as $\mathbf{r}(\mathbf{a}:= \mathbf{x})$. Extensible records are represented internally by cartesian products with an implicit free component δ , i.e. in this case by a triple of the type $T_1 \times T_2 \times \delta$. The third component can be referenced by a *special selector* `more` available on extensible records. Thus, the record \mathbf{T} can be extended later on using the syntax:

```
record ET = T + c::T3
```

The key point is that theorems can be established, once and for all, on \mathbf{T} types, even if future parts of the record are not yet known, and reused in the later definition and proofs over \mathbf{ET} -values. Using this feature, we can model the effect of defining the alphabet of UTP processes incrementally while maintaining the full expressivity of HOL wrt. the types of T_1 , T_2 and T_3 .

2.3 Circus and its UTP Foundation

Circus is a formal specification language [18] which integrates the notions of states and complex data types (in a Z-like style) and communicating parallel processes inspired from CSP. From Z, the language inherits the notion of a schema used to model sets of (ground) states as well as syntactic machinery to describe pre-states and post-states; from CSP, the language inherits the concept of *communication events* and typed communication channels, the concepts of deterministic and non-deterministic choice (reflected by the process combinators $P \square P'$ and $P \sqcap P'$), the concept of concealment (hiding) $P \setminus A$ of events in A occurring in the evolution of process P . Due to the presence of state variables, the *Circus* synchronous communication operator syntax is slightly different from CSP: $P \llbracket n \mid c \mid n' \rrbracket P'$ means that P and P' communicate via the channels mentioned in c ; moreover, P may modify the variables mentioned in n only, and P' in n' only, n and n' are disjoint name sets.

Moreover, the language comes with a formal notion of refinement based on a denotational semantics. It follows the failure/divergence semantics [15], (but coined in terms of the UTP [13]) providing a notion of execution trace \mathbf{tr} , refusals \mathbf{ref} , and divergences. It is expressed in terms of the UTP [11] which makes it amenable to other refinement-notions in UTP. Figure 1 presents a simple *Circus* specification, `FIG`, the fresh identifiers generator.

Predicates and Relations. The UTP is a semantic framework based on an alphabetized relational calculus. An *alphabetized predicate* is a pair (*alphabet*, *predicate*) where the free variables appearing in the predicate are all in the alphabet, e.g. $(\{x, y\}, x > y)$. As such, it is very similar to the concept of a *schema* in Z. In the base theory Isabelle/UTP of this work, we represent alphabetized predicates by sets of (extensible) records, e.g. $\{\mathbf{A}. \mathbf{x} \mathbf{A} > \mathbf{y} \mathbf{A}\}$.

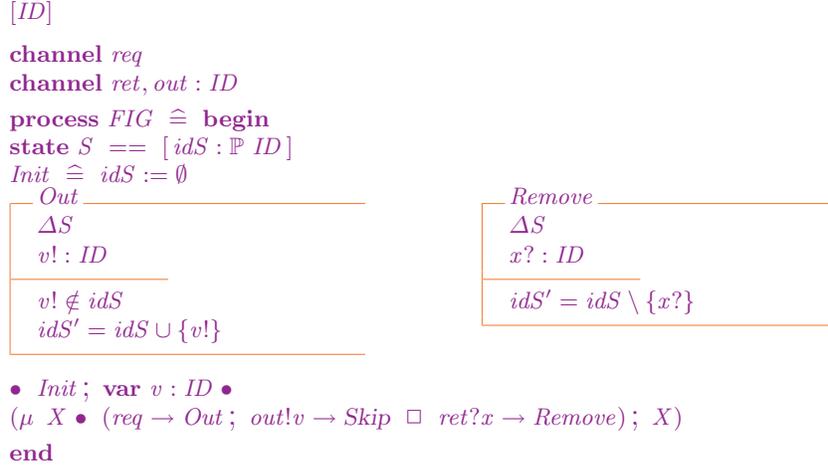


Fig. 1. The Fresh Identifiers Generator in (Textbook) *Circus*

An *alphabetized relation* is an alphabetized predicate where the alphabet is composed of input (undecorated) and output (dashed) variables. In this case the predicate describes a relation between input and output variables, for example $(\{x, x', y, y'\}, x' = x + y)$ which is a notation for: $\{(A, A') . x A' = x A + y A\}$, which is a set of pairs, thus a relation.

Standard predicate calculus operators are used to combine alphabetized predicates. The definition of these operators is very similar to the standard one, with some additional constraints on the alphabets.

Designs and processes. In UTP, in order to explicitly record the termination of a program, a subset of alphabetized relations is introduced. These relations are called *designs* and their alphabet should contain the special boolean observational variable **ok**. It is used to record the start and termination of a program. A UTP design is defined as follows in Isabelle:

$$(P \vdash Q) \equiv \lambda (A, A') . (\text{ok } A \wedge P(A, A')) \longrightarrow (\text{ok } A' \wedge Q(A, A'))$$

Following the way of UTP to describe reactive processes, more observational variables are needed to record the interaction with the environment. Three observational variables are defined for this subset of relations: **wait**, **tr** and **ref**. The boolean variable **wait** records if the process is waiting for an interaction or has terminated. **tr** records the list (trace) of interactions the process has performed so far. The variable **ref** contains the set of interactions (events) the process may refuse to perform. These observational variables defines the basic alphabet of all reactive processes called “**alpha_rp**”.

