Event-B patterns and their tool support

Journal Article**Author(s):**

Hoang, Thai Son; Fürst, Andreas; Abrial, Jean-Raymond

Publication date:

2013-05

Permanent link:

<https://doi.org/10.3929/ethz-b-000069466>

Rights / license:

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

Originally published in:

Software and Systems Modeling 12(2), <https://doi.org/10.1007/s10270-010-0183-7>

Event-B patterns and their tool support

Thai Son Hoang · Andreas Fürst ·
Jean-Raymond Abrial

Received: 15 June 2010 / Revised: 11 November 2010 / Accepted: 17 November 2010 / Published online: 4 January 2011
© Springer-Verlag 2010

Abstract Event-B has given developers the opportunity to construct models of complex systems that are correct-by-construction. However, there is no systematic approach, especially in terms of reuse, which could help with the construction of these models. We introduce the notion of *design patterns* within the framework of Event-B to shorten this gap. Our approach preserves the correctness of the models, which is critical in formal methods and also reduces the proving effort. Within our approach, an Event-B design pattern is just another model devoted to the formalisation of a typical sub-problem. As a result, we can use patterns to construct a model which can subsequently be used as a pattern to construct a larger model. We also present the interaction between developers and the tool support within the associated RODIN Platform of Event-B. The approach has been applied successfully to some medium-size industrial case studies.

Keywords Event-B · Formal methods · Design patterns · Formal modelling · Model reuse

1 Introduction

The purpose of our investigation here is to study the possibility of reusing models in formal modelling. Currently, formal methods are applicable to various domains for constructing models of complex systems. However, often they lack some systematic methodological approaches, in particular, in reusing existing models, for helping the development process. The objective in introducing design patterns within formal methods, in general, and in Event-B, in particular, is to overcome this limitation.

The idea of design patterns in software engineering is to have a general and reusable solution to commonly occurring problems. In general, a design pattern is not necessarily a finished product, but rather a template on how to solve a problem which can be used in many different situations. Design patterns are further populated in object-oriented programming [14]. The idea is to have some predefined solutions, and incorporate them into the development with some modification and/or instantiation. We want to bring this idea into formal methods and, in particular, to Event-B. Moreover, the typical elements that we want to reuse are not only the models themselves, but also (more importantly) their correctness in terms of proofs associated with the models. In our earlier investigations [5, 11, 16] and [10, Sect. 5.4.1], we have already worked on several examples to understand the usefulness and applicability of the approach. We summarise this work and its formalisation in this article.

Our contribution here is the methodology for reusing existing models in Event-B. Our approach allows developers to reuse any existing models (which we call “design patterns”) in a way that preserves the correctness of models, hence we can save effort on not only modelling but also on proving these models correct.

Communicated by Paddy Krishnan, Antonio Cerone, and Dang Van Hung.

This is an extension of an earlier report [16].
Part of this work is supported by the DEPLOY project
(<http://www.deploy-project.eu>).

T. S. Hoang (✉) · A. Fürst
Swiss Federal Institute of Technology (ETH-Zurich),
Zurich, Switzerland
e-mail: htson@inf.ethz.ch

J.-R. Abrial
Marseille, France

The examples that we used in this article are models for communication protocols [23]. Note that, however, the approach is general and its applicability *is not limited* to this domain.

The structure of the article is as follows. Section 2 gives a short introduction to Event-B. Section 3 presents a case study to illustrate the motivation for our approach. Section 4 gives an overview of the formalisation of the approach in Event-B. The list of patterns which are used in our industrial case studies is presented in Sect. 5. Section 6 describes our prototype tool supporting the approach. Finally, in Sect. 7 we review related work and point out future directions.

2 The Event-B modelling method

Event-B [2] represents a further evolution of the B-method [1], which has been simplified and is now centered around the general notion of *events*, also found in Action Systems [6] and TLA [17].

An Event-B [2] model is a collection of modelling elements that are stored in a repository. When presenting our models, we will do so in a pretty-print form, e.g. adding keywords and following a certain layout convention to aid parsing. We proceed like this to improve legibility and help the reader to remember the different constructs of Event-B. The syntax should be understood as a convention for presenting Event-B models in textual form rather than defining a language.

Event-B models are described in terms of the two basic constructs: *contexts* and *machines*. Contexts contain the static part of a model whereas machines contain the dynamic part. Contexts may contain *carrier sets*, *constants* and *axioms*, where carrier sets are similar to types [4]. In this article, we simply assume that there is some context and do not mention it explicitly. Machines are presented in Sect. 2.1, and machine refinement in Sect. 2.2.

2.1 Machines

Machines provide behavioural properties of Event-B models. Machines may contain *variables*, *invariants*, and *events*.¹ Variables v define the state of a machine. They are constrained by invariants $I(v)$. Possible state changes are described by means of events. Each event is composed of a *guard* $G(v)$ and an *action* $S(v)$.² The guard states the necessary condition under which an event may occur, and the action describes how the state variables evolve when the event

occurs. An event can be represented by the following form

$$\text{evt} \triangleq \text{when } G(v) \text{ then } S(v) \text{ end} \quad (1)$$

The short form

$$\text{evt} \triangleq \text{begin } S(v) \text{ end} \quad (2)$$

is used if the guard always holds. A dedicated event of the form (2) is used for *initialisation*.

The action of an event is composed of several *assignments* of the form

$$x := E(v) \quad (3)$$

$$x : \in E(v) \quad (4)$$

$$x : | Q(v, x'), \quad (5)$$

where x are some variables, $E(v)$ expressions, and $Q(v, x')$ a predicate. Assignment form (3) is *deterministic*, the other two forms are *non-deterministic*. Form (4) assigns x to an element of a set, and form (5) assigns to x a value x' satisfying a predicate. The effect of each assignment can also be described by a before–after predicate BAP:

$$\text{BAP}(x := E(v)) \triangleq x' = E(v) \quad (6)$$

$$\text{BAP}(x : \in E(v)) \triangleq x' \in E(v) \quad (7)$$

$$\text{BAP}(x : | Q(v, x')) \triangleq Q(v, x'). \quad (8)$$

A before–after predicate describes the relationship between the state just before an assignment has occurred (represented by unprimed variable names x) and the state just after the assignment has occurred (represented by primed variable names x'). All assignments of an action $S(v)$ occur simultaneously which is expressed by conjoining their before–after predicates, yielding a predicate $A(v, x')$. Variables y that do not appear on the left-hand side of an assignment of an action are not changed by the action. Formally, this is achieved by conjoining $A(v, x')$ with $y' = y$, yielding the before–after predicate of the action:

$$\text{BAP}(S(v)) \triangleq A(v, x') \wedge y' = y. \quad (9)$$

Later, in proof obligations, we represent the before–after predicate $\text{BAP}(S(v))$ of an action $S(v)$ directly by the predicate $\mathbf{S}(v, v')$.

