

Reasoned Modelling Critics: Turning Failed Proofs into Modelling Guidance

A. Ireland¹, G. Grov², M. T. Llano¹ and M. Butler³

¹School of Mathematical and Computer Science, Heriot-Watt University, UK
{a.ireland,mtl4}@hw.ac.uk

²School of Informatics, University of Edinburgh, UK
ggrov@inf.ed.ac.uk

³School of Electronics & Computer Science, Southampton University, UK
mjb@ecs.soton.ac.uk

ABSTRACT. The activities of formal modelling and reasoning are closely related. But while the rigour of building formal models brings significant benefits, formal reasoning remains a major barrier to the wider acceptance of formalism within design. Here we propose *reasoned modelling critics* – a technique which aims to abstract away from the complexities of low-level proof obligations, and provide high-level modelling guidance to designers when proofs fail. Inspired by proof planning critics, the technique combines proof-failure analysis with modelling heuristics. Here, we present the details of our proposal, implement them in a prototype and outline future plans.

Keywords: reasoned modelling; formal methods; formal modelling; formal verification; automated reasoning; artificial intelligence.

1 Introduction

The use of mathematical techniques for system-level modelling and analysis brings significant benefits, as well as challenges. While the rigour of mathematical argument can offer early feedback on design decisions, a key challenge centres on *how* feedback derived from the analysis of complex *proof obligations* (POs) can be used to improve design decisions. Such feedback currently requires user intervention, drawing upon skilled knowledge of the subtle interplay that exists between modelling and reasoning. For example, consider a simple cruise control system with the safety requirement that the brakes (*brake*) cannot be on while the cruise control (*cc*) is enabled. This can be formalised as the following invariant:

$$cc = on \Rightarrow brake = off$$

When the driver applies the brakes, i.e. $brake := on$ it must be shown that the safety requirement (invariant) is preserved, creating a PO of the form:

$$brake = off, cc = on \vdash on = off$$

Note that this PO is false. An (unsafe) solution to this failure would be to restrict the application of the brakes via a guard, i.e. only allow the brakes to be applied when the cruise

control is disabled ($cc = off$). Clearly, the desirable solution requires the introduction of an auto disable mechanism, i.e. the cruise control is disabled if the brakes are applied ($cc := off ; brake := on$). While choosing between the two alternatives do not represent a significant burden to a designer, in general many solutions may arise from proof-failure analysis. In order to make good use of a designer's time, modelling heuristics are needed in order to rank the alternatives.

Here we describe an approach to formal modelling that combines proof-failure analysis with modelling heuristics. The basic idea is to use automatic analysis of failed proofs, constrained by modelling heuristics, to generate modelling suggestions. In this way we abstract away from complexities of low-level POs, and automatically provide designers with high-level decision support oriented towards modelling choices. In turn this will increase the productivity of designers as well as the accessibility of formal modelling tools. The approach has been implemented through what we call *reasoned modelling critics*.

This need for both proof and modelling heuristics is central to our proposal. We aim to identify common modelling and reasoning patterns, and use these to abstract away from complex POs. This will enable us to automatically provide designers with high-level decision support oriented towards modelling choices. In turn this will increase the productivity of designers as well as the accessibility of formal modelling tools. We implement our proposal through what we call *reasoned modelling critics*.

A key inspiration for our proposal is *proof planning*, a technique for automating the search for proofs through the use of high-level proof outlines, known as *proof plans* [7]. Central to proof planning is its proof-failure analysis and proof patching capabilities [12]. Our approach builds directly upon these features. Specifically by combining the proof-failure analysis capabilities with modelling heuristics in order to automatically generate modelling guidance.

Our longer term aim is to combine the proof plan representation with a complementary notion of *modelling patterns*, which we call *reasoned modelling methods*. This, combined with reasoned modelling critics, are the building blocks of what we call *reasoned modelling*: the study of the interplay between formal reasoning and modelling.

While our vision is generic, our starting point is Event-B, which we motivate and introduce together with proof planning in §2. §3 outlines the reasoned modelling critics mechanism and shows its application through examples. §4 discusses the prototype implementation of the reasoned modelling critics mechanism, and §5 shows the results of applying the examples discussed in §3 to the prototype. We discuss related and future work in §6 and conclude in §7.

This paper is an extension of [13]. Here, in addition to provided more explanation of the concepts, we provide a prototype implementation of the ideas, discussed in §and §.

2 Background

2.1 Event-B and the Rodin toolset

Event-B provides a formal framework for modelling discrete complex systems [1], and is mechanised through the Rodin toolset [2]. Event-B promotes an incremental style of formal modelling, where each step of a development is underpinned by formal reasoning. As a

result, there is strong interplay between modelling and reasoning and this is partly supported by the Rodin toolset. This interplay requires skilled user interaction, *i.e.* typically a user will analyse failed proofs, and translate the analysis by hand into corrective actions at the level of modelling. This is exemplified in e.g. [1, 10]. Typical corrective actions include strengthening invariants and guards or modifying actions. Our aim is to provide high-level decision support, by automating the generation, filtering and ranking of modelling suggestions. Event-B models and POs are closely aligned [1], whilst Rodin [3] is an extensible framework. Event-B and the Rodin toolset thus represent a unique opportunity for us to investigate reasoned modelling.

We will use the term *models* for both Event-B *contexts* (the static part) and *machines* (the dynamic part). A machine includes invariants and the Rodin tool generates POs to ensure that all invariants are maintained by each event of a machine. Event-B also supports refinement of machines and Rodin generates POs to ensure the consistency of refinement. For details of the POs see [2]. We will only use the basic features of Event-B and use standard notation:

$$\begin{aligned} \text{EVENT } \langle name \rangle &\hat{=} \text{BEGIN } \langle action \rangle \text{ END} \\ \text{EVENT } \langle name \rangle &\hat{=} \text{WHEN } \langle guard \rangle \text{ THEN } \langle action \rangle \text{ END} \end{aligned}$$

where the event is only executed when the guard holds. Moreover, the term $\langle action \rangle$ is a list of simultaneous assignments to machine variables, while INITIALISATION is a special event without guards defining the initial state.

2.2 Proof planning and proof critics mechanism

Proof planning builds upon the tactic-based tradition of theorem proving, where primitive proof steps are packaged-up into programs known as *tactics*. Starting with a set of general purpose tactics, plan formation techniques are used to construct a customised tactic for a given conjecture. The search for a customised tactic is constrained by a set of *methods*, collectively known as a *proof plan*. Using preconditions, each method specifies the applicability of a general purpose tactic. Having explicit preconditions provides insights when proof planning fails. The proof *critics* mechanism [12] enables “interesting” failures at the level of precondition evaluation to be represented. For a given method a number of critics will typically exist. Each critic represents an alternative generic proof patch, e.g. conjecture generalisation. The analysis of a specific failure, and subsequent proof planning, provides guidance in the selection and instantiation of a generic proof patch. In Figure 1 we present a method and critic schemas, while in Figure 2 we describe in detail the application of an example method and critic. Note that these are deliberately simple – their role is to illustrate the basic critics mechanism. Proof patching has been applied successfully to the problems of inductive conjecture generalisation and lemma discovery [8], as well as loop invariant discovery [14]. Proof critics have also been developed for *abductive reasoning*, *i.e.* the speculation of hypotheses in order to explain a consequence, and applied to the problem of patching faulty conjectures [17]. The proof critics mechanism provides the starting point for reasoned modelling critics. Moreover, in terms of exploiting existing critics, we see the work on abduction playing an important role in our reasoned modelling critics. This is elaborate upon in §5.

method (name)	critic (name)
INPUT: <i>PO</i>	INPUT: <i>POs</i>
PRECONDITIONS:	PRECONDITIONS:
<i>precondition 1</i>	<i>precondition 1</i>
...	...
<i>precondition 2</i>	<i>precondition 2</i>
EFFECTS:	PATCH:
<i>generation of output POs</i>	<i>proof patch</i>
OUTPUT:	
<i>POs</i>	

Methods: A method takes as input a PO. The applicability of a method is determined by evaluating its associated preconditions. If the preconditions of a method succeed, then the output POs (potentially empty) are computed via the effects slot.

Critics: During proof planning, both the success and failure of a method's preconditions are recorded. In particular, when a method fails a set of partial instantiations of the preconditions are associated with the PO. The critics mechanism uses the record of partial successes to search for patchable exceptions to the method. If the pattern represented by a given critic's preconditions matches with a recorded failure pattern then the associated patch is applied. In general, a critic may require multiple failures in order to trigger a patch. In this way, critics can apply both local and global proof-failure analysis.

Figure 1: Proof method and critic schemas

3 From proof critics to reasoned modelling critics

As described above, our reasoned modelling critics will extend the existing proof critics mechanism. This extension will involve two key components:

1. Firstly, in order to combine proof and modelling, our new critics need access to models as well as POs. In addition, critics will now have the option of providing modelling guidance, as well as proof patches.
2. Secondly, both methods and critics are associated with preconditions. We need to extend the meta-language of critics to allow us to represent preconditions which specify modelling heuristics.