Some healthiness conditions are defined over **wait**, **tr** and **ref** to ensure that a reactive process satisfies some properties [6] (see Table 2 in [9]).

A CSP process is a UTP reactive process that satisfies two additional healthiness conditions (all well-formedness conditions can be found in [9]). A process that satisfies these conditions is said to be CSP healthy.

3 Isabelle/*Circus*

```

Process    ::= circusprocess Tpar* name = PParagraph* where Action
PParagraph ::= AlphabetP | StateP | ChannelP | NamesetP | ChansetP | SchemaP
           | ActionP
AlphabetP  ::= alphabet [ vardecl+ ]
vardecl    ::= name :: type
StateP     ::= state [ vardecl+ ]
ChannelP   ::= channel [ chandekl+ ]
chandekl   ::= name | name type
NamesetP   ::= nameset name = [ name+ ]
ChansetP   ::= chanset name = [ name+ ]
SchemaP    ::= schema name = SchemaExpression
ActionP    ::= action name = Action
Action     ::= Skip | Stop | Action ; Action | Action □ Action | Action ⊞ Action
           | Action \ chansetN | var := expr | guard & Action | comm → Action
           | Schema name | ActionName | μ var • Action | var var • Action
           | Action [[ namesetN | chansetN | namesetN ]] Action

```

Fig. 2. Isabelle/*Circus* syntax

The Isabelle/*Circus* environment allows a syntax of processes which is close to the textbook presentations of *Circus* (see Fig. 2). Similar to other specification constructs in Isabelle/HOL, this syntax is “parsed away”, i.e. compiled into an internal representation of the denotational semantics of *Circus*, which is a formalization in form of a shallow embedding of the (essentially untyped) paper-and-pencil definitions by Oliveira et al. [13], based on UTP. *Circus* actions are defined as CSP healthy reactive processes.

In the UTP representation of reactive processes we have given in a previous paper [8], the process type is generic. It contains two type parameters that represent the channel type and the alphabet of the process. These parameters are very general, and they are instantiated for each specific process. This could be problematic when representing the *Circus* semantics, since some definitions rely directly on variables and channels (e.g assignment and communication). In this section we present our solution to deal with this kind of problems, and our representation of the *Circus* actions and processes.

We now describe the foundation as well as the semantic definition of some process operators of *Circus*. A distinguishing feature of *Circus* processes are explicit state variables which do not exist in other process algebras like, e.g., CSP. These can be:

- *global* state variables, i.e. they are declared via alphabetized predicates in the **state** section, or Z-like Δ operations on global states that generate alphabetized relations, or
- *local* state variables, i.e. they are result of the variable declaration statement **var var • Action**. The scope of local variables is restricted to **Action**.

On both kind of state variables, logical constraints may be expressed.

3.1 Alphabets and Variables

In order to define the set of variables of a specification, the *Circus* semantics considers the alphabet of its components, be it on the level of alphabetized predicates, alphabetized relations or actions. We recall that these items are represented by sets of records or sets of pairs of records. The *alphabet of a process* is defined by extending the basic reactive process alphabet (cf. Section 2.3) by its variable names and types. For the example *FIG*, where the global state variable *idS* is defined, this is reflected in Isabelle/Circus by the extension of the process alphabet by this variable, i.e. by the extension of the Isabelle/HOL record:

```
record  $\alpha$  alpha =  $\alpha$  alpha_rp + idS :: ID set
```

This introduces the record type `alpha` that contains the observational variables of a reactive process, plus the variable `idS`. Note that our *Circus* semantic representation allows “built-in” bindings of alphabets in a typed way. Moreover, there is no restriction on the associated HOL type. However, the inconvenience of this representation is that variables cannot be introduced “on the fly”; they must be known statically i.e. at type inference time. Another consequence is that a “syntactic” operation such as variable renaming has to be expressed as a “semantic” operation that maps one record type into another.

Updating and accessing global variables. Since the alphabets are represented by HOL records, i.e. a kind binding “*name* \mapsto *value*”, we need a certain infrastructure to access data in them and to update them. The Isabelle representation as records gives us already two functions (for each record) “select” and “update”. The “select” function returns the value of a given variable name, and the “update” functions updates the value of this variable. Since we may have different HOL types for different variables, a unique definition for select and update cannot be provided. There is an instance of these functions for each variable in the record. The name of the variable is used to distinguish the different instances: for the select function the name is used directly and for the update function the name is used as a prefix e.g. for a variable named “x” the names of the *select* and *update* functions are respectively `x` of type α and `x_update`.