Proof obligations serve to verify certain properties of a machine. Here, a proof obligation is presented in the form of a sequent: “hypotheses” $\vdash$ “goal”. The intuitive meaning of this sequent is that under the assumption of the *hypotheses*, the *goal* holds.

For each event of a machine, the following proof obligation which guarantees *feasibility* must be proved.

¹ Machine can also contain a *variant* for proving convergence properties, but it is not of our interests in this article.

² For simplicity, we do not treat events with *parameters*.

$\begin{array}{l} I(v) \\ G(v) \\ \vdash \\ \exists v' \cdot S(v, v') \end{array}$	FIS
--	------------

By proving feasibility, we achieve that $S(v, v')$ provides an after state whenever $G(v)$ holds. This means that the guard indeed represents the enabling condition of the event.

Invariants are supposed to hold whenever variable values change. Obviously, this does not hold a priori for any combination of events and invariants and, thus, needs to be proved. The corresponding proof obligation is called *invariant preservation*:

$\begin{array}{l} I(v) \\ G(v) \\ \vdash \\ S(v, v') \\ I(v') \end{array}$	INV
--	------------

Similar proof obligations are associated with the initialisation event of a machine. The only difference is that the invariant and guard do not appear in the antecedent of the proof obligations (**FIS**) and (**INV**).

2.2 Machine refinement

Machine refinement provides a mean to introduce more details about the dynamic properties of a model [4]. For more on the well-known theory of refinement, we refer to the Action System formalism [6] that has inspired the development of Event-B. We present some important proof obligations for machine refinement.

A machine CM can refine at most one other machine AM . We call AM the *abstract* machine and CM the *concrete* machine. The state of the abstract machine is related to the state of the concrete machine by a *gluing invariant* $J(v, w)$, where v are the variables of the abstract machine and w the variables of the concrete machine.

Each event ea of the abstract machine is *refined* by one or more concrete events ec . Let abstract event ea and concrete event ec be:

$ea \hat{=} \text{when } G(v) \text{ then } S(v) \text{ end}$
 $ec \hat{=} \text{when } H(w) \text{ then } T(w) \text{ end}$

Somewhat simplified, we can say that ec refines ea if the following conditions hold.

1. The concrete event is feasible. This is formalised by the following proof obligation.

$\begin{array}{l} I(v) \\ J(v, w) \\ H(w) \\ \vdash \\ \exists w' \cdot T(w, w') \end{array}$	FIS_REF
---	----------------

2. The guard of ec is *stronger* than the guard of ea . This is formalised by the following proof obligation.

$\begin{array}{l} I(v) \\ J(v, w) \\ H(w) \\ \vdash \\ G(v) \end{array}$	GRD
--	------------

3. The abstract event can always “simulate” the concrete event and preserve the gluing (concrete) invariant. This is formalised by the following proof obligation.

$\begin{array}{l} I(v) \\ J(v, w) \\ H(w) \\ \vdash \\ T(w, w') \\ \exists v' \cdot S(v, v') \wedge J(v', w') \end{array}$	SIM
--	------------

For the initialisation, the corresponding proof obligations are analogue. The proofs of these above obligations ensure the correctness of the refinement model with respect to the abstract model and the gluing invariant between them.

In the course of refinement, often *new events* ec are introduced into a model. New events must be proved to refine the implicit abstract event **skip** that does nothing.

$\text{skip} \hat{=} \text{begin SKIP end}$

Moreover, it may be proved that new events do not collectively diverge, but this is not relevant here. The new events allow us to observe the system with a finer time grain. This is an analogue of the stuttering principle in TLA [17]: a step that leaves the abstract variables unchanged.

3 Question/Response protocol

In this section, we look at the development of a protocol, namely *Question/Response* in order to understand what we mean by design patterns and how to apply them in system development. Section 3.1 first gives an informal description of the protocol together with its formal specification in Event-B, then identifies *similar fragments* of the formal model that leads to the idea of using patterns. In Sect. 3.2

we formally present a pattern, namely *synchronous multiple message communication*, including its specification and refinement. Finally, we illustrate how the pattern is reused (twice) in our development of the actual Question/Response protocol in Sect. 3.3.

3.1 Description and formal specification

There are two parties participating in this protocol namely the *Questioner* and the *Responder*. The protocol consists of an unbounded number of *rounds*. For each round, there are two steps as follows.

1. The *Questioner* sends a *question* to the *Responder*.
2. After receiving this *question*, the *Responder* sends a *response* back to the *Questioner*.

Formally, we can use two variables to represent the state of the protocol: *quest* to denote the number of questions that have been asked, and *resp* to indicate the number of responses that have been given. The first invariant **QuestResp_0_1** specifies that the number of responses is a natural number and the second invariant, i.e. **QuestResp_0_2** specifies that the communication is synchronous: either the number of questions is the same as the number of responses or it is greater than the number of responses by 1—in the case where a response is expected before another question can be created.

variables: *quest, resp*

invariants:

QuestResp_0_1: $resp \in \mathbb{N}$
QuestResp_0_2: $quest = resp \vee quest = resp + 1$

Initially, there are no questions or responses hence both variables are initialised to 0.

init
begin
 $quest, resp := 0, 0$
end

The dynamic system can be seen in Fig. 1. For each round, the “questioning” phase starts when the number of questions and the number of responses are identical and increases the number of questions by 1. The “responding” phase starts after the “questioning” phase of the same round (when the number of questions and responses are different) and increases the number of responses by 1. This is formalised by the following two events, namely *questions* and *responds*, representing the two phases accordingly.

Fig. 1 Question/Response protocol with two rounds

questions
when
 $quest = resp$
then
 $quest := quest + 1$
end

responds
when
 $quest \neq resp$
then
 $resp = resp + 1$
end

The specification of the above two events are very similar, except for their guards. The two events both correspond to transferring some information from one side to another and can be *repeated*; however, the communication is *synchronous*: a new message can be sent only when the last message has been received. We call this kind of communication *synchronous multiple message communication*. Hence, if we have a development for this type of communication (to be formalised in the next section), we can instantiate it twice: once for the “questioning” phase and once for the “responding” phase.

