

Verification of Control Systems using *Circus*

Ana Cavalcanti
Department of Computer Science
University of York
York - England
Ana.Cavalcanti@cs.york.ac.uk

Phil Clayton
Systems Assurance Group, QinetiQ
Malvern, England
p.clayton@eris.qinetiq.com

Abstract

The design of control systems is usually based on diagrammatic definitions of control laws. The independent use of Z and CSP to verify their implementations has been successful, even for very large applications; high levels of automation have been achieved with tools based on a theorem prover called ProofPower. We have extended this approach to integrate the use of Z and CSP using a notation called Circus; as a result, we can handle a larger set of diagrams and prove more properties of the implementation. In this paper, we show how we can reuse the existing tools and experience to provide automation in the context of the new technique. This gives us confidence in its applicability in industry.

Keywords: Z, CSP, Simulink, refinement.

1. Introduction

Typically, control systems are designed using control law diagrams, which are graphs of blocks. The connections in the graph represent wires that carry signals, and the blocks represent functions on the values of these signals: they determine the values output through outgoing wires in terms of the values input through incoming wires. A diagram can have a continuous or a discrete time model; we work with discrete-time diagrams, in which signals are sampled at fixed intervals, so that input and output occur in cycles.

Numerical modelling and simulation are the main techniques for validation of control laws. More recent work has proposed the use of logic [6, 12]. We use refinement for the verification of implementations, rather than analysis of diagrams. For that, we use a model of discrete control law diagrams based on first order logic (Z [20]) and a process algebra (CSP [11, 17]). Our objective is to extend an approach to verification

of control systems that has been successfully applied in industry. We aim at covering a larger set of diagrams and program properties, but still allow reuse of the expertise on languages and tools already available.

Simulink [1] is practically a standard as a tool to draw and analyse control law diagrams. A translator from discrete-time Simulink diagrams to Z specifications is presented in [3]; it is called ClawZ. Extensive experience with its application for the verification of sequential Ada subprograms (procedures and functions) used in implementations of control systems is already available at QinetiQ [2]. The Z specifications generated by ClawZ are used to formulate refinement conjectures that can be proved using tools based on ProofPower [13]. This is a theorem prover for Z whose powerful tactics allow a high degree of automation.

The Z model of a diagram defined by ClawZ captures the functionality of one cycle; it defines the outputs of the diagram as a function of its inputs. Concurrency is not captured, and, in fact, the computation embedded in the blocks can be carried out concurrently, with order imposed only by the wiring. To verify the architecture of parallel implementations of control law diagrams, QinetiQ employ CSP and model checking.

As an extension of this work, we proposed a *Circus* model for discrete-time Simulink diagrams [8]. *Circus* [19] is a combination of Z, CSP, and the refinement calculus [15]. We defined a translation strategy to convert the output generated by an extended version of ClawZ into a *Circus* specification, using additional graph information about the diagram. In the *Circus* model of a diagram, we capture both its functionality and its concurrent behaviour over any number of cycles.

A calculational strategy to develop concurrent programs based on centralised *Circus* specifications is presented in [9]. In [7], we proposed a different strategy to verify that a *Circus* model of an Ada program refines the *Circus* model of a diagram. In this case, typically,

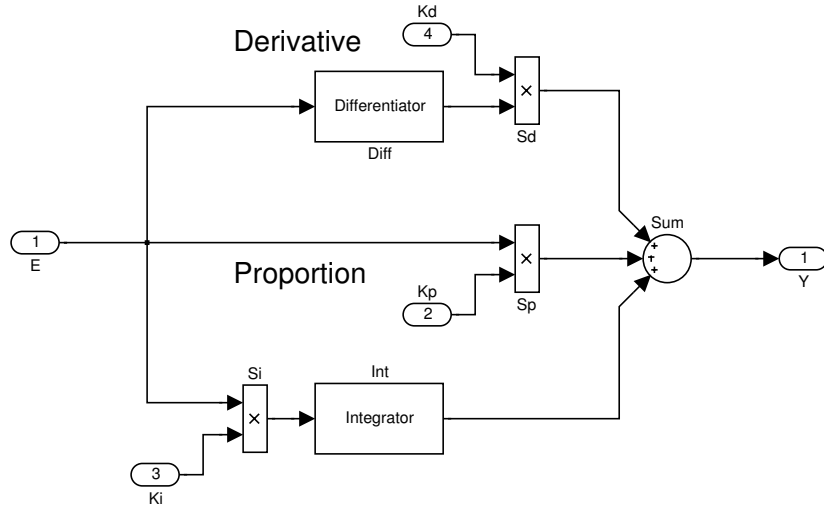


Figure 1. PID (Proportional Integral Derivative) controller

the massive parallelism of the diagram model is reduced to match the architecture of the implementation.

With this technique, independent analyses of the functionality and parallel architecture of an implementation is no longer needed. Yet, in this paper we show that it is possible to reuse the existing technology and expertise on the use of ProofPower and its associated tools in the context of the integrated technique. We provide a strategy based on *Circus* refinement laws to transform the model of the calculations of outputs and of the state updates into appropriate specifications, and show how refinement proof carried out using ProofPower can be used to justify part of the verification entailed by the *Circus* refinement strategy.