To support access to models and modelling guidance, an extended critic schema will be required. Our proposed extension is described in Figure 3. In order to illustrate the new schema and the kind of meta-language extensions we propose for reasoned modelling, we develop below some critics and show their application to simple control system problems.

<p>method (rewrite)</p> <p>INPUT: $\Delta \vdash G$</p> <p>PRECONDITIONS:</p> <ol style="list-style-type: none"> 1. $exp_at(G, Pos, L)$ 2. $rewrite_rule(C \Rightarrow L := R)$ 3. $provable(\Delta \vdash C)$ <p>EFFECTS:</p> <p>$replace_at(Pos, R, G, NewG)$</p> <p>OUTPUT:</p> <p>$[\Delta \vdash NewG]$</p>	<p>critic (casesplit)</p> <p>INPUT: POs</p> <p>PRECONDITIONS:</p> <ol style="list-style-type: none"> 1. $\exists failed_po \in \{ \langle rewrite, _, PO \rangle \in POs \mid failed_proof(PO) \}$. 2. $\langle exp_at(_, _, _), rewrite_rule(C \Rightarrow _ := _), fail \rangle \in Pre$ <p>PATCH:</p> <p>$insert_casesplit(C, PO)$</p>
---	---

The meta-predicates used above are defined in Figure 4. Note that Δ denotes a list of hypotheses, while G and $NewG$ denote single goal formula. Note also the use of Prolog meta-variables, in particular the use of “_” for anonymous variables. The scope of meta-variables extends across all slots of the method and critic schemas.

To illustrate the proof patching aspect of proof planning, consider the following PO:

$$k : \tau, h : \tau, t : list(\tau), l : list(\tau) \vdash k \in [h|append(t, l)]$$

where \in denotes list membership, while $[_|_]$ and $append$ denote the constructor and concatenation function for lists. Moreover, assume the definition of \in gives rise to the following rewrite rules:

$$X = Y \Rightarrow X \in [Z|W] \quad \Rightarrow \quad true \tag{1}$$

$$X \neq Y \Rightarrow X \in [Z|W] \quad \Rightarrow \quad X \in W \tag{2}$$

Now consider the **rewrite** method preconditions more closely i.e.

1. there exists an expression L at position Pos within the goal G ,
2. there exists a rewrite rule such that L matches the left-hand side of the rule, and
3. any condition C attached to the rule is provable within the proof context.

Note that based upon rewrite rule (2), preconditions 1 and 2 succeed, with C , L and R being instantiated to be $k \neq h$, $k \in [h|append(t, l)]$, and $k \in append(t, l)$ respectively. However, precondition 3 fails, i.e. the condition $k \neq h$ is not provable from the given hypotheses. This pattern of partial success triggers the **casesplit** critic, i.e. it associates the failure (*fail*) of precondition 3 with a casesplit patch. The application of the patch gives rise to the following two POs:

$$k : \tau, h : \tau, t : list(\tau), l : list(\tau), k = h \vdash k \in [h|append(t, l)] \tag{3}$$

$$k : \tau, h : \tau, t : list(\tau), l : list(\tau), k \neq h \vdash k \in [h|append(t, l)] \tag{4}$$

Note that the **rewrite** method is applicable to both POs, i.e. rewrite rule (1) applies to (3) while rewrite rule (2) applies to (4).

Figure 2: An example proof method and critic

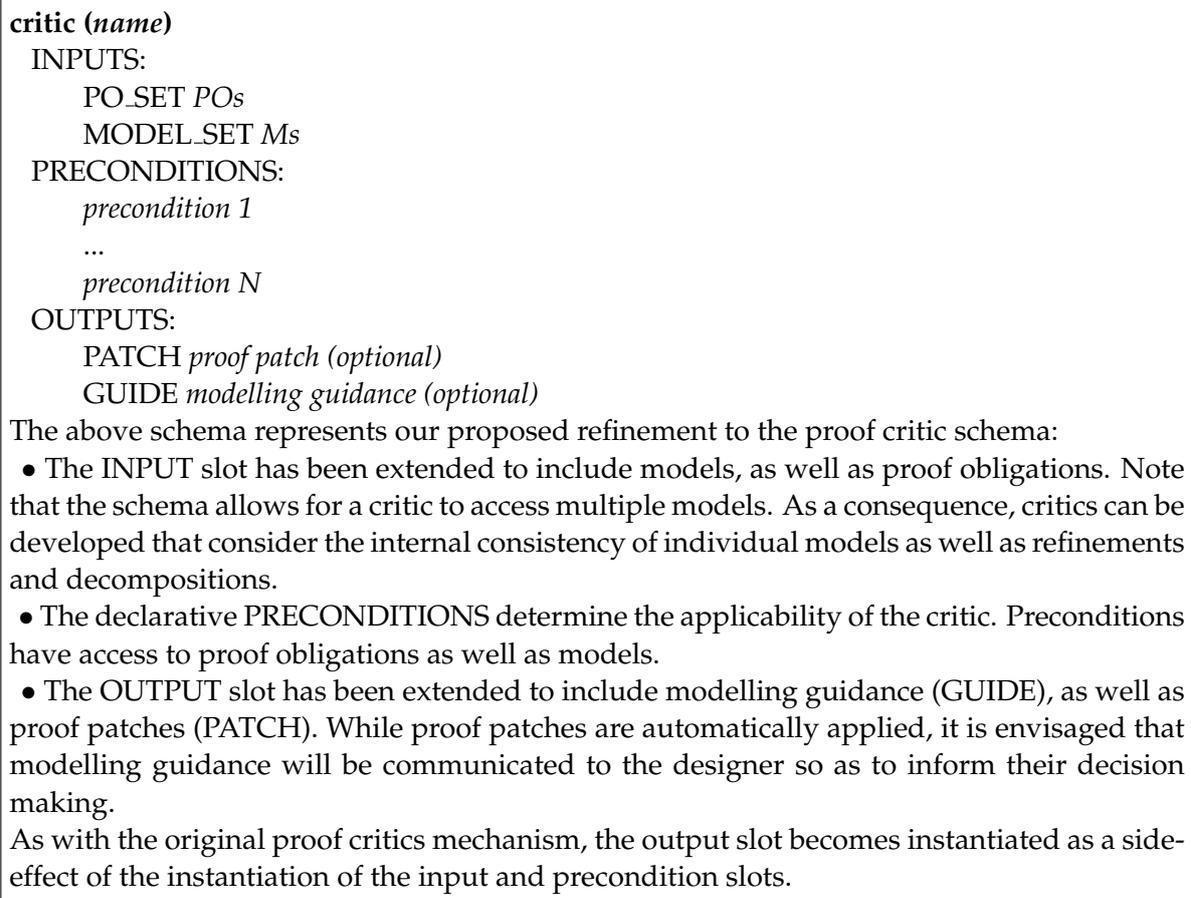


Figure 3: Reasoned modelling critic schema

3.1 Guidance for invariant proof failures via local analysis

Typically, a proof failure arising from an inconsistency in a model may be overcome in a number of ways. For example, consider an invariant which takes the form of an implication, i.e. $X \Rightarrow Y$. If a corresponding invariant proof fails, then the failure may be overcome by revising the model so that either,

1. X is false, or
2. Y is true.

At the level of modelling these effects can be achieved by changing the guards or actions associated with the invariant proof, or the invariant itself. Here we focus on *local* changes to an event, i.e. changes to guards and actions, and delay the discussion of global analysis and invariant changes until §3.3. In general, proof-failure analysis will generate a large number of proof patches. But not all proof patches will make sense in terms of modelling. To address this problem we wish to exploit modelling knowledge in order to rank the proof patches that are offered as guidance. To achieve this we envisage designers providing meta-data in addition to their models. The preconditions of our new critics can then exploit modelling heuristics as well as proof-failure heuristics.

