



Report

Developing control systems with some fragile environment

Author(s):

Hoang, Thai Son; Hudon, Simon

Publication Date:

2010

Permanent Link:

<https://doi.org/10.3929/ethz-a-006906008> →

Rights / License:

[In Copyright - Non-Commercial Use Permitted](#) →

This page was generated automatically upon download from the [ETH Zurich Research Collection](#). For more information please consult the [Terms of use](#).

Developing Control Systems with Some Fragile Environment^{*}

Thai Son Hoang and Simon Hudon

Department of Computer Science,
Swiss Federal Institute of Technology Zurich (ETH-Zurich), Switzerland
htson@inf.ethz.ch and shudon@student.ethz.ch

Abstract. Event-B is a formal method that allows one to model various kinds of systems including control systems working within some fragile environment. However, it is lacking a systematic approach for developing this type of systems and it hinders the applicability of Event-B. Our contribution is such an approach and it is presented in this paper. Our proposed method focuses on a set of elements that should be captured by the formal model and prescribes an order in which they should be introduced. The key aspect of our approach is to first model the required behaviour of the environment, and then to introduce the controller to appropriately influence the environment. It has the advantage that every step of the such a development is dictated by the information available so far, including the requirements. We argue that having a clear development strategy early in the design process will assist the developers in producing high-quality models of the future software systems.

Keywords: Event-B, formal modelling, refinement, development strategy, system development.

1 Introduction

Event-B [1] is a modelling method for discrete transition systems which are correct-by-construction. Its applications range from sequential programs, concurrent programs to distributed systems. In particular, Event-B is one of the few modelling methods having control systems within its scope. More importantly, the development of such systems in Event-B includes the model of the environment which is a necessity for the assurance about the correctness of the future systems.

As a result, developing systems in Event-B is a complex task involving the management of several aspects of the systems, including the environment. A central aspect of Event-B is the use of step-wise refinement to reduce the complexity of system modelling. Abrial suggested in [1] that in practice, before engaging in the actual modelling task, developers should design a *refinement strategy* specifying for each refinement step which details will be introduced into the model. However, coming up with a good and helpful refinement strategy, which aids the system development, is a challenging task. Guidelines are needed in order to design such a refinement strategy.

^{*} Part of this research was carried out within the European Commission ICT project 214158 DEPLOY (<http://www.deploy-project.eu/index.html>).

For developing control systems, Butler has proposed a modelling guideline in what is known as the *cookbook* [2]. An application of the cookbook for developing a cruise control system is reported in [6].

In the present paper, we propose our *development strategy* which differs from that of the cookbook in some key aspects (see Sect. 5). We start in Sect. 2 by offering a summary of the Event-B notation; in Sect. 3, we explain our strategy and apply it to a control problem in Sect. 4. Finally, we discuss our results in Sect. 5.

2 The Event-B Modelling Method

Event-B is supported by a specialized notation for abstract machines, the central object of the development method. It supports both the formulation of formal specifications and their refinement. We give a brief overview of some essential aspects of Event-B in this Section. For a full details of Event-B, we refer our readers to [1].

Specification. In the Event-B notation, a machine is characterized by its state space modelled by some variables v and its transitions modelled by some events. The state variables v are constrained by some invariant $I(v)$. An event evt has the following form: $\text{evt} \hat{=} \mathbf{any } p \mathbf{ where } G(p, v) \mathbf{ then } S(p, v, v') \mathbf{ end}$, where p is the parameters, $G(p, v)$ specifies the enabled condition, and $S(p, v, v')$ is the action. A dedicated event without parameters and guards is used as the initialisation.

Action $S(p, v, v')$ contains several assignments that are supposed to happen simultaneously. Each assignment can take one of the three forms: $v := E(p, v)$, $v \in E(p, v)$, or $v \mid P(p, v, v')$. While the first form deterministically assigns the value of expression $E(p, v)$ to v , the second form non-deterministically assigns to v some value from $E(p, v)$. The last assignment form is the most general. It assigns to v some value satisfying the before-after predicate $P(p, v, v')$.

A machine is consistent if its invariant hold at any given time. In practice, this is guaranteed by proving that the invariant is established by the initialisation and maintained by all its events.

Refinement. Refinement is a well-known technique for reducing the complexity of developing formal models. One starts with an abstract machine capturing some central aspect of the system, and subsequently refines the machine by adding more concrete details to the model. When refining a machine, it is possible to introduce new variables and new events.