Since a variable is characterized essentially by these functions, we define a general type (synonym) called `var` which represents a variable as a pair of its select and update function (in the underlying state σ).

```
types ( $\beta$ ,  $\sigma$ ) var = " $(\sigma \Rightarrow \beta) * ((\beta \Rightarrow \beta) \Rightarrow \sigma \Rightarrow \sigma)$ "
```

For a given alphabet (record) of type σ , $(\beta, \text{the type } \sigma)\text{var}$ represents the type of the variables whose value type is β . One can then extract the select and update functions from a given variable with the following functions:

```
definition select :: " $(\beta, \sigma)$  var  $\Rightarrow \sigma \Rightarrow \beta$ "
  where select f  $\equiv$  (fst f)
```

```
definition update :: " $(\beta, \sigma)$  var  $\Rightarrow \beta \Rightarrow \sigma \Rightarrow \sigma$ "
  where update f v  $\equiv$  (snd f) ( $\lambda$  _ . v)
```

Finally, we introduce a function called `VAR` to implement a syntactic translation of a variable name to an entity of type `var`.

```

syntax "_VAR" :: "id  $\Rightarrow$  ( $\beta$ ,  $\sigma$ ) var" ("VAR _")
translations VAR x => (x, _update_ name x)

```

Note that in this syntactic translation rule, `_update_ name x` stands for the concatenation of the string `_update_` with the content of the variable `x`; the resulting `_update_x` in this example is mapped to the field-update function of the extensible record `x_update` by a default mechanism. On this basis, the assignment notation can be written as usual:

```

syntax
  "_assign" :: "id  $\Rightarrow$  ( $\sigma \Rightarrow \beta$ )  $\Rightarrow$  ( $\alpha$ ,  $\sigma$ ) action" ("_ ':='_ ")
translations
  "x ':='_ E"   => "CONST ASSIGN (VAR x) E"

```

and mapped to the *semantics* of the program variable `(x,x_update)` together with the universal `ASSIGN` operator defined later on, in Section 3.3.

Updating and accessing local variables. In *Circus*, local program variables can be introduced on the fly, and their scopes are explicitly defined, as can be seen in the *FIG* example. In textbook *Circus*, nested scopes are handled by variable renaming which is not possible in our representation due to the implicit representation of variable names. We represent local program variables by global variables, using the `var` type defined above, where selection and update involve an explicit stack discipline. Each variable is mapped to a list of values, and not to one value only (as for state variables). Entering the scope of a variable is just adding a new value as the head of the corresponding values list. Leaving a variable scope is just removing the head of the values list. The select and update functions correspond to selecting and updating the head of the list. This ensures dynamic scoping, as it is stated by the *Circus* semantics.

Note that this encoding scheme requires to make local variables lexically distinct from global variables; local variable instances are just distinguished from the global ones by the stack discipline.

3.2 Synchronization infrastructure: Name sets and channels.

Name sets. An important notion, used in the definition of parallel *Circus* actions, is name sets as seen in Section 2.3. A name set is a set of variable names, which is a subset of the alphabet. This notion cannot be directly expressed in our representation since variable names are not explicitly represented. Thus its definition relies on the characterization of the variables in our representation. As for variables, name sets are defined by their functional characterization. They are used in the definition of the binding merge function *MSt* below:

$$\forall v \bullet (v \in ns1 \Rightarrow v' = (1.v)) \wedge (v \in ns2 \Rightarrow v' = (2.v)) \wedge (v \notin ns1 \cup ns2 \Rightarrow v' = v).$$

The disjoint name sets `ns1` and `ns2` are used to determine which variable values (extracted from local bindings of the parallel components) are used to

update the global binding of the process. A name set can be functionally defined as a binding update function, that copies values from a local binding to the global one. For example, a name set NS that only contains the variable x can be defined as follows in Isabelle/Circus:

```
definition NS lb gb ≡ x_update (x lb) gb
```

where `lb` and `gb` stands for local and global bindings, `x` and `x_update` are the select and update functions of variable `x`. Then the merge function can be defined by composing the application of the name sets to the global binding.

Channels. Reactive processes interact with the environment via synchronizations and communications. A synchronization is an interaction via a channel without any exchange of data. A communication is a synchronization with data exchange. In order to reason about communications in the same way, a datatype `channels` is defined using the channels names as constructors. For instance, in:

```
datatype channels = chan1 | chan2 nat | chan3 bool
```

we declare three channels: `chan1` that synchronizes without data, `chan2` that communicates natural values and `chan3` that exchanges boolean values.

This definition makes it possible to reason globally about communications since they have the same type. However, the channels may not have the same type: in the example above, the types of `chan1`, `chan2` and `chan3` are respectively `channels`, `nat ⇒ channels` and `bool ⇒ channels`. In the definition of some *Circus* operators, we need to compare two channels, and one can't compare for example `chan1` with `chan2` since they don't have the same type. A solution would be to compare `chan1` with `(chan2 v)`. The types are equivalent in this case, but the problem remains because comparing `(chan2 0)` to `(chan2 1)` will state inequality just because the communicated values are not equal. We could define an inductive function over the datatype `channels` to compare channels, but this is only possible when all the channels are known *a priori*.