$disjoint_sub(S_1, S_2)$	$\hat{=}$ $dom(S_1) \cap dom(S_2) = \emptyset$
$failed_proof(P)$	$\hat{=}$ “ P is a PO which is provably false or has failed to be proven”
$provable(P)$	$\hat{=}$ “ P is a PO that is provable”
$max(A)$	$\hat{=}$ $\epsilon x. (x \in A \wedge \forall y \in A. y \leq x)$
$exp_at(X, Y, Z)$	$\hat{=}$ “ Z is the subexpression at position Y within expression X ”
$rewrite_rule(X \Rightarrow Y \Rightarrow Z)$	$\hat{=}$ “A conditional rewrite rule”
$pri_var(V, M)$	$\hat{=}$ “Priority of variable V within model M ”
$priority(S, M)$	$\hat{=}$ $max(\{pri_var(v, M) \mid v \in dom(S)\})$
$replace_at(W, X, Y, Z)$	$\hat{=}$ “ Z is obtained by replacing the subexpression at position W within Y by X ”
$sub2act(\{V \mapsto D\})$	$\hat{=}$ $V := D$
$sub2act(\{V_1 \mapsto D_1, \dots, V_n \mapsto D_n\})$	$\hat{=}$ $sub2act(\{V_1 \mapsto D_1\}) \parallel \dots \parallel sub2act(\{V_n \mapsto D_n\})$
$sub2grd(\{V \mapsto D\})$	$\hat{=}$ $V = D$
$sub2grd(\{V_1 \mapsto D_1, \dots, V_n \mapsto D_n\})$	$\hat{=}$ $sub2grd(\{V_1 \mapsto D_1\}) \wedge \dots \wedge sub2grd(\{V_n \mapsto D_n\})$
$guards(E)$	$\hat{=}$ “Conjunction of guards of event E ”
$sat(P)$	$\hat{=}$ “The predicate P is satisfiable”
$generalisable(H)$	$\hat{=}$ $\exists r, e_1, e_2, g_1, g_2. \langle e_1, g_1 \rangle \in H \wedge \langle e_2, g_2 \rangle \in H$ $\wedge e_1 \neq e_2 \wedge r \Rightarrow g_1 \wedge r \Rightarrow g_2 \wedge sat(r)$
$generalise(H)$	$\hat{=}$ “The weakest r such that $sat(r)$ and for the most $\langle e, g \rangle \in H, r \Rightarrow p$ (for at least 2 events)”
$add_guard(G, E, M)$	$\hat{=}$ “Adds guard G to event E of model M ”
$add_action(A, E, M)$	$\hat{=}$ “Adds action A to event E of model M ”
$add_invariant(I, M)$	$\hat{=}$ “Adds invariant I to model M ”
$insert_casesplit(C, P)$	$\hat{=}$ “Progress the proof of P by a casesplit on X ”

Note that S denotes a substitution, while H denotes a set of event-guard pairs.

Figure 4: Meta-terms for reasoned modelling critics

To illustrate this idea, we develop two critics below. Each critic targets a different pattern of invariant proof failure, as outlined above. In addition, the critics exploit a notion of variable priority – which is based upon the observation that not all variables within a model may carry equal importance. Crucially, such “meta-data” must be supplied by the designer. For example, being able to brake is clearly more important than driving with the cruise control enabled. This notion of *priority* can be expressed as a heuristic, and used to rank or even filter modelling suggestions. Informally, where proof-failure analysis suggests that the value of a variable should be changed within the context of a given event, we adopt the following heuristics:

H1: If the priority of the candidate variable is lower than the priorities of all the variables updated by the event, then it is strongly suggestive that the change should be achieved via a new action.

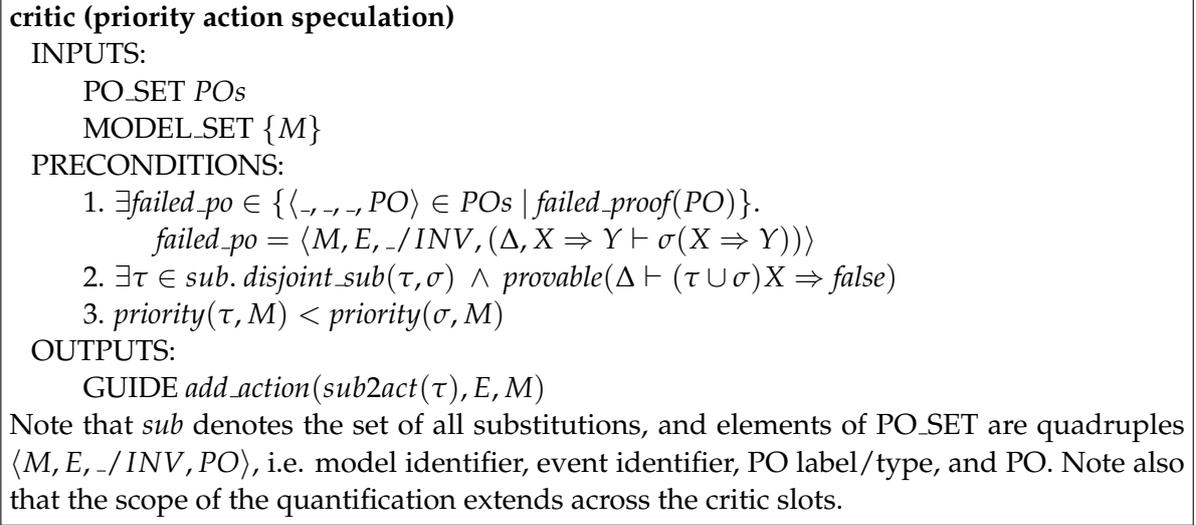


Figure 5: A reasoned modelling critic based on a priority heuristic to modify an action

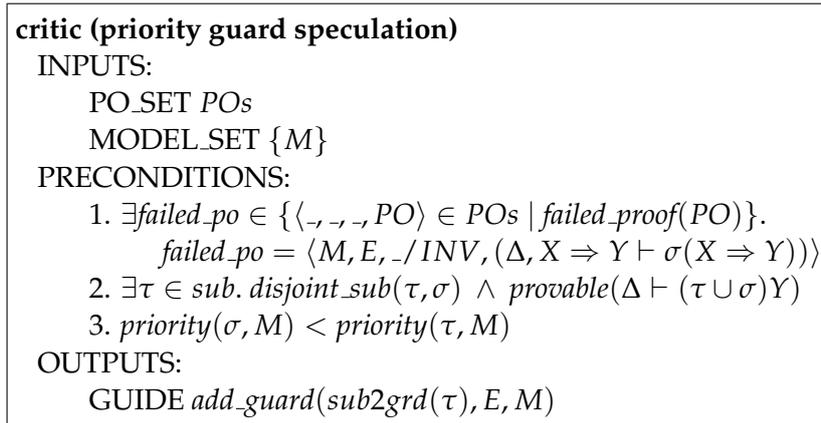


Figure 6: A reasoned modelling critic based on a priority heuristic to modify a guard

H2: If the priority of the candidate variable is higher than the priorities of all the variables updated by the event, then it is strongly suggestive that the change should be achieved via a new guard.

Where priorities are the same, no ranking of the alternatives will be provided. We represent heuristics H1 and H2 as critics in Figure 5 and Figure 6 respectively. The meta-logical terms that appear within the preconditions of these critics are defined in Figure 4. Note that meta-logical predicates, such as *priority*, require input from designers, as mentioned above. Both critics are applicable where an invariant proof has failed, and specifically where the invariant takes the form of an implication. We focus on solutions that involve the addition of either guards or actions, where guards are restricted to equalities. Later we discuss how these basic critics could be generalised.

The critic representing H1 (Figure 5) has three preconditions. The first precondition holds if there exists a PO of the required form that has failed to be proved, i.e. an invariant PO of the form $X \Rightarrow Y$, which is associated with the event E where substitution σ represents the actions corresponding to E . The second precondition holds if there exists a substitution (τ) which falsifies the antecedent, where τ and σ are disjoint. The third precondition expresses a modelling constraint, i.e. the variable(s) introduced by the new substitution have lower priority than the variable(s) updated by the existing actions. If the preconditions hold, then the guidance provided takes the form of a guard – constructed from τ .

The critic representing H2 (Figure 6) again has three preconditions, but illustrates a slightly different failure analysis pattern. That is, instead of making the antecedent false, the analysis tries to make the consequent true. Note that the third precondition ensures that the variable(s) introduced by the new substitution are of a higher priority than the variable(s) updated by the existing actions. Here the guidance provided takes the form of an action – constructed from τ .

3.2 The cruise control system in Event-B

```

MACHINE cruise_control
VARIABLES brake, cc
INVARIANTS inv1: cc = on  $\Rightarrow$  brake = off
METADATA pri_var(cc) < pri_var(brake)
EVENTS
  INITIALISATION  $\hat{=}$ 
    begin
      act1: brake := off
      act2: cc := off
    end
  EVENT enable_cc  $\hat{=}$ 
    begin
      act1: cc := on
    end
  EVENT pressbrake  $\hat{=}$ 
    begin
      act1: brake := on
    end
  ....
END

```

Note that since the machine name can be easily inferred, we write *pri_var*(cc) and *pri_var*(brake) in the **METADATA** field, instead of *pri_var*(cc,cruise_control) and *pri_var*(brake,cruise_control).

Figure 7: The cruise control system in Event-B.

To illustrate the application of critics introduced above, we will first formalise in Event-B the cruise-control system that was outlined in §1. The key fragments of the model are presented in Figure 7 and explained below.