Consistency has to be proved between a concrete machine and its abstract machine. In practice, this is done on a per event basis. An event of the concrete machine is a refinement of an abstract event if the guard is *strengthened* and the action of the concrete event can be *simulated* by the action of the abstract event.

3 Development Strategy

Despite being a powerful modelling method, Event-B lacks a systematic approach for developing different types of systems. We suggest here some guidelines, which we call a *development strategy*, for developing systems with a fragile environment, i.e. systems

consisting of an environment and a controller. The environment and the controller communicate in a bi-directional fashion: the controller receives input from the environment via various *sensors*; reciprocally, the controller produces output to change the environment via various *actuators*.

Our development strategy contains four different stages. Note that each stage can be developed through several refinements.

Stage 1 Models the environments as it should behave.

Stage 2 Models the actuators to command the changes in the environment.

Stage 3 Models the sensors together with the controller.

Stage 4 Models some appropriate scheduler for the controller.

Stage 1 aims at describing an environment with the desirable properties, based on the requirements document. At this stage, we omit the controller completely, focus on global safety properties and how the physical components should work together to achieve such properties.

In Stage 2, actuators are introduced as means by which the controller will affect the environment, such that the physical components interact correctly with each other. This, in turns, puts some constraints on how the actuators can be set.

Up until Stage 2, all the control of the environment via actuators is done with perfect information of the whole system. Since this is unrealistic, Stage 3 aims at interposing sensors between observed components and the controller. This enforces an appropriate specification for the controller.

At Stage 4, we can introduce a scheduling strategy to the controller. The purpose is to optimise its efficiency.

Advantages of the Approach.

- Desirable safety properties of the systems are modelled earlier in the development in Stage 1. We can rely on refinement for the preservation of these properties during the development.
- Moreover, these properties serve as guideline for the design of the actuators (Stage 2) and, in turn, for the design of the sensors and the controller (Stage 3). In our opinion, this leads us more directly to a correct design since the controller is introduced as a *solution* to the *problem* of maintaining safety in the system.
- By deciding to introduce scheduling at the end of the development, we facilitate the design of the controller: the models are not polluted by scheduling details.
- We can have separate models corresponding to different scheduling algorithms.

4 A Signal Control System

In this section, we first present a requirements document of a signal control systems, then subsequently describe our formal development, applying our proposed development strategy.




Fig. 1. A signal control system

4.1 Requirements Document

Our aim is to develop a signal control system at a particular train station. The overview of the system can be seen in Fig. 1.

Our first set of requirements concern the trains and the topology of the network.

ENV0 The station contains a number of *platforms* in between an *entry block* and an *exit block*.

ENV1 A train occupies *no more than one block*.

ENV2 The track is *one-way*, i.e. the train enters the station via the entry block and leaves the station via the exit block.

The next requirements concern the switches located at the two ends of the stations.

ENV3 There are two switches connecting the entry and exit block to some platforms, called in-switch and out-switch accordingly.

ENV4 A train at entry block can only enter some platform block if the in-switch is set to that particular block. Similarly for the out-switch.

The most important property of the system concerns safety: the system must guarantee that trains never collide. This is ensured by precluding the simultaneous presence of two trains on the same block.

SAF5 Two trains cannot be on the same block.

Two (light) signals are installed at the two ends of the station, called entry signal and exit signal respectively.

ENV6 There are two signals which are either red or green.

ENV7 Trains are assumed to stop at red signals.

The controller receives input from various sensors and output its commands via actuators.

ENV8 There are sensors detecting whether a block is occupied.

ENV9 There are sensors detecting the status of the signals.

ENV10 The sensors reflect the current status of the corresponding components.

We design a controller for changing the switch positions, i.e. connecting to which platform and changing the signal from red to green. The signals automatically change from green to red when some train passes by.

ENV11 For each signal, there is an actuator for the controller to command the signal to turn from red to green.

ENV12 The signals change from green to red when a train passes by.

ENV13 For each switch, there is an actuator for the controller to command the switch to change to a specific platform.

4.2 Stage 1. The Model of the Environment

In the first stage, we build a model of the environment. We proceed step-by-step by introducing the details of the system in the following sequence: the blocks, the switches, the signals and the trains.