Thus, we add some constraint to the generic channels type: we require the `channels` type to implement a function `chan_eq` that tests the equality of two channels. Fortunately, Isabelle/HOL provides a construct for this kind of restriction: the type classes (sorts) mentioned in Section 2.1. We define a type class (interface) `chan_eq` that contains a signature of the `chan_eq` function.

```
class chan_eq =
  fixes chan_eq :: "α ⇒ α ⇒ bool"
begin end
```

Concrete channels type must implement the interface (class) “`chan_eq`” that can be easily defined for this concrete type. Moreover, one can use this class to add some definition that depends on the channel equivalence function. For example, a trace equivalence function can be defined as follows:

```
fun tr_eq where
  tr_eq [] [] = True | tr_eq xs [] = False | tr_eq [] ys = False
| tr_eq (x#xs) (y#ys) = if chan_eq x y then tr_eq xs ys else False
```

It is applicable to traces of elements whose type belongs to the sort `chan_eq`.

3.3 Actions and Processes

The *Circus* actions type is defined as the set of all the CSP healthy reactive processes. The type $(\alpha, \sigma)\text{relation_rp}$ is the reactive process type where α is of `channels` type and σ is a record extensions of `action_rp`, i.e. the global state variables. On this basis, we can encode the concept of a process for a family of possible state instances. We introduce below the vital type `action`:

```
typedef(Action)
  ( $\alpha::\text{chan\_eq}, \sigma$ ) action = { $p::(\alpha, \sigma)\text{relation\_rp}$ . is\_CSP\_process p}
proof - {...}
qed
```

As mentioned before, a type-definition introduces a new type by stating a set. In our case it is the set of reactive processes that satisfy the healthiness-conditions for CSP-processes, isomorphic to the new type.

Technically, this construct introduces two constants definitions `Abs_Action` and `Rep_Action` respectively of type $(\alpha, \sigma)\text{relation_rp} \Rightarrow (\alpha, \sigma)\text{action}$ and $(\alpha, \sigma)\text{action} \Rightarrow (\alpha, \sigma)\text{relation_rp}$ as well as the usual two axioms expressing the bijection `Abs_Action(Rep_Action(X))=X` and `is_CSP_process p \implies Rep_Action(Abs_Action(p))=p` where `is_CSP_process` captures the healthiness conditions.

Every *Circus* action is an abstraction of an alphabetized predicate. In [9], we introduce the definitions of all the actions and operators using their denotational semantics. The environment contains, for each action, the proof that this predicate is CSP healthy. In this section, we present some of the important definitions, namely: basic actions, assignments, communications, hiding, and recursion.

Basic actions. `Stop` is defined as a reactive design, with a precondition `true` and a postcondition stating that the system deadlocks and the traces are not evolving.

definition

```
Stop  $\equiv$  Abs_Action (R (true  $\vdash$   $\lambda(A, A')$ . tr A' = tr A  $\wedge$  wait A'))
```

`Skip` is defined as a reactive design, with a precondition `true` and a postcondition stating that the system terminates and all the state variables are not changed. We represent this fact by stating that the `more` field (seen in Section 2.2) is not changed, since this field is mapped to all the state variables. Note that using the `more`-field is a tribute to our encoding of alphabets by extensible records and stands for all future extensions of the alphabet (e.g. state variables).

```
definition Skip  $\equiv$  Abs_Action (R (true  $\vdash$   $\lambda(A, A')$ . tr A' = tr A
 $\wedge$   $\neg$  wait A'  $\wedge$  more A = more A'))
```

The universal assignment action. In Section 3.1, we described how global and local variables are represented by access- and updates functions introduced by fields in extensible records. In these terms, the "lifting" to the assignment action in *Circus* processes is straightforward:

definition

ASSIGN::" $(\beta, \sigma) \text{ var} \Rightarrow (\sigma \Rightarrow \beta) \Rightarrow (\alpha::\text{ev_eq}, \sigma) \text{ action}$ "

where

ASSIGN $x \ e \equiv \text{Abs_Action } (R (\text{true} \vdash Y))$

where

$Y = \lambda (A, A'). \text{tr } A' = \text{tr } A \wedge \neg \text{wait } A' \wedge$
 $\text{more } A' = (\text{assign } x \ (e \ (\text{more } A))) \ (\text{more } A)$

where `assign` is the projection into the update operation of a semantic variable described in section 3.1.