3.2 Synchronous multiple message communication

This section presents the development of a communication between two parties *A* and *B* for transferring some information repeatedly and synchronously from *A* to *B*.

The specification of this protocol contains only one natural number variable *trans* to denote the number of messages that have been transferred.

variables: *trans*

invariants:

SynchMultiCom_0_1: $trans \in \mathbb{N}$

There is only one event in this model to increase the value of variable *trans* denoting the fact that a message has been transferred from *A* to *B*.

```

transfers
begin
  trans := trans + 1
end

```

This synchronous multiple message communication is illustrated in Fig. 2.

However, this is only the abstraction of this protocol (it might be even too abstract in the sense that it does not specify how communication happens, e.g. synchronous vs. asynchronous). In reality, the message needs to be sent via some channel between the two parties. This is illustrated in Fig. 3. Here, the diagram is about different parties (not states) and messages sent between them.

We use three variables to represent the state of the refinement.

- *snds*: the number of messages having been sent by *A*.
- *rcvs*: the number of messages having been received by *B*.
- *chan*: since there is at most one message on the channel, we use a Boolean value to denote the existence of a message on the channel.

At this point, we have a decision to make about refinement of the abstract event **transfers**. It could be refined by the event corresponding to “sends” or it could be refined by the event corresponding to “receives”. We presented here the refinement of event **transfers** when sending, *but the other alternative is also possible*. As a result of this choice, we have the following gluing invariant.

```

invariants:
  SynchMultiCom_1_1: trans = snds

```

We also have additional technical invariants about the properties of the protocol. First, if there is no message on the channel, the number of sent and received messages are the same. Second, if there is a message on the channel, then the number of sent messages is greater than the number of received messages by exactly 1. These two invariants correspond to the

Fig. 2 Synchronous multiple message communication

Fig. 3 Communication via a channel

“synchronous” communication behaviour. Finally, the number of received messages must be a natural number.

```

invariants:
  SynchMultiCom_1_2: chan = F => snds = rcvs
  SynchMultiCom_1_3: chan = T => snds = rcvs + 1
  SynchMultiCom_1_4: rcvs ∈ ℕ

```

Initially, there are no messages that have been sent, received or are in the channel.

```

init
begin
  snds := 0
  rcvs := 0
  chan := F
end

```

Events **sends** and **receives** are straightforward as follows.

```

sends
refines transfers
when
  chan = F
then
  chan := T
  snds := snds + 1
end

```

```

receives
when
  chan = T
then
  chan := F
  rcvs := rcvs + 1
end

```

Event **sends** is enabled if there is no message in the channel. The action of the event specifies that *A* now sent one more message and the message is in the *channel*. Event **receives** is enabled when there is a message in the *channel*. The action of the event removes the message from the *channel* and indicates that *B* has received one more message. Note that event **receives** here is a new event (i.e. it refines **skip**).

3.3 Using the pattern for the protocol

In this section, we see how the pattern developed in Sect. 3.2 is used for developing the Question/Response protocol of Sect. 3.1. There are four steps as follows.

1. We need to “match” the specification of the pattern with the problem.
2. We need to “syntactically check” the matching to see if the pattern is applicable.
3. We have to “rename” those variables and events in the pattern refinement that would lead to a name clash (since we can instantiate the same pattern many times). We can also “rename” non-conflicting variables and events if we like to.
4. Lastly, we “incorporate” the renamed refinement of the pattern to create a refinement of the problem.

As mentioned before, we can instantiate the synchronous multiple message communication pattern twice for the Question/Response protocol: once for the “questioning” phase and a second time for the “responding” phase.

3.3.1 Pattern for “questioning” phase

We follow the different steps to incorporate a synchronous multiple message communication pattern for the “questioning” phase as follows.

1. As a first step we need to identify the “matching” between the specification of the pattern and the problem. The matching here is straightforward with variable *trans* and event *transfers* of the pattern matched with variable *quest* and event *questions* of the problem accordingly.

pattern	$\rightsquigarrow$	problem
<i>trans</i>	$\rightsquigarrow$	<i>quest</i>
<i>transfers</i>	$\rightsquigarrow$	<i>questions</i>

2. The second step is to syntactically check the validity of the pattern. For example, we need to check that given the variable matching *trans* $\rightsquigarrow$ *quest*, the action of event *transfers* is “matched” with the action of *questions*. This should be done automatically by a tool. At the moment, we can assure ourselves that this step is valid. More information about this step can be seen in Sect. 6.2 when we discuss about tool support.
3. The third step is to rename the variables and events of the pattern refinement according to the following rules.

original	$\rightsquigarrow$	renamed as
<i>snds</i>	$\rightsquigarrow$	<i>QQuestSnds</i>
<i>chan</i>	$\rightsquigarrow$	<i>Q2RQuestChan</i>
<i>recv</i>	$\rightsquigarrow$	<i>RQuestRcv</i>
<i>sends</i>	$\rightsquigarrow$	<i>Q_sends_question</i>
<i>receives</i>	$\rightsquigarrow$	<i>R_receives_question</i>

4. In the last step, we incorporate the renamed refinement of the pattern to create a refinement of the problem. The result is the following model.

variables: *resp*,
QQuestSnds,
RQuestRcv,
Q2RQuestChan

invariants:

QuestResp_1_1: *quest* = *QQuestSnds*
QuestResp_1_2: *Q2RQuestChan* = *F* $\Rightarrow$
 QQuestSnds = *RQuestRcv*
QuestResp_1_3: *Q2RQuestChan* = *T* $\Rightarrow$
 QQuestSnds = *RQuestRcv* + 1
QuestResp_1_4: *RQuestRcv* $\in \mathbb{N}$

init
begin
 resp := 0
 Q2RQuestChan := *F*
 RQuestRcv := 0
 QQuestSnds := 0
end

Q_sends_question
refines *questions*
when
 QQuestSnds = *resp*
 Q2RQuestChan = *F*
then
 Q2RQuestChan := *T*
 QQuestSnds = *QQuestSnds* + 1
end

R_receives_question
when
 Q2RQuestChan = *T*
then
 Q2RQuestChan := *F*
 RQuestRcv := *RQuestRcv* + 1
end

responds
refines *responds*
when
 QQuestSnds $\neq$ *resp*
then
 resp := *resp* + 1
end

There are a number of important aspects of the pattern which we want to draw the readers’ attention.