In the next section, we give a brief introduction to Simulink diagrams and their *Circus* models. In Section 3, we discuss the features and *Circus* model of typical Ada implementations of control laws. An overview of our refinement strategy is presented in Section 4. In Section 5 we present our approach to reuse of the existing verification technique. Finally, in Section 6, we indicate some future work.

2. Simulink diagrams and *Circus*

Figure 1 gives an example of a simple Simulink diagram for a PID (Proportional Integral Derivative) controller. The rounded numbered boxes are the inputs, namely, E, Kp, Ki, and Kd, and the output, Y. The circle models a sum. The boxes labelled \times are product blocks. The boxes labelled Differentiator and Integrator are subsystems: blocks defined by other diagrams. For the Integrator, for example, the diagram is that in




Figure 2. PID Integrator

Figure 2. There, the block labelled $1/z$ label is a unit delay, which stores the input, and outputs the value stored in the previous cycle: it has a state.

The Z model of a diagram generated by ClawZ uses the notation of ProofPower. It defines blocks as sets of bindings (records), typically specified by schemas. For the block Int, we have the following schema.

```


$$\begin{array}{l} \text{z} \\ \text{pid\_Int} \\ \hline \text{In1?} : \mathbb{U}; \text{Out1!} : \mathbb{U} \\ \text{Sum} : \text{pid\_Int\_Sum}; \\ \text{UnitDelay} : \text{pid\_Int\_UnitDelay}; \\ \hline \text{Sum.In1?} = \text{In1?}; \\ \text{Sum.In2?} = \text{UnitDelay.Out1!}; \\ \text{Out1!} = \text{UnitDelay.In1?} = \text{Sum.Out1!} \end{array}$$


```

The component $In1?$ represents the input of the block, and $Out1!$, its output. The other components represent the blocks in the Integrator diagram: their types are given by the schemas that specify those blocks. The

predicate defines the connection of the inputs and outputs, and identifies the input and the output of the Integrator diagram with inputs and outputs of the blocks. The type \mathbb{U} is a universal type in ProofPower. In the case of a simple block, that does not define a subsystem, the only components are inputs and outputs, and the predicate defines how the output can be calculated.

There is a library of blocks in Simulink; correspondingly, there is a library of schemas in ClawZ to define some Simulink library blocks. For example, *pid_Int_Sum* is defined as the ClawZ library schema *Sum_P2*, which models a Sum block with two inputs. For blocks with state, ClawZ includes in the schemas components *state*, *state'*, and *initialstate*; they record the value of the state at the beginning and at the end of the cycle, and at the beginning of the first cycle.

Circus is a language for refinement; it includes specification and programming constructs, which can be freely mixed. For specification, we can use Z and Morgan's specification statements. To model communication and concurrency, we can use CSP. To write code, we can use the language of guarded commands. A *Circus* program is a sequence of paragraphs, just like in Z, but there are also channel and process paragraphs; there is an example in Figure 3.

The blocks, as well as the diagram itself, are defined as processes in *Circus*. In the case of the diagram, the process is a parallel composition of the processes for the blocks; in the case of a block, the process lifts the output of ClawZ to *Circus*. Even if the block is a subsystem, the parallelism in its diagram is ignored. The *Circus* model captures the parallelism only of the top level diagram. For the block *Int* block, for example, we have the specification in Figure 3; it uses the schema *pid_Int*, which reflects the functionality of the block structure defined in Figure 2, but not its parallelism.

In the complete *Circus* model, we first declare channels to represent the inputs, outputs, and internal wires of the top diagram. We also declare a channel *end_cycle*; it does not have a type, and is used for synchronisation: each block process, after receiving all its inputs and generating all its outputs, synchronises on *end_cycle* before starting the next cycle. After the channel declarations, the *Circus* model includes the ClawZ library, whose definitions are used in the block processes. These come afterwards, before the definition of the diagram, which concludes the specification.

In *Circus*, parallelism is alphabetised. In the definition of the diagram process, the alphabet of the each block process is the set of channels containing the inputs and outputs of the block, and *end_cycle*; like in CSP, the parallel processes synchronise on the intersection of their alphabets. In the *PID* example, *Si*

and *Int* are required to synchronise on *end_cycle* and *Si_out*: the output channel of *Si*, which is used as input by *Int*. Since *Si_out* is internal, as it is neither an input nor an output of the diagram, then it is hidden. Because all internal channels are hidden, implementations of the diagram do not need to preserve its block structure; during refinement, we can combine or split blocks.

The specifications of the block processes are similar; Figure 3 presents the process *Int* for the block with the same name. It is a basic process that encapsulates state and exhibits behaviour. The state is defined like in Z, using a schema: *Int_State*; its component is the state of the unit delay block in the Integrator diagram.

The main recursive action at the end defines the behaviour of *Int*. It initialises the state, as defined by the schema *Init*, before iterating to execute the cycles. In *Init*, the variable *pid_Diff_UnitDelay_state'* represents the value of the state component after the initialisation. In the predicate of *Init*, this is defined to be the value of the *initialstate* component of a binding *b* from *pid_Int_UnitDelay*. This is a ClawZ schema also included in *Int*; it models the UnitDelay block of the Integrator using the ClawZ library schema *UnitDelay_g*. It takes as parameter a real number that defines the initial value of the state of the block: *initialstate*. In this case, this value is 0, in accordance with information recorded in the diagram. The notation $0 \ e \ 0$ is used to represent 0 as a real number. Besides *pid_Int_UnitDelay*, *Int* includes *pid_Int_Sum* and *pid_Int*; these are the schemas generated by ClawZ for the Integrator.