Communications. The definition of prefixed actions is based on the definition of a special relation `do_I`. In the *Circus* denotational semantics [13], various forms of prefixing were defined. In our theory, we define one general form, and the other forms are defined as special cases.

definition `do_I c x P` $\equiv X \triangleleft \text{wait } o \ \text{fst} \triangleright Y$

where

$X = (\lambda (A, A'). \text{tr } A = \text{tr } A' \wedge ((c \text{ ' } P) \cap \text{ref } A') = \{\})$

and

$Y = (\lambda (A, A'). \text{hd } ((\text{tr } A') - (\text{tr } A)) \in (c \text{ ' } P) \wedge$
 $(c \ (\text{select } x \ (\text{more } A))) = (\text{last } (\text{tr } A')))$

where `c` is a channel constructor, `x` is a variable (of `var` type) and `P` is a predicate. The `do_I` relation gives the semantics of an interaction: if the system is ready to interact, the trace is unchanged and the waiting channel is not refused. After performing the interaction, the new event in the trace corresponds to this interaction.

The semantics of the whole action is given by the following definition:

definition `Prefix c x P S` $\equiv \text{Abs_Action}(R (\text{true} \vdash Y)) ; S$

where

$Y = \text{do_I } c \ x \ P \wedge (\lambda (A, A'). \text{more } A' = \text{more } A)$

where `c` is a channel constructor, `x` is a variable (of type `var`), `P` is a predicate and `S` is an action. This definition states that the prefixed action semantics is given by the interaction semantics (`do_I`) sequentially composed with the semantics of the continuation (action `S`).

Different types of communication are considered:

- Inputs: the communication is done over a variable.
- Constrained Inputs: the input variable value is constrained with a predicate.
- Outputs: the communications exchanges only one value.
- Synchronizations: only the channel name is considered (no data).

The semantics of these different forms of communications is based on the general definition above.

`definition read c x P ≡ Prefix c x true P`

`definition write1 c a P ≡ Prefix c (λs. a s, (λ x. λy. y)) true P`

`definition write0 c P ≡ Prefix (λ_.c) (λ_._, (λ x. λy. y)) true P`

where `read`, `write1` and `write0` respectively correspond to inputs, outputs and synchronization. Constrained inputs correspond to the general definition.

We configure the Isabelle syntax-engine such that it parses the usual communication primitives and gives the corresponding semantics:

`translations`

```

c ? p → P      == CONST read c (VAR p) P
c ? p : b → P  == CONST Prefix c (VAR p) b P
c ! p → P      == CONST write1 c p P
a → P          == CONST write0 (TYPE(_)) a P

```

Hiding. The hiding operator is interesting because it depends on a channel set. This operator $P \setminus cs$ is used to encapsulate the events that are in the channel set cs . These events become no longer visible from the environment. The semantics of the hiding operator is given by the following reactive process:

`definition`

`Hide :: "[(α , σ) action , α set] \Rightarrow (α , σ) action" (infixl "\)`

where

$$P \setminus cs \equiv \text{Abs_Action}(R(\lambda (A, A'). \exists s. (\text{Rep_Action } P)(A, A'(|\text{tr} := s, \text{ref} := (\text{ref } A') \cup cs|)) \wedge (\text{tr } A' - \text{tr } A) = (\text{tr_filter } (s - \text{tr } A) cs))); \text{Skip}$$

The definition uses a filtering function `tr_filter` that removes from a trace the events whose channels belong to a given set. The definition of this function is based on the function `chan_eq` we defined in the class `chan_eq`. This explains the presence of the constraint on the type of the action channels in the hiding definition, and in the definition of the filtering function below:

`fun tr_filter :: "a::chan_eq list \Rightarrow a set \Rightarrow a list" where`

```

tr_filter [] cs = []
| tr_filter (x#xs) cs = (if (¬ chan-in_set x cs)
                        then (x#(tr_filter xs cs))
                        else (tr_filter xs cs))

```

where the `chan-in_set` function checks if a given channel belongs to a channel set using `chan_eq` as equality function.

Recursion. To represent the recursion operator “ μ ” over actions, we use the universal least fix-point operator “`lfp`” defined in the HOL library for lattices and we follow again [13]. The use of least fix-points in [13] is the most substantial

deviation from the standard CSP denotational semantics, which requires Scott-domains and complete partial orderings. The operator *lfp* is inherited from the “*Complete Lattice class*” under some conditions, and all theorems defined over this operator can be reused. In order to reuse this operator, we have to show that the least-fixpoint over functionals that enrich pairs of failure - and divergence trace sets monotonely, produces an **action** that satisfies the CSP healthiness conditions. This consistency proof for the recursion operator is the largest contained in the Isabelle/*Circus* library.

Therefore, we must prove that the *Circus* actions type defines a complete lattice. This leads to prove that the actions type belongs to the HOL “*Complete Lattice class*”. Since type classes in HOL are hierarchic, the proof is in three steps: first, a proof that the *Circus* actions type forms a lattice by instantiating the HOL “*Lattice class*”; second, a proof that actions type instantiates a subclass of lattices called “*Bounded Lattice class*”; third, proof of the instantiation from the “*Complete Lattice class*”. More on these proofs can be found in [9].