- The matching between event *transfers* and event *questions* is not exact.

transfers
begin
 trans := *trans* + 1
end

questions
when
 quest = *resp*
then
 quest := *quest* + 1
end

Taking into account the matching of the variables, i.e. *trans* becomes *quest*, only the actions of those events are matched. The guard of event **questions** does not correspond to any guard of event **transfers**.

- The additional guard of event **questions**, i.e. *quest = resp* is transformed into the guard $QQuestSnds = resp$ of event **Q_sends_question** in the resulting refinement, because variable *quest* is matched with variable *trans* of the pattern and this variable is subsequently refined to $QQuestSnds$, according to the invariant **QuestResp_1_1**.

QuestResp_1_1: $quest = QQuestSnds$

- Similarly, the guard of event **responds**, i.e. *quest ≠ resp*, needs to take into account the fact that variable *quest* now becomes $QQuestSnds$.
- The rewriting of these additional guards is done automatically by the tool support.

3.3.2 Pattern for “responding” phase

We now follow similar steps to use the synchronous multiple message communication pattern for the “responding” phase.

1. The matching is as follows

pattern	↔	problem
<i>trans</i>	↔	<i>resp</i>
transfers	↔	responds

2. Similarly, we assure that the syntax checking for the given matching is successful.
3. We rename the refinement of the pattern according to the following rules.

original	↔	renamed as
<i>snds</i>	↔	<i>RRespSnds</i>
<i>chan</i>	↔	<i>R2QRespChan</i>
<i>rcvs</i>	↔	<i>QRespRcvs</i>
sends	↔	R_sends_response
receives	↔	Q_receives_response

4. We incorporate the renamed pattern refinement with the problem to obtain the following model.

variables: $QQuestSnds$,
 $RQuestRcvs$,
 $Q2RQuestChan$,
 $RRespSnds$,
 $QRespRcvs$,
 $R2QRespChan$

invariants:

QuestResp_2_1: $resp = RRespSnds$
QuestResp_2_2: $R2QRespChan = F \Rightarrow RRespSnds = QRespRcvs$
QuestResp_2_3: $R2QRespChan = T \Rightarrow RRespSnds = QRespRcvs + 1$
QuestResp_2_4: $QRespRcvs \in \mathbb{N}$

```
init
begin
  Q2RQuestChan := F
  RQuestRcvs := 0
  QQuestSnds := 0
  R2QRespChan := F
  QRespRcvs := 0
  RRespSnds := 0
end
```

```
Q_sends_question
refines Q_sends_question
when
  QQuestSnds = RRespSnds
  Q2RQuestChan = F
then
  Q2RQuestChan := T
  QQuestSnds = QQuestSnds + 1
end
```

```
R_receives_question
refines R_receives_question
when
  Q2RQuestChan = T
then
  Q2RQuestChan := F
  RQuestRcvs := RQuestRcvs + 1
end
```

```
R_sends_response
refines responds
when
  QQuestSnds ≠ RRespSnds
  R2QRespChan = F
then
  R2QRespChan := T
  RRespSnds := RRespSnds + 1
end
```

```
Q_receives_response
when
  R2QRespChan = T
then
  R2QRespChan := F
  QRespRcvs := QRespRcvs + 1
end
```

Again, we highlight some important aspects of our pattern application at this step.

- Similar to the previous pattern application in Sect. 3.3.1, the matching between event transfers and event responds are not exact: there is an additional guard in event responds.
- This guard of event responds, i.e. $QQuestSnds \neq resp$ needs to take into account the fact that variable $resp$ is matched with variable $trans$ of the pattern specification and this variable is later refined to $RRespSnds$. This guard is transformed into the guard $QQuestSnds \neq RRespSnds$ of the resulting event $R_sends_response$. Similarly, for the guard of $Q_sends_question$.
- These guards are in fact “cheats” in the model. Event $Q_sends_question$ supposes to be an event of the *Questioner*; however, its guard refers to variable $RRespSnds$ of the *Responder*. The same analysis applies for event $R_sends_response$ and variable $QQuestSnds$. This problem will be handled by a standard refinement step in the next section.

3.3.3 Removing the “cheating” guards

The problem that we mentioned earlier about the “cheating” guards is better known as *local enforceability* [9]. Roughly speaking, on the abstraction level, the global interactions between partners are specified in a way that it might not be enforced during real local implementation without having more additional interactions between the different partners. In our case, it is not possible for the *Questioner* to have access to the information belonging to the *Responder*: currently, event $Q_sends_question$ has access to variable $RRespSnds$ of the *Responder*. In this section, we fix this problem by adding more information on how the two partners interact with each other.

The cheating guards, i.e.

$$QQuestSnds \neq RRespSnds$$

for event $R_sends_response$ can be replaced by the following guard which uses only variables of the *Responder*:

$$RQuestRcvS \neq RRespSnds.$$

The proof for the guard strengthening obligation (GRD) is based on the following invariant **QuestResp_3_1** (which we need to add to the model).

invariants:
QuestResp_3_1: $RQuestRcvS \geq RRespSnds$

The reasoning is as follows:

- From the new guard $RQuestRcvS \neq RRespSnds$ and the new invariant $RQuestRcvS \geq RRespSnds$, we have

$$RQuestRcvS > RRespSnds. \quad (10)$$

- We conclude from the existing invariants **QuestResp_1_2** and **QuestResp_1_3** that

$$QQuestSnds \geq RQuestRcvS. \quad (11)$$

- From (10) and (11), we conclude that $QQuestSnds > RRespSnds$, which ensures $QQuestSnds \neq RRespSnds$, as required.

This step is a standard refinement in Event-B. Intuitively, the new invariant *links* the *questioning* and *responding* phases together and is the core of the Question/Response protocol.

Similarly, the guard $QQuestSnds = RRespSnds$ of event $Q_sends_question$ is replaced by $QQuestSnds = RQuestRcvS$. The refined events $Q_sends_question$ and $R_sends_response$ at their final form are as follows.

```

Q_sends_question
refines Q_sends_question
when
  QQuestSnds = RQuestRcvS
  Q2RQuestChan = F
then
  Q2RQuestChan := T
  QQuestSnds = QQuestSnds + 1
end

```

```

R_sends_response
refines R_sends_response
when
  RQuestRcvS \neq RRespSnds
  R2QRespChan = F
then
  R2QRespChan := T
  RRespSnds := RRespSnds + 1
end

```

Note that we can consider also the guard referring to the channels, i.e. $R2QRespChan = F$ and $Q2RQuestChan = F$ as not locally enforceable, hence should be removed. However, this is not of our interest here.