The Blocks. We specify the set of blocks (*BLOCK*) containing an entry (*ENT*), an exit (*EXT*) and a set of platform blocks (*PLFS*). This corresponds to our requirement **ENV0**. In the initial model, variable *OCC* is used to record the set of occupied blocks. Initially, *OCC* is assigned \emptyset , the empty set. There are four different events, namely *ARRIVE*, *MOVE.IN*, *MOVE.OUT*, and *LEAVE*, to model different cases on how the status of a block can be changed.

ARRIVE : *ENT* becomes occupied (because of an arriving train).

MOVE.IN : *ENT* becomes unoccupied and a platform becomes occupied (because of a train moving into the station).

MOVE.OUT : *EXT* becomes occupied and a platform becomes unoccupied (because of a train moving out of the station).

LEAVE : *EXT* becomes unoccupied (because of a leaving train).

Because of the symmetry between these events, we only present the events *MOVE.IN* and *ARRIVE*, and their subsequent refinements.

$\begin{aligned} \text{ARRIVE} &\hat{=} \text{when } ENT \notin OCC \text{ then } OCC := OCC \cup \{ENT\} \text{ end} \\ \text{MOVE.IN} &\hat{=} \text{any } p \text{ where } p \notin OCC \wedge ENT \in OCC \text{ then } OCC := (OCC \cup \{p\}) \setminus \{ENT\} \text{ end} \end{aligned}$
--

The Switches. In this refinement step, we introduce the variables *IN_SW* and *OUT_SW* to model the two switches located at the two end of the station. The status each switch represents which platform block they are connected to. This corresponds to our requirement **ENV3**. Initially, the switches are set arbitrarily to any platform.

We only need to adjust event *MOVE.IN*. Its parameter *p* is instantiated with the platform the in-switch is connected to. This reflects requirement **ENV4**.

$\text{MOVE.IN} \hat{=} \text{when } IN_SW \notin OCC \wedge ENT \in OCC \text{ then } OCC := (OCC \cup \{IN_SW\}) \setminus \{ENT\} \text{ end}$

The step is finalised by providing a means to change the switches. Events *TURN.IN_SW* and *TURN.OUT_SW* are introduced and we leave the choice arbitrary. We focus on *TURN.IN_SW* and its refinements.

$\text{TURN.IN_SW} \hat{=} \text{begin } IN_SW := \in PLFS \text{ end}$

The Signals. In this refinement step, we introduce the signals ENT_SGN and EXT_SGN located at the two ends of the station. The signals are either RED (meaning that passage is *forbidden*) or GRN (meaning that passage is *allowed*). This corresponds to requirement **ENV6**. Initially, both signals are RED .

We refine $MOVE_IN$ event accordingly, by refining its guards using ENT_SGN instead of referring directly to the status of ENT block. This also reflects the requirement that trains obey the signals (**ENV7**).

However, this guard substitution is only valid if it constitutes a strengthening which we ensure by introducing the following invariant.

$$\mathbf{inv2.0} : ENT_SGN = GRN \Rightarrow IN_SW \notin OCC$$

To preserve **inv2.0**, we make sure that the signal becomes red as soon as the platform designated by the in-switch becomes occupied (**ENV12**).

```

MOVE.IN
  when ENT_SGN = GRN ∧ ENT ∈ OCC then
    OCC := (OCC ∪ {IN_SW}) \ {ENT}
    ENT_SGN := RED
  end

```

There are two new events to change the status of the signals, namely $ALLOW_ENTRY$ and $ALLOW_EXIT$. We show here $ALLOW_ENTRY$ only, taking into account **inv2.0**.

$$ALLOW_ENTRY \cong \text{when } IN_SW \notin OCC \text{ then } ENT_SGN := GRN \text{ end}$$

We also strengthen the guard of event $TURN_IN_SW$ for safety reasons.

$$TURN_IN_SW \cong \text{when } ENT_SGN = RED \text{ then } IN_SW := PLFS \text{ end}$$

The Trains. In the last refinement of the environment model, we introduce the trains into the system. The safety properties concerning the trains all concern their position so this is a good candidate for a new variable. POS is thus introduced to map each train to the only block where it is located (as stated by **inv3.0**), consistently with **ENV1**. To rule out the possibility of collisions, i.e. to enforce **SAF5**, we can now introduce **inv3.1** which states that each train is alone on its block. Finally, for the sake of consistency with the variable OCC , we introduce **inv3.2** so that only trains can occupy a block.