Recall that the status of the brakes is represented by a variable *brake*, while the status of the cruise control is represented by the variable *cc*. Both variables are of type $Val = \{on, off\}$, where $on \neq off$. The variable *brake* is *on* if the user currently presses the brake (*pressbrake*), and *off* otherwise. While the variable *cc* is *on* if the cruise control is enabled (*enable_cc*), and *off* otherwise. The system invariant is represented as follows:

$$inv1 : cc = on \Rightarrow brake = off$$

and initially both variables are *off*:

$$\text{INITIALISATION} \hat{=} \mathbf{BEGIN} \textit{brake} := \textit{off} \parallel \textit{cc} := \textit{off} \mathbf{END}$$

We assume that the following priority relationship between the variables has been specified by the designer:

$$pri_var(cc, \textit{cruise-control}) < pri_var(\textit{brake}, \textit{cruise-control})$$

Failure and modelling suggestion 1

critic (priority guard speculation)

INPUT

PO_SET *pos*

MODEL_SET { *cruise-control* }

PRECONDITIONS:

1. $\langle \textit{cruise-control}, \textit{enable_cc}, \textit{inv1} / INV, \textit{po} \rangle \in \textit{pos} \wedge$
 $\textit{po} = (cc = on \Rightarrow brake = off \vdash \{cc \mapsto on\}(cc = on \Rightarrow brake = off)) \wedge$
 $\textit{failed_po}((cc = on \Rightarrow brake = off \vdash brake = off)$
2. $\textit{disjoint_sub}(\{brake \mapsto off\}, \{cc \mapsto on\}) \wedge$
 $\textit{provable}(cc = on \Rightarrow brake = off \vdash \{brake \mapsto off, cc \mapsto on\}(brake = off))$
3. $\textit{priority}(\{cc \mapsto on\}, \textit{cruise-control}) < \textit{priority}(\{brake \mapsto off\}, \textit{cruise-control})$

OUTPUTS:

GUIDE $\textit{add_guard}(brake = off, \textit{enable_cc}, \textit{cruise-control})$

Figure 8: An instantiation of the ‘priority guard speculation’ critic.

The model contains events that control the brakes and cruise control. The consistency proofs of invariant *inv1* over these events fails in two cases. Firstly, the event to enable the cruise control is defined as follows:

$$\mathbf{EVENT} \textit{enable_cc} \hat{=} \mathbf{BEGIN} \textit{cc} := \textit{on} \mathbf{END}$$

This gives rise to the following proof obligation:

$$cc = on \Rightarrow brake = off \vdash \{cc \mapsto on\}(cc = on \Rightarrow brake = off)$$

which can be reduced to the unprovable goal:

$$cc = on \Rightarrow brake = off \vdash brake = off.$$

However, the preconditions of the ‘priority guard speculation’ critic given in Figure 6 hold, and an instantiation of this critic for this example is shown in Figure 8 (after some simplification and unfolding). This instantiation suggests the addition of a new guard of the form $brake = off$. If the suggestion is accepted by the designer, then the updated event takes the form:

EVENT enable_cc $\hat{=}$ **WHEN** $brake = off$ **THEN** $cc := on$ **END**

Failure and modelling suggestion 2

<p>critic (priority action speculation)</p> <p>INPUT</p> <p>PO_SET pos</p> <p>MODEL_SET { $cruise-control$ }</p> <p>PRECONDITIONS:</p> <ol style="list-style-type: none"> 1. $\langle cruise-control, pressbrake, inv1/INV, po \rangle \in pos \wedge$ $po = (cc = on \Rightarrow brake = off \vdash \{brake \mapsto on\}(cc = on \Rightarrow brake = off)) \wedge$ $failed_po((cc = on, brake = off \vdash false))$ 2. $disjoint_sub(\{cc \mapsto off\}, \{brake \mapsto on\}) \wedge$ $provable(cc = on \Rightarrow brake = off \vdash \{brake \mapsto on, cc \mapsto off\}((cc = on) \Rightarrow false))$ 3. $priority(\{cc \mapsto off, cruise-control\}) < priority(\{brake \mapsto on\}, cruise-control)$ <p>OUTPUTS:</p> <p>GUIDE $add_action(cc := off, pressbrake, cruise-control)$</p>

Figure 9: An instantiation of the ‘priority action speculation’ critic.

The second failure with respect to $inv1$ occurs when a driver presses the brakes, as defined by the event:

EVENT pressbrake $\hat{=}$ **BEGIN** $brake := on$ **END**

This creates the following proof obligation:

$$cc = on \Rightarrow brake = off \vdash \{brake \mapsto on\}(cc = on \Rightarrow brake = off)$$

which, when simplified, becomes false:

$$cc = on, brake = off \vdash false.$$

In this case the ‘priority action speculation’ critic given in Figure 5 triggers, and Figure 9 shows the corresponding instantiation. As a result, the following updated *pressbrake* event is suggested to the designer:

$$\text{EVENT } \textit{pressbrake} \hat{=} \text{BEGIN } \textit{brake} := \textit{on} \parallel \textit{cc} := \textit{off} \text{END}$$

Generalisations of the critics

An alternative definition of the invariant is

$$\textit{inv2} : \textit{brake} = \textit{on} \Rightarrow \textit{cc} = \textit{off}$$

which is the contrapositive of *inv1*. To keep our presentation as simple as possible, our critics will not trigger on POs arising from *inv2*. However, this limitation could be overcome by generalising the failure analysis, i.e. allowing the critics to consider the alternatives of falsifying an antecedent and validating a consequent. This could be achieved by changing precondition 2 as follows:

$$\textit{provable}(\Delta \vdash (\tau \cup \sigma)X \Rightarrow \textit{false}) \vee \textit{provable}(\Delta \vdash (\tau \cup \sigma)Y)$$

Another limitation is that currently we only consider the addition of guards that take the form of equalities. The ‘priority guard speculation’ critic could be extended to capture a more general notion of a guard. However, this would require a new notion of priority, which is currently not defined for arbitrary terms (only variables and substitutions). Finally, global consistency issues have been ignored. When an action is modified, the validity of other invariants may change, or a deadlock may arise. The critics should be updated with such consistency checks. Depending on how global analysis is dealt with – as discussed below – the ‘priority action speculation’ critic of Figure 5 may have to also include such a *consistency* predicate.

3.3 From local to global modelling suggestions

As mentioned earlier, we envisage that our critics will support global analysis, as well as the local analysis illustrated via the cruise control example. For instance, when attempting to prove the internal coherence of a machine, local analysis may suggest the need to add guards across multiple events. Taking a more global perspective, these local suggestions may reveal a more general modelling suggestion. For instance, if a formula can be found that is logically weaker than all the suggested guards, then this points to a global rather than a local problem – and the need for a global solution, i.e. an invariant.

A concrete example of this kind of global analysis is found in Abrial’s ‘Cars on a Bridge’ example [1]*. In this example a single-laned bridge between an island and the mainland is modelled. Since it is single-laned, cars can only travel in a given direction at a given time – and this is controlled by two traffic-lights: *ml_tl* is the traffic light on the mainland; and *il_tl* is the traffic light on the island. The lights can have two colours – *green* and *red* – and it

*The example is modified slightly.

is assumed that $green \neq red$. Only a finite number of cars can be on the bridge and on the island, and three variables are used to count the cars (using sensors in later refinements): a is the number of cars on the bridge travelling towards the island; b is the number of cars on the island; c denotes the number of cars on the bridge leaving the island, while d is a constant denoting the total number of cars. Our discussion will focus on the four events where cars enter the bridge. $ML.out1$ and $ML.out2$ captures cars entering the bridge from the mainland, and is formalised as follows:

EVENT $ML.out1 \hat{=}$ WHEN $ml_tl = green$ $a + b + 1 < d$ THEN $a := a + 1$ END	EVENT $ML.out2 \hat{=}$ WHEN $ml_tl = green$ $a + b + 1 = d$ THEN $a := a + 1$ $ml_tl := red$ END
---	--

$IL.out1$ and $IL.out2$ captures cars entering the the bridge from the island, and is formalised as follows:

EVENT $IL.out1 \hat{=}$ WHEN $il_tl = green$ $b > 1$ THEN $b := b - 1$ $c := c + 1$ END	EVENT $IL.out2 \hat{=}$ WHEN $il_tl = green$ $b = 1$ THEN $b := b - 1$ $c := c + 1$ $il_tl := red$ END
---	--

Without giving all the details[†], the proofs of

$$il_tl = green \Rightarrow a = 0$$

fail over events $ML.out1$ and $ML.out2$, while the proofs of invariant

$$ml_tl = green \Rightarrow c = 0$$

fail over events $IL.out1$ and $IL.out2$. The corresponding POs have the following form:

$$\dots, il_tl = green \vdash \sigma_1(ml_tl = green \Rightarrow \dots) \text{ where } ml_tl \notin dom(\sigma_1) \quad (5)$$

$$\dots, ml_tl = green \vdash \sigma_2(il_tl = green \Rightarrow \dots) \text{ where } il_tl \notin dom(\sigma_2) \quad (6)$$