Overall, this (standard) refinement step where we impose the policy for local enforceability cannot be done automatically by a tool: this corresponds to how the protocol is constructed and is usually protocol dependent.

4 Pattern incorporation in Event-B

In this section, we summarise the idea of incorporating patterns into Event-B developments. The process can be seen in Fig. 4.

First of all, in our notion, a pattern is just a development in Event-B including specification p_0 and a refinement p_1 .³

³ In general, this can be extended to multiple refinement level.

Fig. 4 Using patterns in Event-B

During a normal development in Event-B, at refinement m_n , developers can match part of the model with the pattern specification p_0 . As a result of this matching, the refinement p_1 can be incorporated to create the refinement m_{n+1} of m_n (with possible “renaming” to avoid name clashes).

Moreover, we have presented here the incorporation of each synchronous multiple message communication pattern separately. However, it is possible that they could be incorporated at the same time. In other words, there can be more than one pattern that can be matched at the same time with the problem at hand. There are side conditions to guarantee that the patterns do not interfere with each other, e.g. there should be no matching to the same variable.

4.1 Formalisation of the approach

We assume that we have the following patterns containing a specification p_0 and its refinement p_1 . We further assume that the pattern specification p_0 has some variables v with invariant $J(v)$. We consider a particular event p with guard $L(v)$ and some actions $v :| T(v, v')$.

In the refinement p_1 of p_0 , variable v is data refined by variable w with gluing invariant separated into $v = X(w)$ and $K(v, w)$. Here, we make the assumption that the gluing invariant can be functionally expressed as $v = X(w)$ with some other extra invariants $K(v, w)$. This assumption is valid for all our examples so far and could be relaxed later. Event

p is refined by event q with concrete guard $M(w)$ and some actions $w :| U(w, w')$.

We assume that we have arrived at a refinement level in a particular development which we call problem specification m_n . The machine has some variables b which we intend to match with the above pattern. Moreover, this problem specification could have some other variables c which we have to keep when incorporating the pattern into the development. We do not need to consider the invariant for this machine hence this is left out.

Without loss of generality, we consider two events of the problem specification: event e which is going to be matched with event p of the pattern specification, and event f which is not going to be matched. Event e is separated into the parts which are matched with event p of the pattern specification, taken into account the decision that variable b is matched with variable v of the pattern specification. Here, we say that every variable in the pattern need to be matched with some variable in the problem. However, this condition can be relaxed to make the approach more flexible (see future work in Sect. 7.3). Hence, the guard of the event is separated into $H(b)$ and $N(b, c)$, where $H(b)$ is matched with guard $L(v)$ of event p . Similarly, the action is separated into $b :| R(b, b')$ —which is a match of $v :| T(v, v')$ —and $c :| S(b, c, c')$. The validity of this matching can be syntactically checked and/or even be “discovered” by a tool. For the unmatched event f , we require that it must not change variable b , hence its action is of the form $c :| P(b, c, c')$. However, it can refer to b in the guard and in the action (only as reference to the before state). The preservation of this restriction will be checked by the supporting tool (more information in Sect. 6.2). The matching and the extraction from the gluing invariant can be summarised as follows.

pattern	$\rightsquigarrow$	problem
v	$\rightsquigarrow$	b
p	$\rightsquigarrow$	e
$L(v)$	$\rightsquigarrow$	$H(b)$
$v : T(v, v')$	$\rightsquigarrow$	$b : R(b, b')$

The refinement m_{n+1} of m_n is generated by combining the problem specification and the pattern refinement as follows.

We must guarantee that the constructed machine m_{n+1} is indeed a refinement of the specification m_n . The detailed proofs are in [11, Sect. 4.5]. Intuitively, the proofs assume the correctness of the problem specification m_n , the pattern specification p_0 and the pattern refinement p_1 in order to prove the correctness of the problem refinement m_{n+1} . The obligation list includes feasibility, guard strengthening and simulation for both events **e** and **f**.

As an example, we sketch the proof for guard strengthening obligation of event **e** which is stated as follows.

$$\begin{array}{l}
 b = X(w) \\
 K(b, w) \\
 J(b) \\
 M(w) \\
 N(X(w), c) \\
 \vdash \\
 H(b) \wedge N(b, c)
 \end{array}$$

The proof of the above sequent can be split into two parts since the goal is a conjunction.

$$\begin{array}{l}
 b = X(w) \\
 K(b, w) \\
 J(b) \\
 M(w) \\
 N(X(w), c) \\
 \vdash \\
 H(b)
 \end{array}
 \quad (12)$$

$$\begin{array}{l}
 b = X(w) \\
 K(b, w) \\
 J(b) \\
 M(w) \\
 N(X(w), c) \\
 \vdash \\
 N(b, c)
 \end{array}
 \quad (13)$$

The second part of the proof (13) for proving $N(b, c)$ follows from the assumptions $b = X(w)$ and $N(X(w), c)$. The first part (12) of the proof relies on the fact that event **q** is a

refinement of event **p** in the pattern, hence we have proved the guard strengthening obligation for **q**, namely.

$$\begin{array}{l}
 J(v) \\
 v = X(w) \\
 K(v, w) \\
 M(w) \\
 \vdash \\
 L(v)
 \end{array}$$

Moreover, from the matching information v is matched with b and guard $H(b)$ is matched with $L(v)$ (i.e. H and L are syntactically the same), we can derive (with renaming variable from v to b) the following.

$$\begin{array}{l}
 J(b) \\
 b = X(w) \\
 K(b, w) \\
 M(w) \\
 \vdash \\
 H(b)
 \end{array}$$

and from there we can conclude the proof for (12).

4.2 What we gain with the pattern approach

So far, it seems that we have to do more work in order to apply patterns: we have to develop the pattern separately and incorporate it into the main development. But we do have the following advantages.