In each cycle, *Int* calculates and produces its output as defined by *Execute_Int_out*, and in parallel updates the state, as defined by *Int_StateUpdate*. These actions synchronise on *Si_out*, since the input is needed to define the value of the output and of the state.

In a parallelism of actions, to avoid conflict, the disjoint set of variables which each of the parallel actions can modify is fixed. In the parallelism between *Execute_Int_out* and *Int_StateUpdate*, *Execute_Int_out* does not change the state, and *Int_StateUpdate* changes the only state component; this is explicitly indicated by associating them to the empty set of names of variables ($\{\}$) and to the set containing only *pid_Int_UnitDelay_state*. Both actions can access the initial value of *pid_Int_UnitDelay_state*, but only *Int_StateUpdate* can change this value. The same concerns arise for interleaving of actions.

Since *pid_Int* is not an operation over the state of *Int*, we use it to define a state operation: *Calculate_Int*. There, we use ΔInt_State to declare *pid_Int_UnitDelay_state* and *pid_Int_UnitDelay_state'* to represent the value of the state component before and after the operation. We also declare the input and out-

process $Int \hat{=} \text{begin}$

state $Int_State \hat{=} [pid_Int_UnitDelay_state : \mathbb{U}]$

$pid_Int_Sum \hat{=} Sum_P2$

$pid_Int_UnitDelay \hat{=} UnitDelay_g(X0 \hat{=} 0 \ e \ 0)$

pid_Int $In1? : \mathbb{U}; Out1! : \mathbb{U}$ $Sum : pid_Int_Sum; UnitDelay : pid_Int_UnitDelay$
$Sum.In1? = In1? \wedge Sum.In2? = UnitDelay.Out1! \wedge UnitDelay.In1? = Sum.Out1!$ $Out1! = UnitDelay.In1?$

$Init$ Int_State'
$\exists b : pid_Int_UnitDelay \bullet pid_Int_UnitDelay_state' = pid_Int_UnitDelay.initialstate$

$Calculate_Int$ ΔInt_State $In1?, Out1! : \mathbb{U}$
$\exists b : pid_Int \bullet b.In1? = In1? \wedge b.UnitDelay.state = pid_Int_UnitDelay_state \wedge$ $b.UnitDelay.state' = pid_Int_UnitDelay_state' \wedge b.Out1! = Out1!$

$Calculate_Int_out \hat{=} Calculate_Int \setminus (pid_Int_UnitDelay_state') \wedge \Xi Int_State$

$Execute_Int_out \hat{=} \text{var } In1 : \mathbb{U} \bullet Si_out?x \rightarrow In1 := x; \text{var } Out1 : \mathbb{U} \bullet Calculate_Int_out; Int_out!Out1 \rightarrow Skip$

$Calculate_Int_state \hat{=} Calculate_Int \setminus (Out1!)$

$Int_StateUpdate \hat{=} \text{var } In1 : \mathbb{U} \bullet Si_out?x \rightarrow In1 := x; Calculate_Int_state$

$\bullet (Init; \mu X \bullet (Execute_Int_out [\{ \} | \{E\} | \{pid_Int_UnitDelay_state\}] Int_StateUpdate); end_cycle \rightarrow X)$

end

Figure 3. Circus process for the block Int

put variables; the use of the decorations ? and ! in their names is a Z convention. In the predicate, $In1$ and $pid_Int_UnitDelay_state$ are used to identify a binding b of pid_Int . This binding defines the values of $pid_Int_UnitDelay_state'$ and $Out1!$.

In the specification of $Execute_Int_out$, we need an operation that only calculates the output $Out1$, but does not affect the state. We use $Calculate_Int_out$, which is specified by hiding $pid_Int_UnitDelay_state'$ in the definition of $Calculate_Int$, and conjoining the result with the schema $\Xi Diff_State$ that specifies that $pid_Int_UnitDelay_state$ is not modified. The action $Calculate_Int_state$ is used in $Int_StateUpdate$. It is defined by hiding $Out1!$ in the definition of $Calculate_Diff$, so that it defines a value for $pid_Int_UnitDelay_state'$, but does not have an output.

This is a simple example. More of the complexity of real diagrams is considered in [8].

3. Circus model for Ada implementations

The implementation of a control law diagram is usually composed of subprograms that implement sequentially the functionality of a (group) of blocks, and

schedulers for these subprograms. Typically, the cycle is split in time frames and the schedulers allocate the execution of the subprograms to these frames.

In Figure 4 we present the package (module) architecture of an Ada implementation of the PID. The packages $Exec_0$, $Exec_1$, $Exec_2$, and $Exec_3$ are the main programs that are executed in parallel. In each of them, there is an initialisation of variables followed by a loop whose iterations last for the duration of a time frame, and schedule some subprograms.

Each main program is associated with a scheduler: F_Sch , $Task_1$, $Task_2$, or $Task_3$. The subprograms that implement the blocks are in the package PID , which is implemented using another package $Discrete$. The program $Exec_0$ does not use PID because it only maintains timing information: its schedule, F_Sch keeps a frame counter. The other schedulers actually allocate the subprograms of PID to frames.