A local analysis of each of these failures suggests making the antecedent false, by finding a substitution τ such that:

$$disjoint_sub(\tau, \sigma_{\{1,2\}}) \wedge provable(\Delta \vdash (\tau \cup \sigma_{\{1,2\}})X \Rightarrow false)$$

In (5), τ will be instantiated to $\{ml_tl \mapsto red\}$ while in (6), τ will be instantiated to $\{il_tl \mapsto red\}$. However, there is a common solution to these failures. Firstly, we assume that the substitutions are turned into guards (by *sub2grd*). Now, if an event is constrained by the guards G_1, \dots, G_n , then adding an additional guard G_{n+1} to patch a failure is equivalent to adding the weaker guard $G_1 \wedge \dots \wedge G_n \Rightarrow G_{n+1}$ (via modus ponens). Thus the suggested

[†]These details can be found in [1] and <http://www.event-b.org>.

additional guards can be “weakened” in this way with the existing guards of the events. For this particular example, we ignore all guards except those shown in (5,6) – the other guards are irrelevant for this failure. For example, from (5) the guard $ml_tl = red$ is derived and “weakened” to give $il_tl = green \Rightarrow ml_tl = red$. We then try to find a common solution for (5,6), i.e. an r such that:

$$\begin{aligned} r &\Rightarrow il_tl = green \Rightarrow ml_tl = red \\ r &\Rightarrow ml_tl = green \Rightarrow il_tl = red. \end{aligned}$$

Moreover, we require that r is satisfiable, i.e. the trivial solution $r = false$ is not considered. The only valid such generalisation r (modulo equivalences) is:

$$il_tl = red \vee ml_tl = red$$

which is suggested as an invariant. This represents a key safety requirement – at all times at least one traffic-light is red.

critic (multiple failure invariant speculation)

INPUTS:

PO_SET POs

MODEL_SET $\{M\}$

PRECONDITIONS:

1. $\exists h \subseteq hs. h = \{ \langle E, G \rangle \mid \exists po \in pos. \langle _, E, _ / INV, po \rangle \in POs \wedge failed_proof(po) \wedge po = (\Delta, X \Rightarrow Y \vdash \sigma(X \Rightarrow Y)) \wedge \exists \tau \in sub. disjoint_sub(\tau, \sigma) \wedge provable(\Delta \vdash (\tau \cup \sigma)X \Rightarrow false) \wedge G = guards(E) \Rightarrow sub2grd(\tau) \}$
2. $generalisable(h) \wedge \exists i \in forms. i = generalise(h)$

OUTPUTS:

GUIDE $add_invariant(i, M)$

Note that hs denotes the set of all event-guard pairs, while pos and $forms$ denote the sets of all proof obligations (sequents) and formulae respectively.

Precondition 1 holds if a local analysis suggests that proof failures can be overcome by the introduction of additional guards.

Precondition 2 holds if the suggested guards can be generalised to give an invariant.

Figure 10: A global reasoned modelling critic suggesting an invariant from multiple failures

With regards to realising such global reasoning, we are considering two distinct alternatives via our proposed critic schema:

1. Combine both local and global analysis of the available POs within a single critic, as illustrated in Figure 10.
2. Firstly use critics to perform local analysis of the available POs, recording any modelling suggestions. Secondly, use separate critics to perform a global analysis, exploiting the results of the local analysis.

From a scientific point of view, the approaches are not very different – and should be seen more as an implementation issue. Experimentation will be required in order to determine the most practical approach.

4 Implementation

As with most theories, the notion of reasoned modelling will only be fruitful with proper tool support. Here, we will describe the REMO[‡] tool, which implements a prototype for the reasoned modelling critic mechanism introduced above. Since the technique proposed here has been illustrated for Event-B, the tool is necessarily integrated within the Eclipse-based Rodin toolset — implemented in Java. However, as is the case with the ProB Rodin plug-in [15], the main part of the program has been implemented in another programming language. In our case, we have chosen OCaml (*Objective Caml Language*)[§]. This has been motivated by the the following factors:

- this is a prototype implementation which we expect will change (frequently). OCaml is a declarative functional programming language which reduces the size of the source code, thus making change to the code (relatively) easy;
- a large part of the work is checking precondition, which can be handled very elegantly using inductive data types and pattern matching, which is supported by OCaml;
- compared to other languages in the ML-family, OCaml embodies an object-oriented layer, which provides a good encapsulation mechanism to represent models, and particularly critics;
- OCaml is used in the CORE project[¶], which is a sister project of ours. The use of OCaml enables us to take advantage of the reasoning techniques developed within the CORE project. In particular, term synthesis and integration with the CVC3 SMT solver, discussed below.

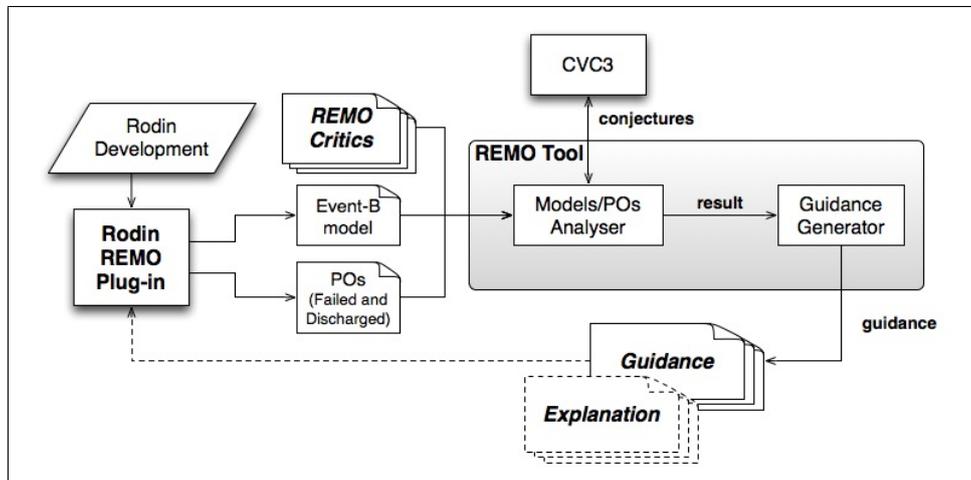


Figure 11: Tool architecture

The overall process embodied within REMO involves an analysis of the applicability of the critics with respect to Event-B models and their failed proof obligations. The results of the analysis are presented to the user in the form of guidance – it is the user that ultimately

[‡]The REMO acronym follows from REasoned MOdelling

[§]See <http://caml.inria.fr/ocaml/>

[¶]See <http://www.macs.hw.ac.uk/core/>

decides whether or not to accept the guidance. An overview of the architecture of the tool can be found in Figure 11. The labels with **bold** font indicate the main components of the tool chain; the stippled lines indicate work in progress; the square boxes  represent processes; the document boxes  represent files that are used as inputs and/or outputs of a process; while the rectangle box  represents input data used in a process.

4.1 Rodin REMO Plug-in

The goal of the Rodin REMO plug-in is to provide an interface between the Rodin toolset and our REMO tool. Specifically the plug-in has two main roles:

1. translate Event-B developments (models, contexts and proof obligations) from Rodin into the REMO tool; and
2. communicate modelling guidance and explanation back from the REMO tool into Rodin.

However, at the current state only 1 has been implemented – 2 is in our plan of future work, and communication is currently at the level of printed messages on a terminal window. However, a description of how we envisage the guidance will be presented to the user by the Rodin plug-in is given throughout this section.

The input of the Rodin REMO plug-in is a Rodin development, from which the Plug-in generates two files that serve as the input of the REMO tool. The first file contains a string representation of the Event-B models (machines and contexts) of a development, while the second file stores a string representation of all the POs (status, type, hypotheses and goal) associated to each model. These files are then translated into the REMO internal representation, i.e. OCaml objects.

Currently, the PO generator of Rodin applies some simplification before the POs are generated. This simplification consists of replacing the substitutions given by the events to the variables in the POs. As a consequence, the file generated by our plug-in contains the POs after the simplification has taken place. However, in the analysis of the critics the pure PO (i.e. before the substitutions are applied) is also needed in order to do the matching between the required POs (given by the preconditions) and the ones of the models. Currently, this information is hand-crafted in the file; however, we plan to update the Rodin PO generator to overcome this problem.

4.2 REMO Critics

The reasoned modelling critics represent another input to the REMO tool. Each critic is implemented as a class in OCaml that inherits from the virtual class *reasonedModellingCritic*. This super class encapsulates the basic elements of a critic as shown in Figure 12.

In particular, each critic overrides the methods *checkPreconditions* and *generateGuidance*. The first method evaluates the applicability of the critic through the checking of the preconditions while the second method analyses the output of this check and produces the guidance that will be suggested to the user. This use of inheritance allows new critics to be added relatively easily to the REMO tool.