- We do not need to prove that m_{n+1} is a refinement of m_n . This is because we have already done this proof when developing patterns.
- Moreover, we can reuse the pattern more than once. For example, in the development of the Question/Response protocol, we use the synchronous multiple message communication pattern twice, so we save doing proofs for one pattern.
- Since the pattern is just a normal Event-B development, the meaning of the pattern is also intuitive. Moreover, we can use any development as pattern in our approach.

The proof statistics related to the synchronous multiple message communication and Question/Response protocol is given in Table 1. As we can see, by developing the synchronous multiple message communication pattern separately, we have to prove 15 obligations. However, we do not need to prove the model “Question/Response 1” and “Question/Response 2” (which has a total of 32 obligations) since it is correct by construction. Hence, in total, we save $32 - 15$, i.e. 17 proofs. Note that the number of proof obligations for each model “Question/Response 1” and “Question/Response 2” is roughly the same as that of “Synch. Multi. Com. 1”, since in

Table 1 Proof statistics

Models	Total	Auto. (%)	Man. (%)
Synch. Multi. Com. 0	2	2 (100)	0 (0)
Synch. Multi. Com. 1	13	12 (92)	1 (8)
Question/Response 0	6	5 (83)	1 (17)
Question/Response 1	16	15 (94)	1 (6)
Question/Response 2	16	15 (94)	1 (6)
Question/Response 3	5	4 (80)	1 (20)

each model we apply the pattern once. The development of the two protocols is available on-line [13].

5 Patterns used in industrial case studies

Our approach has been applied to formalise communication protocols from SAP. The examples are *Buyer/Seller B2B* as described in [23] and *Ordering/Supply Chain A2A Communications* as described in [10, Section 5.3.3]. Table 2 shows the proof statistics comparing the developments without patterns and with patterns for the two case studies. More importantly, our approach save on average of the two case studies 33% of the manual proofs (those that need interactive efforts to discharge).

In this section, we give the description of other patterns that have been used in these protocols.

- Section 5.1 presents the *Single Message Communication* pattern.
- Section 5.2 presents the *Request/Confirm* pattern.
- Section 5.3 presents the *Request/Confirm/Reject* pattern.
- Section 5.4 presents the *Asynchronous Multiple Message Communication* pattern.
- Section 5.5 presents the *Asynchronous Multiple Message Communication with Repetition* pattern.

Table 2 Case studies' proof statistics (with vs. without pattern)

Models/savings	Total	Auto. (%)	Man. (%)
A2A (without pattern)	281	249 (89%)	32 (11%)
A2A (with pattern)	184	164 (89%)	20 (11%)
Savings	97	85 (88%)	12 (12%)
Savings percentage	35%	34%	38%
B2B (without pattern)	498	427 (86%)	71 (14%)
B2B (with pattern)	342	291 (85%)	51 (15%)
Savings	156	136 (87%)	20 (13%)
Savings percentage	31%	32%	28%

Fig. 5 Single Message Communication pattern

5.1 Single Message Communication pattern

The description of the pattern is as follows. There are two parties involved in the protocol, namely *Sender* and *Receiver*. There is a message sent from the *Sender* to the *Receiver*. If we denote the status of the protocol by a single variable *trans*, the (abstract) protocol can be seen in Fig. 5. In the refinement, the message is transferred via a channel between the *Sender* and the *Receiver*.

5.2 Request/Confirm pattern

The description of the protocol is as follows. There are two parties involved in the protocol, namely *Sender* and *Receiver*. The protocol contains two phases:

1. In the first phase, the *Sender* sends a request to the *Receiver*.
2. In the second phase, upon receiving the request, the *Receiver* sends a confirmation back to the *Sender*.

Using two Boolean variables *req* and *conf* to represent the state, the protocol can be illustrated as in Fig. 6. The development of this pattern used the single message communication pattern (described in Sect. 5.1) twice. These two patterns are used as illustrative examples in our earlier report [16].

5.3 Request/Confirm/Reject pattern

The description of the protocol is as follows. There are two parties involved in the protocol, namely *Sender* and *Receiver*. The protocol also contains two phases:

1. In the first phase, the *Sender* sends a request to the *Receiver*.
2. In the second phase, after receiving this request, the *Receiver* can either send a “confirmation” back to the

Fig. 6 Request/Confirm pattern

Sender if he agrees; or the *Receiver* sends a “rejection” back to the *Sender* if he does not agree.

Using three Boolean variables *req*, *conf* and *rej* to represent the state, the protocol can be seen in Fig. 7.

The development of this pattern used the single message communication pattern (described in Sect. 5.1) three times.

5.4 Asynchronous Multiple Message Communication pattern

The description of the protocol is as follows. There are two parties involved in this protocol, namely *Sender* and *Receiver*.

1. The *Sender* can send many messages (multiple message) to the *Receiver*.
2. The messages are different, in other words, there is no resend.
3. To distinguish the freshness of the message, each message is stamped with a sequence number.
4. The *Receiver* can only receive new messages.
5. The *Receiver* can discard any message.

5.5 Asynchronous Multiple Message with Repetition Communication pattern

The description of the protocol is as follows. There are two parties involved in this protocol, namely *Sender* and *Receiver*.

1. The *Sender* can send many messages (multiple message) to the *Receiver*.
2. The messages can be the same, in other words, messages could be resent.
3. To distinguish the freshness of the message, each message is stamped with a sequence number.
4. The *Receiver* can receive any message which is not old.

Fig. 7 Request/Confirm/Reject pattern

5. The *Receiver* can discard any message.

The only difference compared to the asynchronous multiple message communication (no repetition) pattern is that here messages can be resent.

6 Tool support

We have implemented our prototype for supporting our approach as a plug-in for the RODIN Platform [3] which is an open source platform based on Eclipse. The plug-in provides a wizard taking users through different steps of applying patterns, namely, *matching*, *syntax checking*, *renaming* and *incorporating*.

6.1 Matching

The tool assists developers in inputting the matching between the problem and the specification. This includes a dialog for the developers to choose the matching between variables and events. Moreover, in some cases, we need to also match the context information, i.e. carrier sets and constants which can also be chosen through the wizard page (in fact, this “matching context” is better known as generic instantiation in Event-B [4]). Information about this matching can be persistently saved for reuse later. A screen-shot of the wizard page for this step is in Fig. 8.

6.2 Syntax checking

In this step, the tool needs to check the consistency of the matching provided by the user in the previous steps. The consistency checking at this step includes:

- For events matched with some events in the pattern, we need to check the signature of these events against the corresponding pattern events.
- For remaining (unmatched) events, we need to check that they do not modify the matched variables (as mentioned earlier in Sect. 4.1).