The structure of packages of the Ada program is not preserved in its *Circus* model because the module construct of *Circus*, process, represents an independent flow of execution. In the *Circus* model of the PID implementation, for example, we have a process for each of its main programs: $Exec_0$, $Exec_1$, $Exec_2$, and $Exec_3$.




Figure 4. Architecture of the Ada implementation

In each of them, inputs and outputs are communicated via the channels identified in the model of the diagram. Also, new internal channels are created to communicate shared variables; specifically, we need channels *Dsh* and *Ish* to communicate two shared variables *D* and *I*. Finally, the scheduling is modelled using synchronisation on *end_cycle* and on an extra channel *frame*, even though in the Ada program the scheduling uses variables to record time periods and a delay command.

Except for the need to match inputs and outputs to channels, and to identify shared and time variables, most of the *Circus* model can be calculated from the program, but we leave this as future work. Identifying the correspondence between the program variables and the wires of the diagram, and consequently, the channels of the *Circus* model is an activity already undertaken as part of the current verification process.

In Figure 5 we present *Exec₀*. Its state component *cur_f* corresponds to the Ada variable that records the number of the current frame. Its type, *FrameIndex*, contains only the numbers 1 and 2. Since *cur_f* is shared, it is output through the new channel *frame*. The structure of the main action is determined by the number of frames, and the scheduling: after the initialisation, in each frame, its number is output, then a conditional uses that information to invoke the actions to be executed. (In *Exec₀*, no action is scheduled.) In the end of the second frame, there is a synchronisation on *end_cycle*. Finally, the frame number is updated.

The processes *Exec₂* and *Exec₃* implement the groups of blocks *Si* and *Int*, and *Diff* and *Sd* (Figure 1). The process *Exec₂* is in Figure 6. Its state components are the variables that are used in *Exec₂*; and the actions correspond to the procedures used in *Exec₂*. The frame scheduling is factored out in a procedure *Step*. In contrast with *Exec₀*, the number of the current frame is input, rather than output. Before the procedure *Calc_Integral* is called, its inputs are received in interleaving. Since the variable *I* is shared, its value is communicated at the beginning of the second frame.

The process *Exec₁* is similar to *Exec₂* and *Exec₃*.

```

process Exec0  $\hat{=}$  begin
  state [cur_f : FrameIndex]
  Init_F  $\hat{=}$  cur_f := 1
  Next_F  $\hat{=}$  cur_f := (cur_f mod 2) + 1
  • Init;  $\mu$  X • frame!cur_f  $\rightarrow$ 
    if (cur_f = 1)  $\rightarrow$  Skip
    || (cur_f = 2)  $\rightarrow$  end_cycle  $\rightarrow$  Skip
  fi;
  Next_F; X
end

```

Figure 5. Circus model of *Exec₀*

The main difference is that its *Step* procedure inputs rather than outputs the values of the shared variables; also, it produces the output at the end of the cycle.

The model of the complete program is the parallel composition of *Exec₀*, *Exec₁*, *Exec₂*, and *Exec₃*, synchronising on their common channels; the new channels *frame*, *Dsh*, and *Ish* are hidden. The structure of this rather small example is representative of real control systems in its approach to scheduling and sharing.

4. Refinement strategy

A Simulink diagram has a hierarchical structure. We can assume that, at the top level, there is always a diagram with a single block that takes all the inputs and produces all the outputs. For all interesting examples, this is a single subsystem block that is specified by another diagram. Like for the PID, this may include further subsystem blocks, and, in this case, yet another level of diagrams is used. There is no limit to the number of diagrams that can be used to specify a control law; and they can be organised across several levels.

The *Circus* model of a top level diagram is a single process that lifts the ClawZ output. For the verification of a sequential implementation, this is a suitable model. For the verification of a parallel implementation, however, the *Circus* model of the diagram in the second level is likely to be more adequate. In fact, the best starting point is a model of a diagram in which all blocks that are implemented by different subprograms are explicitly drawn, rather than hidden in subsystem blocks. Typically, to draw this diagram, we need to rearrange the hierarchy of diagrams.

In the case of the PID, the model of the diagram in Figure 1 is adequate to verify the parallel implementation presented in Section 4. Since the implementations of *Diff* and *Int* are sequential, there is no need to expand their diagrams. If we did so, the *Circus* model of the diagram would include extra processes and channels corresponding to the blocks and wires in those diagrams. No doubt, this model would be equivalent to that of the smaller diagram in Figure 1, since the ex-

```

process Exec2  $\hat{=}$  begin
  state [Error, Ki, I :  $\mathbb{U}$ ]
  Init_Integral  $\hat{=}$  I := 0.0
  Integ  $\hat{=}$  (val Input, K :  $\mathbb{U}$ ; res Output :  $\mathbb{U}$  • Output := Output + K × Input)
  Calc_Integral  $\hat{=}$  Integ(Error, Ki, I)
  Step  $\hat{=}$  frame?f → if(f = 1) → ((E?x → Error := x) [| {Error} | {Ki}]) (Ki?x → Kd := x); Calc_Integral
    || (f = 2) → Ish!I → end_cycle → Skip
  fi
  • Init_Integral;  $\mu$  X • Step; X
end

```

Figure 6. Circus model of `Exec_2`

tra channels would be internal. The verification based on it, however, would require more effort to remove the parallelism that is not present in the implementation.