$$\begin{aligned} \mathbf{inv3.0} : POS \in TRAIN \Rightarrow BLOCK \\ \mathbf{inv3.1} : \forall t_1, t_2. t_1 \in \text{dom}(POS) \wedge t_2 \in \text{dom}(POS) \wedge t_1 \neq t_2 \Rightarrow POS(t_1) \neq POS(t_2) \\ \mathbf{inv3.2} : \text{ran}(POS) = OCC \end{aligned}$$

In this model, events such as $TURN_IN_SW$ and $ALLOW_ENTRY$ stay unchanged since it does not directly interact with train positions. We refine events $MOVE_IN$ and $ARRIVE$ accordingly to include how the train position are updated.

```

MOVE.IN
  any t where
    ENT_SGN = GRN
    t ∈ dom(POS)
    POS(t) = ENT
  then
    OCC := (OCC ∪ {IN_SW}) \ {ENT}
    ENT_SGN := RED
    POS(t) := IN_SW
  end

```

```

ARRIVE
  any t where
    ENT ∉ OCC
    t ∉ dom(POS)
  then
    OCC := OCC ∪ {ENT}
    POS(t) := ENT
  end

```

Finally, **inv3.2** enables us to rewrite the guard of $MOVE_IN$ as shown.

4.3 Stage 2. The Actuators

At the end of the first stage we have an idealised model of the environment specifying how physical components should be working together. We introduce some actuators, i.e. output of the future controller, to commands the adaptation of the state of the environment component, in such a way that the normal behavior of the environment is coerced into the modelled behavior.

The *switch actuators* are used to send commands to the switches to change to a specific platform (**ENV13**). We focus on the actuator of the in-switch. Two new variables act_in_sw and $act_in_sw_plf$ are used to model the actuator: the former is a boolean to indicate whether there is a pending command for the device, the latter specifies which platform the switch should change to.

Similarly, the *signal actuators* are used for sending commands to set the signals to *GRN* (**ENV11**). The signal actuators are modelled as one boolean each, act_ent_sgn for the entry signal and act_ext_sgn for the exit signal, respectively.

Event TURN_IN_SW is refined accordingly using the command from the actuator. The actuator is reset after the switch changes.

```
(abs)TURN_IN_SW
when
  ENT_SGN = RED
then
  IN_SW := PLFS
end
```

```
(cnr)TURN_IN_SW
when
  act_in_sw = TRUE
then
  IN_SW := act_in_sw_plf
  act_in_sw := FALSE
end
```

Event ALLOW_ENTRY is refined similarly. We now introduce invariants to make sure that the substitution of guards is indeed a strengthening.

```
inv4.0 : act_out_sw = TRUE ⇒ EXT_SGN = RED
inv4.1 : act_ent_sgn = TRUE ⇒ IN_SW ∉ OCC
```

Finally, we create new controller events responsible for sending commands via the actuators to the in-switch ($ctrl_trigger_in_sw$) and to the entry signal ($ctr_chg_ent_sgn$).

```
ctrl_trigger_in_sw
any p where
  act_in_sw = FALSE
  ENT_SGN = RED
  act_ent_sgn = FALSE
  p ∈ PLFS
then
  act_in_sw := TRUE
  act_in_sw_plf := p
end
```

```
ctr_chg_ent_sgn
when
  act_ent_sgn = FALSE
  IN_SW ∉ OCC
  act_in_sw = FALSE
then
  act_ent_sgn := TRUE
end
```

Notice that the guards of these events guarantee that the newly introduced invariants are maintained.

4.4 Stage 3. The Sensors and the Controller

We are now ready to introduce the sensors together with the assumption that they reflect the status of the actual physical components. This is straightforward and we add

variables for the sensors: $snsr_occ$, $snsr_ent_sgn$ and $snsr_ext_sgn$ corresponding respectively to the sensors of the blocks, the entry signal and the exit signal (**ENV8**, **ENV9**). These new variables are glued with the old variables using invariants, such as $snsr_occ = OCC$, corresponding to requirement **ENV10**.

Within the controller events such as $ctrl_trigger_out_sw$ references to the status of a physical component such as OCC are replaced by the corresponding sensor, in this case $snsr_occ$.