Circus Processes. A *Circus* process is defined in our environment as a local theory by introducing qualified names for all its components. This is very similar to the notion of *namespaces* popular in programming languages. Defining a *Circus* process locally makes it possible to encapsulate definitions of alphabet, channels, schema expressions and actions in the same namespace. It is important for the foundation of Isabelle/*Circus* to avoid the ambiguity between local process entities definitions (e.g. `FIG.Out` and `DFIG.Out` in the example of Section 4).

4 Using Isabelle/*Circus*

We describe the front-end interface of Isabelle/*Circus*. In order to support a maximum of common *Circus* syntactic look-and-feel, we have programmed at the SML level of Isabelle a compiler that parses and (partially) pretty prints *Circus* process given in the syntax presented in Figure 2.

4.1 Writing specifications

A specification is a sequence of paragraphs. Each paragraph may be a declaration of alphabet, state, channels, name sets, channel sets, schema expressions or actions. The main action is introduced by the keyword `where`. Below, we illustrate how to use the environment to write a *Circus* specification using the `FIG` process example presented in Figure 1.

```
circusprocess FIG =
  alphabet = [v::nat, x::nat]
  state = [idS::nat set]
  channel = [req, ret nat, out nat]
  schema Init = idS := {}
  schema Out =  $\exists a. v' = a \wedge v' \notin idS \wedge idS' = idS \cup \{v'\}$ 
```

```

schema Remove = x  $\notin$  idS  $\wedge$  idS' = idS - {x}
where var v  $\cdot$  Schema Init; ( $\mu$  X  $\cdot$ (req  $\rightarrow$ Schema Out; out!v  $\rightarrow$ Skip)
     $\square$  (ret?x  $\rightarrow$ Schema Remove); X)

```

Each line of the specification is translated into the corresponding semantic operator given in Section 3.3. We describe below the result of executing each command of FIG:

- the compiler introduces a scope of local components whose names are qualified by the process name (FIG in the example).
- **alphabet** generates a list of record fields to represent the binding. These fields map names to value lists.
- **state** generates a list of record fields that corresponds to the state variables. The names are mapped to single values. This command, together with **alphabet** command, generates a record that represents all the variables (for the FIG example the command generates the record FIG_alphabet, that contains the fields v and x of type nat list and the field idS of type nat set).
- **channel** introduces a datatype of typed communication channels (for the FIG example the command generates the datatype FIG_channels that contains the constructors req without communicated value and ret and out that communicate natural values).
- **schema** allows the definition of schema expressions represented as an alphabetized relation over the process variables (in the example the schema expressions FIG.Init, FIG.Out and FIG.Remove are generated).
- **action** introduces definitions for *Circus* actions in the process. These definitions are based on the denotational semantics of *Circus* actions. The type parameters of the action type are instantiated with the locally defined channels and alphabet types.
- **where** introduces the main action as in **action** command (in the example the main action is FIG.FIG of type (FIG_channels, FIG_alphabet)action).

4.2 Relational and Functional Refinement in Circus

The main goal of Isabelle/*Circus* is to provide a proof environment for *Circus* processes. The “shallow-embedding” of *Circus* and UTP in Isabelle/HOL offers the possibility to reuse proof procedures, infrastructure and theorem libraries already existing in Isabelle/HOL. Moreover, once a process specification is encoded and parsed in Isabelle/*Circus*, proofs of, e. g., refinement properties can be developed using the ISAR language for structured proofs.

To show in more details how to use Isabelle/*Circus*, we provide a small example of action refinement proof. The refinement relation is defined as the universal reverse implication in the UTP. In *Circus*, it is defined as follows:

definition A1 \sqsubseteq_c A2 \equiv (Rep_Action A1) \sqsubseteq_{utp} (Rep_Action A2)

where A1 and A2 are *Circus* actions, \sqsubseteq_c and \sqsubseteq_{utp} stands respectively for refinement relation on *Circus* actions and on UTP predicate.

This definition assumes that the actions A1 and A2 share the same alphabet (binding) and the same channels. In general, refinement involves an important data evolution and growth. The data refinement is defined in [16, 5] by backwards and forwards simulations. In this paper, we restrict ourselves to a special case, the so-called *functional* backwards simulation. This refers to the fact that the abstraction relation R that relates concrete and abstract actions is just a function:

definition Simulation (" \preceq ") where

$$A1 \preceq A2 = \forall a \ b. (\text{Rep_Action } A2)(a, b) \longrightarrow (\text{Rep_Action } A1)(R \ a, R \ b)$$

where A1 and A2 are *Circus* actions and R is a function mapping the corresponding A1 alphabet to the A2 alphabet.

4.3 Refinement Proofs

We can use the definition of simulation to transform the proof of refinement to a simple proof of implication by unfolding the operators in terms of their underlying relational semantics. The problem with this approach is that the size of proofs will grow exponentially with the size of the processes. To avoid this problem, some general refinement laws were defined in [5] to deal with the refinement of *Circus* actions at operators level and not at UTP level. We introduced and proved a subset of these laws in our environment (see Table 1).