Our verification strategy comprises four phases that progressively collapse the parallelism of the diagram model to match the architecture of the implementation.

- P1** Write the main action of each block process as a recursion that iteratively takes the inputs in interleaving, calculates the outputs and updates the state, communicates the outputs in interleaving, and synchronise on *end_cycle*.
- P2** Collapse the parallelism between the processes of the blocks that are implemented by a single subprogram, and write the main action of the resulting processes in the form described in phase **P1**.
- P3** For each of the processes created in phase **P2**, introduce the action that specifies the corresponding subprogram, and prove that the calculations can be refined by a call to that action.
- P4** Collapse the parallelism between the processes that are scheduled in the same task.

Each phase of the strategy can be accomplished applying *Circus* refinement laws. In [7] we explain the procedure. Here, we explain how we can use ProofPower and its associated tools to carry out phase **P3**.

For the PID example, in phase **P1**, the main action of *Int* is written as follows.

```

Init;  $\mu$  X • Si_out?x → In1 := x;
  Calculate_Int_out; Calculate_Int_state;
  Int_out!Out1 → Skip; end_cycle → X

```

The variables *In1* and *Out1* are now state components. Since there is only one input and one output, no interleaving is needed.

In phase **P2**, we identify the subprograms that implement blocks. For each that implements more than one, we create a single process by collapsing the parallelism between the processes for the blocks. In `Exec_2`, for example, we identify the procedure `Calc_Integral` which implements *Si* and *Int*. We

create a process *SiInt* by collapsing *Si* and *Int*, and refine its main action to be written as follows.

```

Init;
 $\mu$  X • (E?x → pid_Si_In1 := x) ||| (Ki?x → pid_Si_In2 := x);
  pid_Si; pid_Int_In1 := pid_Si_Out1;
  Calculate_Int_out; Calculate_Int_state;
  Sd_out!pid_Sd_out → end_cycle → X

```

Inputs are taken from *E* and *Ki* in interleaving; for conciseness, we omit the name sets. Afterwards, the calculations of both *Si* and *Int* are performed, and the output of *Sd* is produced. The output of *Si* is the input of *Int*; since this is an internal communication, it is removed.

In phase **P3**, we refine the main action of each process that corresponds to a subprogram: we introduce the declaration and a call to the action that models the subprogram. For that, we match the variables of the program to those of the model. Since the model variables correspond to wires and block states, this means matching the program variables to diagram components.

In the case of *SiInt*, we introduce the definition of `Calc_Integral` as the following call.

```

Integ(pid_Si_In1, pid_Si_In2, pid_Int_UnitDelay_state)

```

Actually, first we introduce *Integ* so that it can be used in `Calc_Integral`. Since *Integ* does not refer to program variables, we define it just as shown in Figure 6.

What we need to prove is the following refinement.

$$\left(\begin{array}{l} \textit{pid_Si}; \textit{pid_Int_In1} := \textit{pid_Si_Out1}; \\ \textit{Calculate_Int_out}; \textit{Calculate_Int_state} \end{array} \right) \sqsubseteq \textit{Calc_Integral} \quad (1)$$

In the next section, we show how we carry out this proof using ProofPowerZ. With this result, we can write the main action of *SiInt* as follows.

```

Init;
 $\mu$  X • (E?x → pid_Si_In1 := x) ||| (Ki?x → pid_Si_In2 := x);
  Calc_Integral;
  Sd_out!pid_Sd_out → end_cycle → X

```

At this point, we can simplify *SiInt* by removing the

actions that are no longer used: the ClawZ schemas.

In phase **P4**, we group the subprogram processes in task processes, and rewrite their main actions to use the frame definition.

5. Using ProofPower

Using ProofPower, we can prove that a specification statement that specifies the functionality of a (group of) blocks in terms of the program variables can be refined by Ada commands. The specification statement is constructed using the output of ClawZ for the group of blocks and identifying how the variables of the program correspond to components of the diagram. The Ada commands are those in the body of the subprogram that implements the group of blocks.

For *Integ*, the following conjecture is provable.

$$\Delta Output : \left[\begin{array}{l} \exists Integ_System \bullet \\ Si.In1? = Input \wedge Si.In2? = K \wedge \\ Int.UnitDelay.state = Output_0 \wedge \\ Int.UnitDelay.state' = Output \wedge \\ Int.Out1! = Output \end{array} \right] \quad (2)$$

⊆

$$Output := Output + K \times Input$$

In the specification, $\Delta Output$ defines that the parameter *Output* can be modified. The predicate is a quantification over a schema *Integ_System*.

<i>Integ_System</i>
<i>Int</i> : <i>pid_Int</i> ; <i>Si</i> : <i>pid_Si</i>
<i>Int.In1?</i> = <i>si.Out1!</i>

This schema is generated by ClawZ if we indicate that *Si* and *Int* form an artificial subsystem: a group of blocks that define a subsystem, even though they do not correspond to a subsystem block. In the predicate of the specification statement in (2), the inputs, output, and state components of the artificial subsystem are identified with parameters of *Integ*. Like in the refinement calculus, *Output₀* refers to the initial value of *Output*. The refinement conjecture states that such a system can be implemented by the body of *Integ*. Our aim is to use this result to prove the refinement in (1). For that, we carry out the following steps.