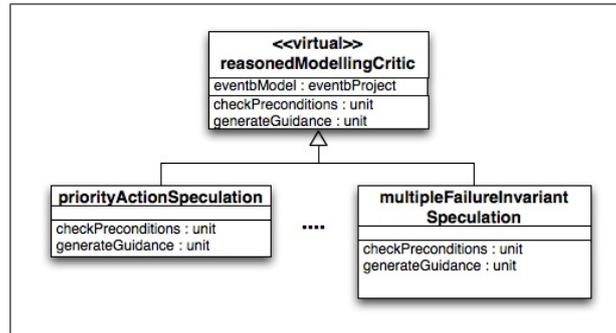


Figure 12: Critics implementation

4.3 The REMO Tool

The REMO tool is at the heart of over implementation. It contains two components, *Models/POs Analyser* and *Guidance Generator*, which are applied sequentially and discussed separately below.

Models/POs Analyser

The inputs of the Model/POs Analyser are the models, the POs and the critics. The first task of the process is to parse the model and PO files generated by the Rodin REMO plugin. Then the analyser evaluates the Event-B models and their POs in order to find which critics are applicable (if any). Note that we assume that all unproven POs are in fact unprovable. Obviously, since the underlying logic of Event-B is undecidable, this is not a valid assumption for all models – and we plan to restrict this in future work, e.g. by incorporating a proof planner (see §6).

The process of evaluating the applicability of a critic with respect to a model consists of the following steps (this process is further illustrated with the cruise control example in §5):

1. First, the analyser searches for matches between each of the preconditions of the critic and the models. If more than one match is found, two things can happen. If the precondition being evaluated is one that identifies different failures in the model (e.g. failed POs), each match would be considered as a separate instantiation of the critic. However, if the precondition being checked is one that identifies different solutions to a same failure (e.g. finding a substitution that is disjoint with the other substitutions of an event), all the matches will be part of the set of all possible solutions of a failure.
2. Second, the subsequent preconditions are evaluated for each element of the set of possible solutions of each instance of a critic. When all the preconditions have been checked, the analyser verifies which instances were successful in all the preconditions.
3. Finally, the analyser collects the results of each successful instance and sends them to the guidance generator.

The REMO tool uses the CVC3 SMT solver [5] to check if a match is found for a specific precondition. This choice of SMT solver follows from an existing OCaml interface developed within the CORE project. Currently, we only support first-order predicates with

equality, as well as arithmetic. Thus, the full underlying (set-theoretic) Event-B mathematical notation is not supported. In the future, we plan to take advantage of both work with SMT solvers for Event-B^{||}, as well the use of a proof planner (see above), and possibly a generic purpose theorem provers.

Guidance Generator

The guidance generator takes the raw results from the Analyser and produces a formatted, ranked (ordered) list of alternative guidance suggestions, which will then be sent to the Rodin plug-in and presented to the user.

Two processes take place here. The first process consists of interpreting the guidance given as an input in order to produce written explanations. These are linked to the critic preconditions of each suggestion. The purpose behind this, is to help the user understanding the nature of the guidance, to (hopefully) make a more informed decision about which modification to apply. We have classified the guidance produced by the generator into three categories:

1. **Global guidance:** This refers to a global solution that addresses multiple local failures (e.g. the addition of an invariant to the model instead of guards in multiple events). This type of guidance is given alongside an explanation of how the local suggestions are used to propose the general one.
2. **Local guidance:** This is concerned with a solution that addresses a local failure (e.g. a guard is missing from an event). An explanation of how the failure is overcome with the solution is provided.
3. **Conditional guidance:** This type of guidance is related to solutions that are subject to the fulfilment of certain conditions (e.g. priorities of variables that were not defined in the model). A list of conditions together with explanations of why these conditions are needed is given with the suggested solution.

The second process, which is currently under development, deals with the linking of the guidance and explanation to the Rodin plug-in.

Guidance & Explanation

The result of the guidance generator is a set of files capturing both the guidance and explanations for each critic instance. These files will then be sent to the Rodin plug-in for presentation to the user.

From the user perspective, the guidance will consist of the fragment of Event-B model that is relevant to the suggested modification. For instance, in the case of the priority action speculation critic, the guidance will contain the current event plus the new action as shown in Figure 13. The instantiation of this schema for our running cruise control example is shown in Figure 16.

The explanations will be represented as templates, i.e. a natural language explanation of the critic preconditions, with slots to denote instance specific details. To illustrate, the

^{||}See <http://www.cprover.org/SMT-LIB-LSM/>

<p>Guidance (priority action speculation critic)</p> <p>EVENT <i>event_name</i> WHERE <i>guards</i> THEN <i>old_actions</i> <div style="border: 1px solid black; display: inline-block; padding: 2px;"><i>new_action</i></div> END</p>

Figure 13: Guidance provided by the priority action speculation critic analysis.

proposed template associated to the priority action speculation critic is shown in Figure 14, with an instantiation of this template is given in Figure 16.

<p>Explanation template (priority action speculation critic)</p> <ol style="list-style-type: none"> 1. There exists an unproven invariant PO _____. 2. There exists a variable update _____ that does not interfere with the existing actions _____ of event _____ and which makes the PO provable. 3. Update preserves the variables priority relationship _____. <p>Therefore: Failure can be overcome by adding action _____ to event _____.</p>
--

Figure 14: Explanation template associated to the priority action speculation critic.

Regarding conditional guidance, consider again precondition 3 of the priority action speculation critic (see Figure 5). Whether or not this precondition holds depends upon if the user has provided the relevant meta-data, i.e. priority information with respect to their model. If a critic fails to apply because of the lack of such meta-data, the user can still be presented with the relevant guidance, but conditional with respect to the missing meta-data.

5 Results - Examples Revisited

Here, we apply the ‘cruise control’ and ‘cars on a bridge’ examples to the REMO tool, and demonstrate step by step the analysis process. These examples are discussed separately below.

5.1 The cruise-control example – the (local) priority critics

The initial step, is the evaluation of the preconditions. The analysis of the critic’s preconditions involves two parts, a syntactic check followed by a semantic check. The syntactic checking is a filtering process that consists of searching for elements in the models and the

POs that have specific structures. After this filtering is done, the semantic checking consists of the analysis of specific properties that these elements must satisfy in order to be an instantiation of the precondition. Both of these checks rely heavily on pattern matching, thus illustrating the advantages of functional languages such as OCaml.

To demonstrate, we analyse the priority action speculation critic in the ‘cruise control’ model. Precondition 1 (PC_1) of this critic states:

$$\mathbf{PC}_1 : \exists \text{failed_po} \in \{ \langle _, _, _, _ \rangle \in \text{POs} \mid \text{failed_proof}(\text{PO}) \}. \\ \text{failed_po} = \langle M, E, _ / \text{INV}, (\Delta, X \Rightarrow Y \vdash \sigma(X \Rightarrow Y)) \rangle$$

We need to search for all failed POs, whose type is *INV* (i.e. invariant preservation POs) with the form $X \Rightarrow Y$ (since invariant proofs have an inductive structure, this form will be found in both the goal and the hypothesis). The POs are represented as objects in the tool – where they are stored with their attributes, status and type. Using this meta-information we can directly check if the type is *INV*. The next part checks that the form of the goal/hypothesis is (syntactically) $X \Rightarrow Y$. This is achieved by pattern matching. Two (failed) POs succeed this check:

$$\mathbf{po}_1 : \dots \wedge cc = on \Rightarrow brake = off \vdash \{ cc \mapsto on \} cc = on \Rightarrow brake = off \\ \mathbf{po}_2 : \dots \wedge cc = on \Rightarrow brake = off \vdash \{ brake \mapsto on \} cc = on \Rightarrow brake = off$$

Each of these POs represents a different failure, as a result, each of them will be treated as an instance of the critic. Let $instance_1$ be the instance of the critic associated to po_1 , and $instance_2$ be the instance associated to po_2 . In the following precondition we look for a substitution to falsify the antecedent, and the precondition is formalised as follows:

$$\mathbf{PC}_2 : \exists \tau \in \text{sub. disjoint_sub}(\tau, \sigma) \wedge \text{provable}(\Delta \vdash (\tau \cup \sigma)X \Rightarrow \text{false})$$

The checking of this precondition is broken into three steps:

- $\mathbf{PC}_{2.1}$: Find the possible substitutions that can be applied to the antecedent of the goal.
- $\mathbf{PC}_{2.2}$: From $\mathbf{PC}_{2.1}$ find the substitutions that are disjoint with the current substitutions of the event.
- $\mathbf{PC}_{2.3}$: From $\mathbf{PC}_{2.2}$ find the substitutions that are provable.