A screen-shot of the relevant wizard page is in Fig. 9.

6.3 Renaming

The tool assists developers in inputting renaming patterns. This includes a dialog for the developers to give renaming pattern of variables and events. Consistency (e.g. name clash) for this renaming is verified at this step. A screen-shot of the renaming wizard page is in Fig. 10.

Fig. 8 First step. Matching

Fig. 9 Second step. Syntax checking

6.4 Incorporating

Finally, the tool generates the refinement of the problem according to the input in the previous steps. In order to incorporate the refinement of the pattern into the development, the tool needs to extract information from the gluing invariant on how the abstract variables v in the pattern are refined.

Fig. 10 Third step. Renaming

Fig. 11 Fifth step. Incorporating

Usually, the information is of the form $v = X(w)$. At the moment this information is also entered manually by the user in the wizard. A screen-shot of the wizard page for the incorporating step is in Fig. 11.

7 Conclusion

We have presented an approach for reusing formal models as patterns in Event-B. During a development, patterns can be discovered by either identifying the part of the model matched by existing patterns, or by recognising similar elements of the model which could be developed separately as a new pattern themselves.

Even though we presented in Sect. 4.1 a formalisation of our approach when there is only a single refinement step in the pattern development, the approach is also valid when there are multiple refinement steps. This is the same as applying patterns step by step for each level of refinement. Since refinement is monotonic, the final resulting model will be a refinement of the original model. Practically, only the last refinement model of the pattern's refinement-chain is incorporated in the development. This is already supported by our tool presented in Sect. 6. This feature allows us to reuse our formal models more flexibly, for example, using the question/response protocol in the development of the A2A Communications [10, Section 5.3.3].

7.1 Scalability

We have applied our approach to two medium-size case studies from SAP, namely the Buyer/Seller B2B [23] and Ordering/Supply Chain A2A Communications [10, Section 5.3.3]. However, our approach is general and is not restricted to this specific domain. The efforts on modelling and proving are replaced by specifying how patterns are identified and incorporated into the development. Our experiments show that this process is scalable. In particular, the patterns can be nested, i.e., a pattern can be used to develop another pattern, which then can be reused in a larger development.

So far, our patterns are quite specific since they arose from some domain-specific problems that we are trying to solve. More general patterns can be “parameterised” by some carrier sets and constants, which can be “instantiated” upon application to a problem (see our discussion on future work in Sect. 7.3). This makes the patterns more reusable in distinct problems within different contexts.

Finally, tool support is important for making our approach scalable. Our aim is to have as less interaction from the user as possible by providing different assistances for users when using the tool. Our initial experiments with the implemented tool support is encouraging.

7.2 Related work

Design patterns are well-known concepts in object-oriented programming, in particular, in the work of the Gang-of-Four (GoF) [14]. In their work, each pattern is usually represented by some informal description and some diagram in UML. There is no formal semantics associated with patterns, hence the meanings of these patterns are imprecise. There is some work on formalising these classic software design patterns in different formal notations, e.g., using predicate logic [7], using TLA+ [22], using DisCo [18]. In these papers, the first step is to give some formal meaning to the pattern before the verification of its correctness can take place. This also needs to be done for any newly defined pattern. To overcome this

problem, one needs to give some formal semantics to the diagrams used to define patterns. LePUS3 [15] is designed precisely for this purpose. However, verification in LePUS3 emphasises on the consistency between a specification (diagram) and a program. In our opinion, this is quite different from using patterns consistently to design the future system.

Our approach is related to decomposition [8,4] where developers can separate a model into sub-models and can subsequently refine these sub-models independently. The similarity with our approach is when some of the sub-models already exist as some off-the-shelf components (patterns). In this case the advantage of reusing is similar; however, decomposition is not intended for reusing.

Another related work to ours is the “automatic refinement tool” [19]. Our patterns are just formal models which encode some design decisions about refining some abstract models. However, the automatic refinement tool still requires proofs in order to make sure that the proposed refinement is correct. This approach does not necessarily preserve correctness.

Comparing with classical B [1], reusing of components is facilitated by the *INCLUDES* clause in the specification level and *IMPORTS* clause at the implementation level to compose different components. In order to reuse the same components several times, classical B supports a renaming mechanism by prefixing the name of the included/imported components with some certain identifier. In our approach, we allow the user to specify the renaming of the pattern, but it could also be done systematically with a prefixing mechanism. The main difference between our approach and the including/importing mechanism is that the including/importing mechanism does not support incorporation refinement, i.e. only reuse of the specification of the pattern is possible.

In Z [21], schemas can be reused conveniently by combining together using operators of the *schema calculus*. Moreover, *instances* of schema can be created by *schema referencing* mechanisms which include both *generic constructions* and *renaming*. Similar to classical B, this technique allows reusing of a single specification component only.

7.3 Future work

As for future work, we intend to implement the missing features from the current prototype plug-in for the RODIN Platform, e.g. syntax checking and support for extracting information from the pattern refinement. The current documentation for tool support is at the Event-B wiki documentation system [12]. At the same time, we are going to investigate more examples in other domains that could benefit from our approach.

Furthermore, we also need to “instantiate” the context of the pattern development. In our examples so far, the contexts of the pattern and the problem are the same. However, we would like to use the patterns in a more general context. For

example, the model of the communication for transferring a certain (abstract) message should be instantiated for any kind of (concrete) message, e.g., if the message is just a Boolean value, or if the message contains some numbers or some complicated data structure. This requires the context of the pattern to be instantiated accordingly. Generic instantiation [4] is a more general concept and could be used in association with other applications, for example with shared-event composition as shown in [20].

As mentioned before, it is not necessarily the case that all the variables of the pattern need to be matched with some variables in the problem. It could be the case that only a part of the variables needs to be matched or even none of them, which corresponds to the case where we do superposition refinement [4]. This makes the approach more flexible.

Moreover, we have specifically chosen to have the “syntax checking” rather than raising proof obligations when applying patterns. In the future, if this turns out to be too restrictive, we can choose to generate the corresponding proof obligations, again for more flexibility. Note that if a pattern matching can be syntactically checked successfully, the proof obligations generated should be trivial to be discharged.

Acknowledgments We would like to thank Matthias Schmalz for his comments on early drafts of the article. We also like to thank anonymous reviewers for their constructive comments to improve the quality of the article.