1. **Match the procedure parameters to model variables.** Implicitly, we are matching wires and block states in the diagram to procedure parameters.

In the example of the *Integ* procedure, *Input* corresponds to *pid_Si_In1*, *K* to *pid_Si_In2*, and *Output* to *pid_Int_UnitDelay_state* and *pid_Int_Out1*.

2. **Introduce the parameter declarations, taking the corresponding model variables as arguments.** In *Circus*, the treatment of parametrised procedures is based on Back's parametrised commands [4], which have been incorporated in a refinement calculus for Z [10]. In this context, we can refine any program to an instantiation of a parametrised command as long as proper renamings are carried. For, the specification in (1), which from now on we name *SiIntC* for conciseness, we proceed as follows.

$$(\text{val } Input, K : \mathbb{U}; \text{vres } Output : \mathbb{U} \bullet \\ SiIntC \left[\begin{array}{l} Input / pid_Si_In1?, \\ K / pid_Si_In2?, \\ Output / pid_Int_UnitDelay_state, \\ Output' / pid_Int_UnitDelay_state', \end{array} \right] \\)(pid_Si_In1?, pid_Si_In2?, pid_Int_UnitDelay_state)$$

The renamings rewrite *SiIntC* in terms of the parameters, but the instantiation of arguments means that the program acts on the model variables, just as *SiIntC* does. We observe that *Output* is only matched to *pid_Int_UnitDelay_state*. Later, we have to record the extra relationship to *pid_Int_Out1*.

3. **Remove state components that are not matched to procedure parameters from the state.** A state component is a variable that is in scope in the main action of a process; we can remove its declaration from the state and introduce a corresponding variable block in the main action.

In our example, we need to remove *pid_Si_Out1*, *pid_Int_In1*, and *pid_Int_Out1* from the state. The first two components record the value communicated internally by *Si* to *Int*. This communication is eliminated in a sequential implementation. The last component matches *Output*, but since *Output* matches two variables, one is left out for now.

4. **Reduce the scope of the variable block to include only the calculations and updates.** This can be achieved with standard variable block laws. In our example, the body of the block is reduced to *SiIntC*[*Input*/*pid_Si_In1?*, ...].

5. **Reduce the calculations and updates to a schema.** This requires specific laws to flatten sequences. One of the them is presented below.

Law 5.1 (schema-assign-seq)

$$[\Delta S \mid p]; x := e \\ = \\ [\Delta S \mid \exists x_0 \bullet p[x_0/x'] \wedge x' = e[x_0/x]']$$

provided $x, x' \in \alpha S$

It flattens a sequence of a schema that defines an operation over a state *S* that includes a variable *x*, and

an assignment to x . The value defined for x by the original schema is given the name x_0 and hidden in the predicate of the resulting schema.

In this step, we use this law and others like this in the following way.

- (a) **Join $_out$ and $_state$ schemas back together.** In our specification, we have the sequence $Calculate_Int_out; Calculate_Int_state$. The sequence operator is not the schema calculus operator, but the command constructor. By removing it, we get the schema $Calculate_Int$ (Figure 3), which lifts pid_Int .
- (b) **Eliminate assignments.** For that, we use the law above, or another similar law that applies when the assignment comes first in the sequence. The assignments correspond to passages of values from a block to another through the internal wires. The equality that arises from the assignments in the resulting schema should be kept: they should not, for instance, be used to eliminate quantifiers using the one-point rule. In our example, we keep the equality $pid_Int_In1' = pid_Si_Out1'$.
- (c) **Flatten remaining sequences.** The left sequences involve only schemas. The ClawZ schemas should be incorporated, but not expanded. In our example, we do not expand the schema pid_Si : the ClawZ model of Si .

The final result of this step for our example is the schema shown below.

$$\begin{array}{l}
 \text{Step5} \\
 \hline
 pid_Int_In1, pid_Int_In1', \\
 pid_Int_UnitDelay_state, \\
 pid_Int_UnitDelay_state' \\
 pid_Si_In1?, pid_Si_In2?, \\
 pid_Si_Out1!, pid_Int_Out1! : \mathbb{U} \\
 \hline
 pid_Si \wedge pid_Int_In1' = pid_Si_Out1' \\
 \left(\begin{array}{l}
 \exists b : pid_Int \bullet \\
 b.In1? = pid_Int_In1' \\
 b.UnitDelay.state = \\
 pid_Int_UnitDelay_state \\
 b.UnitDelay.state' = \\
 pid_Int_UnitDelay_state' \\
 b.Out1! = pid_Int_Out1!
 \end{array} \right)
 \end{array}$$

It acts on a state that includes all the variables in scope in the main action. We do not decorate pid_Int_In1 with a $?$ because its final value is changed, so it is not an input variable in this schema.

- 6. **Remove renamings and variable block.** This can be achieved by turning them into an existential quan-

tification: for the renamings, we use the one-point-rule, and for variable blocks, we use the law below.