The antecedent of the goals of po_1 and po_2 is $cc = on$; thus, in the analysis of $\mathbf{PC}_{2.1}$ we need to find all possible values that the variable cc can take. In order to do this, we need to identify the type of the variable, which is given by the invariant $cc \in \text{state}$, and then we look for the members of the set *state*. This is analysed by looking for a ‘‘type axiom’’ of the form ‘ $\text{state} = \{ _, _, \dots, _ \}$ ’, or in the case this is not given, by axioms of the form ‘ $_ \in \text{state}$ ’. From this, in the cruise control model, *on* and *off* are defined as the members of the set *state*. Therefore, two possible substitutions are found for each instance of the critic: $cc \mapsto on$ and $cc \mapsto off$. These are considered as two potential solutions for each failed PO.

The next step, $\mathbf{PC}_{2.2}$, is the verification of the disjointedness of the substitutions found in $\mathbf{PC}_{2.1}$. We find that for $instance_1$, the suggested substitutions are not disjoint with $cc \mapsto on$, which is the substitution of event *enable_cc*. As a result, both solutions are considered as failures of this instance of the critic. On the contrary, in $instance_2$, the suggested substitutions are both disjoint with $brake \mapsto off$, which is the substitution of event *pressbrake*.

Finally, in $\mathbf{PC}_{2.3}$ we need to verify if the statement $\Delta \vdash (\tau \cup \sigma)X \Rightarrow \text{false}$ is provable for the substitutions found in $\mathbf{PC}_{2.2}$. The instantiation of this statement generates the following conjectures:

conj₁: $cc = on \Rightarrow brake = off \vdash \{cc \mapsto on, brake \mapsto on\} cc = on \Rightarrow false$

conj₂: $cc = on \Rightarrow brake = off \vdash \{cc \mapsto off, brake \mapsto on\} cc = on \Rightarrow false$

Both conjectures are validated through CVC3. The result is that conj₂ is valid but conj₁ is not. This rules out $cc \mapsto on$ as a possible solution of the critic. Thus, up to this point, we have *instance₂* as the only instance of the critic, and $cc \mapsto off$ as the only potential solution for that instance.

Finally, the last precondition verifies the priorities. The precondition is:

$$\mathbf{PC}_3: priority(\sigma, M) < priority(\tau, M)$$

In this case, the priority of the new substitution should be less than the priority of the current substitutions (actions) of the event. The priority of a substitution is the priority of the variable being modified; therefore, in order to perform the analysis of the precondition, we must compare the priorities of the variables. Event *pressbrake* has only one substitution $brake \mapsto off$; then, we need to check that

$$pri_var(cc, cruise-control) < pri_var(brake, cruise-control)$$

. This is evaluated by the REMO tool through pattern matching of the meta-data defined in the model. The precondition succeeds, creating as guidance the addition of the action $cc := off$ in event *pressbrake*.

Figure 15 sums up the steps taken by the tool for the analysis of the priority action speculation critic. In the figure, we present the evaluation of the two instances of the critic for the cruise control system. For each instance, the branches represent the different alternatives to solve the failure while each node represents the result (or matches) of the analysis of each precondition with respect to the specific alternative. Each branch terminates either when a failure of a precondition is found or when all the preconditions have succeeded in that branch.

Finally, Figure 16 shows the guidance that would be provided by the tool. The analysis of the priority guard speculation critic is performed in a similar fashion.

5.2 The cars on a bridge example – the (global) multiple failure critic

The multiple failure invariant speculation critic is implemented by combining both local and global analysis within a single critic. This methodology has been chosen since it facilitated the collection of the results of the local analysis, as well as avoiding duplication of information since all the processes are achieved within the same critic, instead of being distributed over many critics.

In the ‘cars on a bridge’ example, the analysis of the first precondition involves a very similar process to the process described for the priority action speculation critic. Thus, we will focus on how the global analysis was implemented – i.e. the second precondition of Figure 10. This precondition is used to generalise the solutions found in the first precondition into an invariant which ideally captures all the failures – however, we only require it to capture more than 1 failure over more than 1 event. The precondition is:

$$\mathbf{PC}_2: generalisable(h) \wedge \exists i \in forms. i = generalise(h)$$

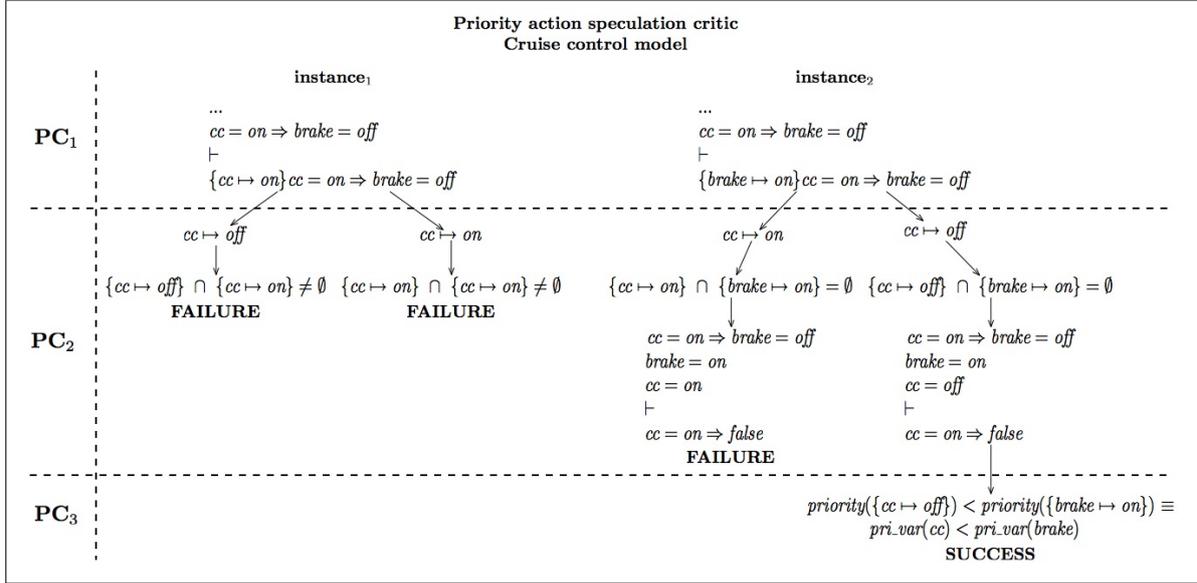


Figure 15: Sequence of steps performed by REMO in the analysis of the priority action speculation critic for the cruise control example.

with the following definitions:

$$\begin{aligned}
 sat(P) &\hat{=} \text{“The predicate } P \text{ is satisfiable”} \\
 generalisable(H) &\hat{=} \exists r, e_1, e_2, g_1, g_2. \langle e_1, g_1 \rangle \in H \wedge \langle e_2, g_2 \rangle \in H \\
 &\quad \wedge e_1 \neq e_2 \wedge r \Rightarrow g_1 \wedge r \Rightarrow g_2 \wedge sat(r) \\
 generalise(H) &\hat{=} \text{“The weakest } r \text{ such that } sat(r) \text{ and for the} \\
 &\quad \text{most } \langle e, g \rangle \in H, r \Rightarrow p \text{ (for at least 2 events)”}
 \end{aligned}$$

The definition of *generalise* is (deliberately) under-defined, and finding good implementation of it is a hard problem. Below, we have implemented a rather specialised version of it, but in the future we plan to rely on *abduction* [17] (see §2.2), possibly implement using *term-synthesis* [16], which using forward chaining, generates formulas which are then checked. An efficient implementation of term-synthesis would require finding good heuristics to avoid generating too many uninteresting formulas.

However, before that we will expand the ideas behind precondition 1 and 2 of this critic. In the first step we first find a guard which will cause the PO to be proven. We then weaken the solution as much as we can with the existing guards of the event (it can be even further weakened by existing invariants – but this would have complicated the presentation too much). For example, a solution G for an event with guards G_1 and G_2 becomes $G_1 \wedge G_2 \Rightarrow G$. In the second precondition, we then try to find the weakest strengthening r which is as least as strong as two of the weakened solutions (of two different events).