References

1. Abrial, J.-R.: *The B-Book: Assigning Programs to Meanings*. Cambridge University Press, Cambridge (1996)
2. Abrial, J.-R.: *Modeling in Event-B: System and Software Engineering*. Cambridge University Press, Cambridge (2010)
3. Abrial, J.-R., Butler, M., Hallerstede, S., Hoang T.S., Mehta, F., Voisin, L.: RODIN: an open toolset for modelling and reasoning in event-B. *Int. J. Softw. Tools Technol. Transf.* **12**(6), 447–466 (2010)
4. Abrial, J.-R., Hallerstede, S.: Refinement, decomposition, and instantiation of discrete models: application to event-B. *Fundam. Inf.* **77**(1–2), 1–28 (2007)
5. Abrial, J.-R., Hoang, T.S.: Using design patterns in formal methods: an event-B approach. In: Fitzgerald, J.S., Haxthausen, A.E., Yenigün, H. (eds.) *ICTAC*, Lecture Notes in Computer Science, vol. 5160, pp. 1–2. Springer, Berlin (2008)
6. Back, R.-J.: Refinement calculus II: parallel and reactive programs. In: deBakker, J.W., deRoever, W.P., Rozenberg, G. (eds.) *Stepwise refinement of distributed systems*. Lecture Notes in Computer Science, vol. 430, pp. 67–93. Springer, The Netherlands (1989)
7. Bayley, I.: Formalising design patterns in predicate logic. In: SEFM, pp. 25–36. IEEE Computer Society, USA (2007)
8. Butler, M.: Decomposition Structures for Event-B. In: *Integrated Formal Methods*. Lecture Notes in Computer Science, vol. 5423, pp. 20–38. Springer, Berlin (2009). <http://www.springerlink.com/content/3202127567642301/>
9. Decker, G., Weske, M.: Local enforceability in interaction petri nets. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) *BPM*. Lecture Notes in Computer Science, vol. 4714, pp. 305–319. Springer, Berlin (2007)
10. DEPLOY Project: Deliverable JD1—Report on Knowledge Transfer (2009). <http://www.deploy-project.eu/pdf/fv-d5-jd1-reportonknowledgetransfer.zip>
11. Fürst, A.: Design patterns in Event-B and their tool support. Master’s thesis, Department of Computer Science, ETH Zurich, March (2009). <http://e-collection.ethbib.ethz.ch/view/eth:41612>
12. Fürst, A.: Documentation on tool support for Event-B design patterns (2010). <http://wiki.event-b.org/index.php/Pattern>
13. Fürst, A., Hoang, T.S.: Rodin platform archive of question/response protocol (2010). <http://deploy-eprints.ecs.soton.ac.uk/230/>
14. Gamma, E., Helm, R., Johnson, R.E., Vlissides, J.: *Design Patterns. Elements of Reusable Object-Oriented Software*. Addison-Wesley, Reading (1995). ISBN: 10: 0201633612; 13: 978-0201633610.
15. Gasparis, E., Nicholson, J., Eden, A.H.: Lepus3: an object-oriented design description language. In: Stapleton, G., Howse, J., Lee, J. (eds.) *Diagrams*. Lecture Notes in Computer Science, vol. 5223, pp. 364–367. Springer, Berlin (2008)
16. Hoang, T.S., Fürst, A., Abrial, J.-R.: Event-B patterns and their tool support. In: Hung, D.V., Krishnan, P. (eds.) *SEFM*, pp. 210–219. IEEE Computer Society, USA (2009)
17. Lamport, L.: The temporal logic of actions. *Trans Progr Lang Syst* **16**(3), 872–923 (1994)
18. Mikkonen, T.: Formalizing design patterns. In: *ICSE*, pp. 115–124 (1998)
19. Requet, A.: BART: a tool for automatic refinement. In: Börger, E., Butler, M.J., Bowen, J.P., Boca, P. (eds.) *ABZ*. Lecture Notes in Computer Science, vol. 5238, pp. 345. Springer, Berlin (2008)
20. Silva, R., Butler, M.: Supporting reuse of event-B developments through generic instantiation. In: Breitman, K., Cavalcanti, A. (eds.) *ICFEM*. Lecture Notes in Computer Science, vol. 5885, pp. 466–484. Springer, Berlin (2009)
21. Spivey, M.: *The Z Notation: A Reference Manual*, 2nd edn. Prentice Hall International, Englewood Cliffs (1992)
22. Taibi, T., Herranz-Nieva, Á., Moreno-Navarro, J.J.: Stepwise refinement validation of design patterns formalized in TLA+ using the TLC model checker. *J. Object Technol.* **8**(2), 137–161 (2009)
23. Wiczkorek, S., Roth, A., Stefanescu, A., Charfi, A.: Precise steps for choreography modeling for SOA validation and verification. In: *Proceedings of the Fourth IEEE International Symposium on Service-Oriented System Engineering*, December (2008). <http://deploy-eprints.ecs.soton.ac.uk/41/>

Author Biographies

Thai Son Hoang is currently a senior researcher at Swiss Federal Institute of Technology, Zurich (ETH-Zurich), Switzerland. He studied undergraduate at the University of New South Wales (UNSW), Australia. He received his PhD at UNSW in 2006. He has worked initially as an academic guest, then as a post-doc and a lecturer at Department of Computer Science, ETH Zurich since 2005. He was a member of the team developing the RODIN Platform for Event-B at ETH Zurich from 2005 to 2007. His research interests and competence include formal modelling, formal verification and developing tool supports.

Andreas Fürst is a PhD student in the Information Security Group at the Department of Computer Science, ETH Zurich, where he also received his MSc degree in Computer Science. His research interests include the applicability of formal methods for software development, building secure and reliable systems and programming methodology.

Jean-Raymond Abrial is the co-inventor of Z, B and Event-B. He is the author of the “B-book” (CUP 1996), which presents the B-Method. He published recently a new book “Modeling in Event-B: System and Software Engineering” (CUP 2010). He was a guest Professor at ETH Zurich from 2004 to 2007 where he led the team developing the Rodin Platform for Event-B (funded by the European Project RODIN). After that, he was a researcher also at ETH Zurich, working on a new European Project called DEPLOY till May 2009. Before being in Zurich, he was a consultant for more than 20 years working in close contact with industrial companies but also with various universities around the world.