$\frac{P \preceq_S Q \quad P' \preceq_S Q'}{P; P' \preceq_S Q; Q'} \text{SeqI}$	$\frac{P \preceq_S Q \quad g_1 \simeq_S g_2}{g_1 \& P \preceq_S g_2 \& Q} \text{GrdI}$
$\frac{P \preceq_S Q \quad x \sim_S y}{\text{var } x \bullet P \preceq_S \text{var } y \bullet Q} \text{VarI}$	$\frac{P \preceq_S Q \quad x \sim_S y}{c?x \rightarrow P \preceq_S c?y \rightarrow Q} \text{InpI}$
$\frac{P \preceq_S Q \quad P' \preceq_S Q'}{P \sqcap P' \preceq_S Q \sqcap Q'} \text{NdetI}$	$\frac{P \preceq_S Q \quad x \sim_S y}{c!x \rightarrow P \preceq_S c!y \rightarrow Q} \text{OutI}$
$\frac{[X \preceq_S Y] \quad \dots \quad P X \preceq_S Q Y \quad \text{mono } P \quad \text{mono } Q}{\mu X \bullet P X \preceq_S \mu Y \bullet Q Y} \text{MuI}$	$\frac{P \preceq_S Q \quad P' \preceq_S Q'}{P \sqcap P' \preceq_S Q \sqcap Q'} \text{DetI}$
$\frac{[Pre \ sc_1(S \ A)] \quad [Pre \ sc_1(S \ A) \quad sc_2(A, A')] \quad \dots \quad Pre \ sc_2 \ A \quad sc_1(S \ A, S \ A')}{\text{schema } sc_1 \preceq_S \text{schema } sc_2} \text{SchI}$	$\frac{P \preceq_S Q}{a \rightarrow P \preceq_S a \rightarrow Q} \text{SyncI}$
$\frac{P \preceq_S Q \quad P' \preceq_S Q' \quad ns_1 \sim_S ns'_1 \quad ns_2 \sim_S ns'_2}{P \llbracket ns_1 \mid cs \mid ns_2 \rrbracket P' \preceq_S Q \llbracket ns'_1 \mid cs \mid ns'_2 \rrbracket Q'} \text{ParI}$	$\frac{}{Skip \preceq_S Skip} \text{SkipI}$

Table 1. Proved refinement laws

In Table 1, the relations “ $x \sim_S y$ ” and “ $g_1 \simeq_S g_2$ ” record the fact that the variable x (repectively the guard g_1) is refined by the variable y (repectively by the guard g_2) w.r.t the simulation function S .

These laws can be used in complex refinement proofs to simplify them at the *Circus* level. More rules can be defined and proved to deal with more complicated statements like combination of operators for example. Using these laws, and exploiting the advantages of a shallow embedding, the automated proof of refinement becomes surprisingly simple.

Coming back to our example, let us consider the DFIG specification below, where the management of the identifiers via the set `idS` is refined into a set of removed identifiers `retidS` and a number `max`, which is the rank of the last issued identifier.

```
circusprocess DFIG =
  alphabet = [w::nat, y::nat]
  state = [retidS::nat set, max::nat]
  schema Init = retidS' = {} ^ max' = 0
  schema Out = w' = max ^ max' = max+1 ^ retidS' = retidS - {max}
  schema Remove = y < max ^ y ∉ retidS ^ retidS' = retidS ∪ {y}
                  ^ max' = max
  where var w · Schema Init; (μ X · (req → Schema Out; out!w → Skip)
                              □ (ret?y → Schema Remove); X)
```

We provide the proof of refinement of FIG by DFIG just instantiating the simulation function R by the following abstraction function, that maps the underlying concrete states to abstract states:

```
definition Sim A = FIG_alphabet.make (w A) (y A)
  ({a. a < (max A) ^ a ∉ (retidS A)})
```

where A is the alphabet of DFIG, and `FIG_alphabet.make` yields an alphabet of type `FIG_Alphabet` initializing the values of `v`, `x` and `idS` by their corresponding values from `DFIG_alphabet`: `w`, `y` and `{a. a < max ^ a ∉ retidS}`.

To prove that DFIG is a refinement of FIG one must prove that the main action `DFIG.DFIG` refines the main action `FIG.FIG`. The definition is then simplified, and the refinement laws are applied to simplify the proof goal. Thus, the full proof consists of a few lines in ISAR:

```
theorem "FIG.FIG ≤Sim DFIG.DFIG"
  apply (auto simp: DFIG.DFIG_def FIG.FIG_def mono_Seq
            intro!: VarI SeqI MuI DetI SyncI InpI OutI SkipI)
  apply (simp_all add: SimRemove SimOut SimInit Sim_def)
done
```

First, the definitions of `FIG.FIG` and `DFIG.DFIG` are simplified and the defined refinement laws are used by the `auto` tactic as introduction rules. The second step replaces the definition of the simulation function and uses some proved lemmas to finish the proof. The three lemmas used in this proof: `SimInit`, `SimOut` and `SimRemove` give proofs of simulation for the schema `Init`, `Out` and `Remove`.