A straightforward such strengthening is to remove one or more of the guards, e.g. creating $G_2 \Rightarrow G$, and this approach is followed here, where the use of common (free) variables in the goal is used to guide the elimination of guards. From this, we divide the analysis for this precondition into four steps:

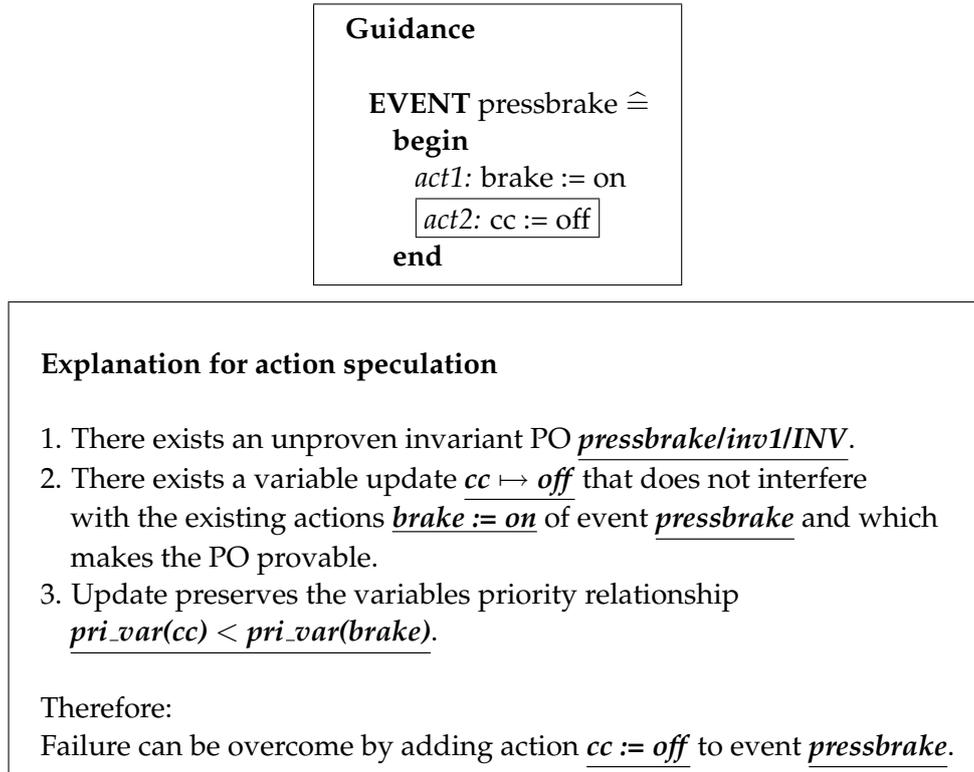


Figure 16: Proposed guidance and explanation provided by REMO for the cruise control system.

1. First, the set of free variables of the goals of the events (FV) is found.
2. Then all conjuncts (each representing a guard of the event) that use free variables not in FV are omitted.
3. We then have four conjectures – however two and two conjectures are (syntactically) equivalent, given the following two (different) conjectures (resulting from POs of different events):

$$il_tl = green \Rightarrow ml_tl = red \quad \text{and} \quad ml_tl = green \Rightarrow il_tl = red$$

4. Finally, we check for equivalence between the resulting conjectures. If two conjectures (from different events) are equivalent then one of them are picked. The (equivalent) solutions are then ranked according to which capture most different events – followed by the number of conjectures it captures (in a lexicographical order). This highest ranked solution will then form the weakest such generalisation, and is returned. These steps can therefore be seen as an implementation of *generalise*. In this case, they are all equivalent, and can be rewritten to

$$ml_tl = red \vee il_tl = red.$$

This is then presented as a global alternative to address the multiple failures.

6 Related and future work

Our notion of *reasoned modelling* represents a new paradigm for exploring the interplay between reasoning and modelling. As evident from the above discussion, our general ideas on reasoned modelling are strongly influenced by the proof planning paradigm [7], while the particular work discussed here follows from the notion of *proof critics* [12]. Currently, Event-B users have to manually analyse failed proof attempts and patch their models, e.g. [1, 10]. This form of user interaction represents a significant barrier to the accessibility of the Event-B toolset.

Previously, the application of *design patterns* [11] has been suggested for Event-B [4]. The ability to reuse design patterns, and their associated proofs, has obvious benefits over the conventional notion of design patterns. Our work with failure-analysis is more closely aligned with *anti-patterns* [6], i.e. common patterns of bad design, coupled with solutions. Bad design, however, does not necessarily mean incorrect design – i.e. the POs may still be provable. For this reason the failure-driven nature of the critics presented here differ from the conventional notion of an anti-pattern. Another related area is ontology repair plans [9] – proof failure is explored using proof planning techniques to repair ontologies – i.e. proof planning is used to describe evolving models.

In terms of future work, our short-term aim is to complete and fully integrate the REMO tool within Rodin. Moreover, we aim to investigate modelling heuristics in addition to variable priorities. While we have relied upon existing external provers, in the longer term we aim to incorporate a proof planner. This will enable us to also develop critics that analyse successful proofs and provide guidance along the lines of conventional anti-patterns.

7 Conclusion

We have motivated the need to abstract away from the complexities of proof obligations and proof-failure analysis. To address this need, we have proposed a technique of proof management that attempts to turn proof-failure analysis into modelling guidance via *reasoned modelling critics*. In doing so we aim to increase the productivity of designers as well as the accessibility of formal modelling. Our technique is motivated by proof planning and, in particular, proof critics. The work differs from previous work in that it provides a uniform framework in which proof-failure heuristics can be combined with modelling heuristics. It is this combined approach that will enable us to abstract away from the complexity of proof obligations. While our initial critics have focused on variable priorities, our framework will be extensible, allowing designers to record meta-data relating to other kinds of modelling decisions. Our REMO prototype represents a proof of concept, and the first step towards a reasoned modelling plug-in for the Rodin toolset.

Acknowledgements

This research is supported by EPSRC grants EP/F037058 and EP/E005713. Gudmund Grov is supported by EPSRC grant EP/H024204/1 “AI4FM: the use of AI to automate proof search in Formal Methods” (www.ai4fm.org). Butler’s involvement is part of the EU research project ICT 214158 DEPLOY (www.deploy-project.eu). Our thanks go to

Alan Bundy, who motivated the use of variable priorities. Thanks also go to Jean-Raymond Abrial, Cliff Jones and Ewen Maclean for their feedback and encouragement with this work. Finally we thank the three anonymous ABZ 2010 reviewers for their constructive feedback.

References

- [1] J.-R. Abrial. *Modelling in Event-B: System and Software Engineering*. Cambridge University Press, 2009. To be published.
- [2] J.-R. Abrial, M. Butler, S. Hallerstede, T. S. Hoang, F. Mehta, and L. Voisin. Rodin: An open toolset for modelling and reasoning in event-b. *International Journal on Software Tools for Technology Transfer (STTT)*, April 2010.
- [3] J.-R. Abrial, M. Butler, S. Hallerstede, and L. Voisin. An Open Extensible Tool Environment for Event-B. In *ICFEM'06*, volume 4260 of *LNCS*, pages 588–605. Springer, 2006.
- [4] J.-R. Abrial and T. S. Hoang. Using Design Patterns in Formal Methods: An Event-B Approach. In *ICTAC'08*, volume 5160 of *LNCS*, pages 1–2. Springer, 2008.
- [5] C. Barrett and C. Tinelli. CVC3. In *CAV '07*, volume 4590 of *LNCS*, pages 298–302. Springer, 2007.
- [6] W. Brown, R. Malveau, H. W. S. McCormick, and T. Mowbray. *AntiPatterns: Refactoring Software, Architectures, and Projects in Crisis*. John Wiley & Sons, 1998.
- [7] A. Bundy. The Use of Explicit Plans to Guide Inductive Proofs. In *CADE'88*, volume 310 of *LNCS*, pages 111–120. Springer, 1988.
- [8] A. Bundy, D. Basin, D. Hutter, and A. Ireland. *Rippling: Meta-level Guidance for Mathematical Reasoning*. Cambridge University Press, 2005.
- [9] A. Bundy and M. Chan. Towards ontology evolution in physics. In *WoLLIC*, volume 5110 of *LNCS*, pages 98–110. Springer, 2008.
- [10] M. Butler and D. Yadav. An Incremental Development of the Mondex System in Event-B. *Formal Aspects of Computing*, 20(1):61–77, 2008.
- [11] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1995.
- [12] A. Ireland. The Use of Planning Critics in Mechanizing Inductive Proofs. In *LPAR'92*, number 624 in *LNCS*, pages 178–189. Springer, 1992.
- [13] A. Ireland, G. Grov, and M. Butler. Reasoned Modelling Critics: Turning Failed Proofs into Modelling Guidance. In *ABZ'10*, number 5977 in *LNCS*, pages 189–202. Springer, 2010.
- [14] A. Ireland and J. Stark. Proof Planning for Strategy Development. *Annals of Mathematics and Artificial Intelligence*, 29(1-4):65–97, Feb. 2001.
- [15] M. Leuschel and M. Butler. ProB: an Automated Analysis Toolset for the B Method. *Journal Software Tools for Technology Transfer*, 10(2):185–203, March 2008.
- [16] E. Maclean, A. Ireland, R. Atkey, and L. Dixon. Refinement and term synthesis in loop invariant generation. In *WING'09*, pages 72–86, 2009.
- [17] R. Monroy, A. Bundy, and A. Ireland. Proof Plans for the Correction of False Conjectures. In *LPAR'94*, number 822 in *LNAI*, pages 54–68. Springer, 1994.