Law 5.2 (var-schema-hiding)

$$(\text{var } x : T \bullet SExp) = SExp \setminus (x, x')$$

A schema hiding is an existential quantification. For our example, the result is sketched below.

$$\begin{array}{l}
 \text{Step6} \\
 \hline
 Output, Output', Input, K : \mathbb{U} \\
 \hline
 \exists pid_Int_In1, pid_Int_In1', pid_Si_Out1!, \dots \bullet \\
 pid_Si_In1? = Input \wedge pid_Si_In2? = K \wedge \\
 \dots \wedge \\
 \dots \text{predicate of Step5} \dots
 \end{array}$$

- 7. **Rename the dashed In variables to use the $?$ decoration.** This is possible because the undashed variables are not free in the quantified predicate, so we can eliminate their quantification, and change the name of the dashed variables. In the example, the quantification becomes $\exists pid_Int_In1?, pid_Si_Out1!, \dots : \mathbb{U} \bullet \dots$. We cannot rename pid_Int_In1' without eliminating pid_Int_In1 beforehand because, in *Circus*, pid_Int_In1 and pid_Int_In1' are the same variable. There are no input and output variables, only state components and local variables; the $?$ and $!$ decorations are used for compatibility with *Z*.
- 8. **Express quantification using simplified ClawZ schemas.** We introduce schemas that contain just the declaration part of the ClawZ schemas for the blocks. For pid_Int , we declare the schema below.

$$\begin{array}{l}
 pid_IntD \\
 \hline
 In1?, Out1! : \mathbb{U} \\
 Sum : pid_Int_Sum \\
 UnitDelay : pid_Int_UnitDelay
 \end{array}$$

We also declare a schema pid_SiD .

The quantification can be expressed in terms of these schemas: there is a direct correspondence between the quantified variables and their components. After all, it was these schemas that were used to construct the model in the first place. For example, the component $In1?$ of pid_IntD corresponds to $pid_Int_In1?$.

Part of the rewritten predicate is shown below.

$$\begin{aligned} & \exists Int : pid_IntD; Si : pid_SiD \bullet \\ & Si.In1? = Input \wedge Si.In2? = K \wedge \\ & \dots \wedge \\ & pid_Si[Si.In1?/pid_Si_In1?, \dots] \wedge \\ & Int.In1? = Si.Out1! \wedge \\ & \left(\begin{array}{l} \exists b : pid_Int \bullet \\ b.In1? = Int.In1? \\ b.UnitDelay.state = Int.UnitDelay.state \\ b.UnitDelay.state' = Int.UnitDelay.state' \\ b.Out1! = pid.Int.Out1! \end{array} \right) \end{aligned}$$

9. **Check that the quantified variables are elements of the proper ClawZ schema.** In the example, we need to check that *Int* and *Si* are bindings of *pid_Int* and *pid_Si*, and not simply *pid_IntD* and *pid_SiD* as declared. We need to consider two cases.

(a) **The block has a state.** In this case, an existential quantification over a binding of the ClawZ schema defines the variable. This is due to the way in which state is handled in the *Circus* model using the state in the ClawZ schemas. In our example, *Int* gives an illustration.

Since the blocks are deterministic, the quantified predicate allows us to conclude that the binding is equal to the quantified variable: $b = Int$ in the example. With that, the one-point rule can be used to eliminate the existential quantification and conclude that $Int \in pid_Int$.

(b) **The block does not have a state.** In this case, instead of an existential quantification, a renaming of the original schema defines the quantified variable: *Si* and *pid_Si* illustrate the situation. The renaming explicitly states that the quantified variable satisfies the property of the ClawZ schema. A more concise way of expressing this is to say that the variable belongs to the type defined by the ClawZ schema. In the example, we can write $Si \in pid_Si$.

As a result, we can simplify the predicate of the schema, by giving stronger declarations for the variables. In our example, the result is as follows.

$$\begin{aligned} & \exists Int : pid_Int; Si : pid_Si \bullet \\ & Si.In1? = Input \wedge Si.In2? = K \wedge \\ & Int.UnitDelay.state = Output \wedge \\ & Int.UnitDelay.state' = Output' \wedge \\ & Int.In1? = Si.Out1! \end{aligned}$$

10. **Transform schema into a specification statement.** This can be accomplished with a law of the Z refinement calculus; they are all valid in *Circus*. The

result is almost exactly the specification statement used in the current verification technique (see (2)), but the components and the predicate of the artificial subsystem schema are explicitly included.

11. **Declare the artificial subsystem schema and use it in the specification.** This is a simple application of predicate calculus to change the quantification.

12. **Record missing correspondences between parameters and model variables.** If any of the parameters represented more than one diagram component (or equivalently, model variable), only one of the correspondences was used. We can now strengthen the postcondition of the specification statement to record the others.

In the example, we did not record the correspondence between *Output* and *pid_Int_Out1*. If we strengthen the postcondition of the specification statement with the conjunct $Int.Out1! = Output$, we record the relationship.

13. **Use ProofPower.** The specification statement is now exactly that constructed using the current approach. So, the refinement can be proved using ProofPower; it justifies the replacement of the specification statement with the procedure body.

14. **Introduce the procedure call.** This is now a direct application of the copy rule.

Except for the actual use of ProofPower, and the need to match program variables to diagram components, all the steps of this procedure can be automated.

6. Conclusions

We have presented an approach to reuse existing tools and expertise in the verification of control systems in the context of a technique that covers both the sequential subprograms and the schedulers of an implementation. It is based on a notation that integrates Z and CSP, and, as such, can formalise the functionality and the parallelism in diagrams and their implementations. We use a single model of the diagram, and a unified verification technique.