5 Conclusions

We have shown for the language *Circus*, which combines data-oriented modeling in the style of Z and behavioral modeling in the style of CSP, a semantics in form of a shallow embedding in Isabelle/HOL. In particular, by representing the somewhat non-standard concept of the *alphabet* in UTP in form of extensible records in HOL, we achieved a fairly compact, typed presentation of the language. In contrast to previous work based on some deep embedding [19], this shallow embedding allows arbitrary (higher-order) HOL-types for channels, events, and state-variables, such as, e.g., sets of relations etc. Besides, systematic renaming of local variables is avoided by compiling them essentially to global variables using a stack of variable instances. The necessary proofs for showing that the definitions are consistent — i.e. satisfy altogether `is_CSP_healthy` — have been done, together with a number of algebraic simplification laws on *Circus* processes.

Since the encoding effort can be hidden behind the scene by flexible extension mechanisms of the Isabelle, it is possible to have a compact notation for both specifications and proofs. Moreover, existing standard tactics of Isabelle such as `auto`, `simp` and `metis` can be reused since our *Circus* semantics is representationally close to HOL. Thus, we provide an environment that can cope with combined refinements concerning data and behavior. Finally, we demonstrate its power — w.r.t. both expressivity and proof automation — with a small, but prototypic example of a process-refinement.

In the future, we intend to use Isabelle/*Circus* for the generation of test-cases, on the basis of [4], using the HOL-TestGen-environment [2].

6 Acknowledgement

We warmly thank Markarius Wenzel for his valuable help with the Isabelle framework. Furthermore, we are greatly indebted to Ana Cavalcanti for her comments on the semantic foundation of this work.

References

1. P. B. Andrews. *Introduction to Mathematical Logic and Type Theory: To Truth through Proof*. Kluwer Academic, 2nd edition, 2002. now published by Springer.
2. A. D. Brucker and B. Wolff. On theorem prover-based testing. *Formal Aspects of Computing*, 2012. To appear.
3. M. Butler. CSP2B: A practical approach to combining CSP and B. *Formal Aspects of Computing*, 12:182–196, 2000.
4. A. Cavalcanti and M.-C. Gaudel. Testing for refinement in Circus. *Acta Informatica*, 48(2):97–147, 2011.
5. A. L. C. Cavalcanti, A. C. A. Sampaio, and J. C. P. Woodcock. A Refinement Strategy for *Circus*. *Formal Aspects of Computing*, 15(2 - 3):146 — 181, 2003.

6. A. L. C. Cavalcanti and J. C. P. Woodcock. A Tutorial Introduction to CSP in Unifying Theories of Programming. In *Refinement Techniques in Software Engineering*, volume 3167 of *LNCS*, pages 220 – 268. Springer-Verlag, 2006.
7. A. Church. A formulation of the simple theory of types. *Journal of Symbolic Logic*, 5(2):56–68, June 1940.
8. A. Feliachi, M.-C. Gaudel, and B. Wolff. Unifying theories in Isabelle/HOL. In *UTP 2010*, volume 6445 of *LNCS*, pages 188–206. Springer Verlag, 2010.
9. A. Feliachi, M.-C. Gaudel, and B. Wolff. Isabelle/Circus : a process specification and verification environment. Technical Report 1547, LRI, <http://www.lri.fr/srubrique.php?news=33>, Univ. Paris-Sud XI, Nov. 2011.
10. C. Fischer. How to combine Z with process algebra. In *11th Int. Conf. of Z Users on The Z Formal Specification Notation*, pages 5–23. Springer-Verlag, 1998.
11. C. A. R. Hoare and He Jifeng. *Unifying Theories of Programming*. Prentice Hall International Series in Computer Science, 1998.
12. T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL—A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer-Verlag, 2002.
13. M. Oliveira, A.L.C. Cavalcanti, and J.C.P. Woodcock. A denotational semantics for Circus. *Electron. Notes Theor. Comput. Sci.*, 187:107–123, 2007.
14. M. Roggenbach. CSP-CASL: a new integration of process algebra and algebraic specification. *Theor. Comput. Sci.*, 354:42–71, 2006.
15. A. W. Roscoe, C. A. R. Hoare, and Richard Bird. *The Theory and Practice of Concurrency*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1997.
16. A. C. A. Sampaio, J. C. P. Woodcock, and A. L. C. Cavalcanti. Refinement in Circus. In *FME 2002*, volume 2391 of *LNCS*, pages 451—470. Springer, 2002.
17. K. Taguchi and K. Araki. The state-based CCS semantics for concurrent Z specification. In *ICFEM'97*, pages 283–292. IEEE, 1997.
18. J. C. P. Woodcock and A. L. C. Cavalcanti. The semantics of Circus. In *ZB 2002*, volume 2272 of *LNCS*, pages 184—203. Springer-Verlag, 2002.
19. F. Zeyda and A.L.C. Cavalcanti. Encoding Circus programs in ProofPowerZ. In *UTP 2008*, volume 5713 of *LNCS*. Springer-Verlag, 2009.