Finally, since IN_SW is used in the guard of $ctr_chg_ent_sgn$, the controller needs to know the status of the in-switch when sending the command for changing the entry signal. The controller keeps a copy of status of the in-switch with its variable $ctrl_in_sw$. Note that variable $ctrl_in_sw$ does not necessarily reflect the current value of IN_SW . Indeed, we only need them to be the same when there is no actuator command for the in-switch. $ctrl_in_sw$ is updated when the controller commands the corresponding switch to change with event $ctrl_trigger_in_sw$.

4.5 Stage 4. Scheduling

At the end of Stage 3, we have a model of the signal control system including its working environment which guarantees to satisfy our safety requirement. We can then impose extra scheduling algorithm for our controller for optimising its execution. In our Event-B model, it is done by merely strengthening guards of the controller events. As an example, we show here the optimisation for event $ctrl_trigger_in_sw$ so that the in-switch

- changes only to a new free platform, i.e. $p \notin snsr_occ \wedge p \neq ctrl_in_sw$, and
- only when the entry block is occupied, i.e. $ENT \in snsr_occ$.

5 Conclusion

We have presented our development strategy for modelling control systems working within some fragile environment. Our strategy starts with the modelling of the environment, followed by the introduction of the actuators, before the controller and sensors are modelled. Finally, further scheduling details are imposed on the controller as an optimisation step for the system. Applying our development strategy reduces the difficulty in modelling this type of systems, results in models which are easy to understand and verify.

Our development strategy is initially inspired by the development of an elevator system by Laurent Voisin, which has been used as a student project for a course on Event-B at ETH Zurich. We have applied the approach to several systems of this type, including a re-development of “Cars on a bridge” example from Abrial [1, Chapter 2]. Our approach is fundamentally different from the *inside-out* approach taken by Abrial. In contrast to our approach, Abrial starts by first modelling the controller and the environment is introduced after. Even though both approaches are possible for developing this type of systems, our *outside-in* approach is more constructive: instead of defining a controller and then proving that it fits the environment, we use the requirements to deduce constraints that the controller must fulfill and we go on to build it accordingly.

Compared to the guideline proposed by Butler in [2], our development strategy is similar in that it focuses initially on the model of the environment. The main difference is that we introduce the actuators before the sensors. This is influenced by our *backward* reasoning: we want to deduce the design of our controller and its input from constraint imposed on its output. We believe that this approach is simpler and gives stronger guidance for the design, similar to the reasoning using weakest-precondition [3].

The validation of control systems have been studied using other formal methods. Hansen validated a railway interlocking model using VDM [4]. However, the paper only establishes a model of the environment without the controller. Haxthausen and Peleska presented an approach using RAISE for developing a distributed railway control system [5]. Their approach consists of two stages. In their first stage, the model of the environment and controllers are developed globally together. Their second stage focuses on the design of a distributed controller corresponding to the model in the first stage. Our development strategy can be seen as a guideline for developing the model in their first stage.

One of the aspects that has not been captured in our example is the assumptions. Typically, they concern the speed of communication and response of the controller. It can be shown that using our development strategy, these assumptions arise naturally during the formal developments which otherwise will be difficult to find *a priori*. Furthermore, we have focused on the development of a system with some critical safety properties. Developing a systems with some liveness properties, e.g. all trains must eventually leave the station will require additional modelling guidelines.

References

1. Jean-Raymond Abrial. *Modeling in Event-B: System and Software Engineering*. Cambridge University Press, May 2010.
2. Michael Butler. Towards a Cookbook for Modelling and Refinement of Control Problems. Working paper, <http://deploy-eprints.ecs.soton.ac.uk/108/>, May 2009.
3. Edsger Dijkstra. *A Discipline of Programming*. Prentice Hall International, Englewood Cliffs, N.J., 1976.
4. Kirsten Mark Hansen. Validation of a railway interlocking model. In Maurice Naftalin, B. Tim Denvir, and Miquel Bertran, editors, *FME*, volume 873 of *LNCS*, pages 582–601. Springer, 1994.
5. Anne Elisabeth Haxthausen and Jan Peleska. Formal development and verification of a distributed railway control system. *IEEE Trans. Software Eng.*, 26(8):687–701, 2000.
6. S. Yeganehfar, M. Butler, and A. Rezazadeh. Evaluation of a guideline by formal modelling of cruise control system in Event-B. In *Proceedings of NFM 2010*, pages 182–191, 2010.