Our approach is a procedure for application of refinement laws; using a tactic language [16], we can automate it. All the steps of this procedure are justified by refinement laws. If the application of any of them fails, there is an error in the program, or in our interpretation of how the procedures and variables correspond to components of the diagram.

The majority of the proof obligations generated are related to the implementation of the functionality of the blocks. We have shown how these can be discharged using the tools and proof tactics already available. Ex-

perience shows that they can discharge 95% to 98% of the proof obligations automatically.

The needed information concerning the relationship between wires of the diagram and variables of the program and of the model, and between blocks and procedures, is also required by the existing verification technique. Therefore, the feasibility of producing this information is already confirmed for real examples.

In [5], a weakest precondition semantics is used to validate Simulink diagrams; a PID controller is also considered as an example. Assertion reasoning is also used in [18] to analyse Stateflow diagrams. The work in [14] considers dataflow networks with feedback loops, which are similar to control diagrams, although parallelism has to be explicitly indicated. The reasoning technique is also based on predicate transformers.

Our example is simple, and further case studies are in our plans for future work. In any case, the PID verification illustrates the issues involved in the use of ProofPower and its associated tools. They are applied in the third phase of our refinement strategy to processes generated in the second phase by collapsing two or more block processes. Our approach handles sequences of assignments and schemas that define the functionality of a group of blocks. The processes generated in the second phase of our strategy are always of this form.

Automation is essential for the success of a verification technique. Diagrams can span over hundreds of pages, and it is not feasible to handle them without tool support. We plan to automate the rest of our strategy using a refinement tool and tactics based on ProofPower.

We are also working on tools for *Circus*. Of particular relevance to verification of control systems is the *Circus* model checker. It follows the approach adopted by FDR for CSP, but integrates the use of theorem proving to handle proof obligations generated by the data part of the *Circus* programs. An important piece of future work is to investigate its application to prove the refinement conjectures generated by our strategy. Currently, the use of FDR is limited by the size of the models of the control diagrams and of the schedulers. The use of *Circus* guides the combined use of traditional techniques of model checking and theorem proving and is likely to lead to a high level of automation.

Acknowledgements

This work is funded by the Royal Society. We discussed it with Colin O'Halloran, and Jim Woodcock.

References

[1] *The MathWorks Site. Simulink Reference Manual.*

- www.mathworks.com/products/simulink.
- [2] M. M. Adams and P. B. Clayton. Cost-Effective Formal Verification for Control Systems. In *ICFEM 2005*, volume 3785 of *LNCS*, pages 465 – 479. Springer, 2005.
- [3] R. Arthan, P. Caseley, C. O'Halloran, and A. Smith. ClawZ: Control laws in Z. In *ICFEM 2000*, pages 169 – 176. IEEE Press, 2000.
- [4] R. J. R. Back. Procedural Abstraction in the Refinement Calculus. Technical report, Department of Computer Science, Åbo, Finland, 1987. Ser. A No. 55.
- [5] J. Blow and A. Galloway. Generalised Substitution Language and Differentials. In *ZB 2002*, volume 2272 of *LNCS*, pages 396 – 415. Springer-Verlag, 2002.
- [6] R. J. Boulton, R. Hardy, and U. Martin. A Hoare-Logic for Single-Input Single-Output Continuous-Time Control Systems. In *HIS 2003*, volume 2623 of *LNCS*, pages 113 – 125. Springer-Verlag, 2003.
- [7] A. L. C. Cavalcanti and P. Clayton. From Control Law Diagrams to Ada via *Circus*. In *FM 2006*, 2006. Submitted.
- [8] A. L. C. Cavalcanti, P. Clayton, and C. O'Halloran. Control Law Diagrams in *Circus*. In *FM 2005*, volume 3582 of *LNCS*, pages 253 – 268. Springer-Verlag, 2005.
- [9] 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.
- [10] A. L. C. Cavalcanti and J. C. P. Woodcock. ZRC—A Refinement Calculus for Z. *Formal Aspects of Computing*, 10(3):267—289, 1999.
- [11] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall International, 1985.
- [12] M. Jersak, D. Ziegenbein, F. Wolf, K. Richter, and R. Ernst. Embedded System Design using the SPI Workbench. In *FDL 2000*, 2000.
- [13] D. J. King, R. D. Arthan, and I. C. L. Winnersh. Development of Practical Verification Tools. *ICL Systems Journal*, 11(1), 1996.
- [14] B. Mahony. First International Workshop on Formalising Continuous Mathematics. In *The DOVE Approach to the Design of Complex Dynamic Processes*, pages 167 – 187, 2002.
- [15] C. C. Morgan. *Programming from Specifications*. Prentice-Hall, 2nd edition, 1994.
- [16] M. Oliveira, A. L. C. Cavalcanti, and J. C. P. Woodcock. ArcAngel: A Tactic Language for Refinement. *Formal Aspects of Computing*, 15(1):28 – 47, 2003.
- [17] A. W. Roscoe. *The Theory and Practice of Concurrency*. Prentice-Hall Series in Computer Science. Prentice-Hall, 1998.
- [18] I. Toyn and A. Galloway. Proving Properties of Stateflow Models using ISO Standard Z and CADiZ. In *ZB 2005*, volume 3455 of *LNCS*, pages 104 – 123. Springer-Verlag, 2005.
- [19] 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.
- [20] J. C. P. Woodcock and J. Davies. *Using Z*. Prentice-Hall, 1996.