

Modularisation in Event-B

Alexei Iliasov, Alexander Romanovsky
Elena Troubitsyna, Linas Laibinis

Newcastle University
Åbo Academy

Contents

- ▶ About the Modularisation
- ▶ Example
- ▶ Patterns of decomposition
- ▶ Experience summary

Installation

The plugin works with the platform version 2.0 or higher.
The modularisation plugin is installed as follows:

- ▶ go to Install New Software
- ▶ in the software sites, select *Modularisation*
- ▶ check and click to install

Alternatively,

- ▶ click Add Site, the site url is
`http://iliasov.org/modplugin`
- ▶ then proceed as above

What the plugin does

- ▶ The plugin extends the Event B modelling language with the concept of a module
- ▶ A module is a parametrised Event B development associated with a module **interface**
- ▶ An interface defines a number of **operations**
- ▶ A specification is decomposed by including a module in a machine and connecting the two using operation calls and gluing invariants

What the plugin provides

- ▶ a new type of Event B component - a module interface (editor, pretty-printer and proof obligations generator)
- ▶ new machine constructs: **IMPLEMENTS** and **USES**
- ▶ new event attributes: **group** and **final**
- ▶ the ability to write operation calls in event actions
- ▶ additional proof obligations for operation calls
- ▶ additional proof obligations for implementation machines

Parking Lot

A popular parking lot requires an access control and payment collection mechanisms. The following main requirements were identified:

1. no car may enter when there is no space left in the parking lot
2. a fare must be paid when a car leaves the parking lot
3. each time a car leaves the parking lot, the fare to be paid is determined by multiplying the total length of stay since the midnight (that is, including any previous stay(s)) by the cost of parking per unit of time
4. the amount paid in any single transaction is capped
5. at midnight, the accumulated parking time of all cars is reset to zero


Parking Lot

Solution overview:

1. two gates are placed to control entry and exit
2. a payment collection machine is placed near to the exit gate in such a manner that a driver may use it before going through the exit gate
3. the exit gate does not open until the full payment is collected
4. the entrance gate does not open if the car park is full

Abstract model

the initial model describes the phenomena of cars entering and leaving the parking lot. It addresses the capacity restrictions although without exhibiting a concrete mechanism for controlling the number of cars entering the parking lot.



Model variables

- ▶ *LOT_SIZE* - the parking lot capacity (constant)
- ▶ *entered* - the number of cars that have entered the parking lot
- ▶ *left* - the number of cars that have left the parking lot
- ▶ hence, *left* - *entered* is the current number of cars in the parking lot

INVARIANT

entered $\in \mathbb{N}$

left $\in \mathbb{N}$

entered - *left* $\in 0 \dots LOT_SIZE$

Model events

a new car appears:


```
enter  =  WHEN
          entered - left < LOT_SIZE
        THEN
          entered := entered + 1
        END
```

a car leaves:

```
leave  =  WHEN
          entered - left > 0
        THEN
          left := left + 1
        END
```

First refinement

In the first refinement the entrance is controlled by a **gate**. The gate prevents a car from entering when there is no free space and also records the registration plate of an entering car.



Gate Module

The logic controlling a gate is easily decoupled from the main model. We decompose the model into the controller part and an entry gate

The first step of this decomposition is to define a gate module interface.

Gate variables

- ▶ *CAR* - car id (registration plate)
- ▶ *mcars* - the number of cars that has passed through the gate
- ▶ *current* - the id of the car in the front of the gate

INVARIANT

mcars $\in \mathbb{N}$

current $\in CAR$

Gate operations

when there is no car in front of the gate, a driver may press the gate button to try to open the gate:

```
carid  $\leftarrow$  Button    =    PRE
                           current = empty
                           POST
                           current'  $\in$  CAR  $\setminus$  {empty}
                           carid' = current
                           END
```

Gate operations

the car park controller orders the gate to open; the gate has sensors to observe whether the car has moved through the gate ($moved = \text{TRUE}$) or stayed in front of the gate:

```
moved  $\leftarrow$  OpenGate  =  PRE
                         $current \neq empty$ 
                        POST
                        ( $moved' = \text{TRUE} \wedge mcars' = mcars + 1 \wedge$ 
                           $current' = empty$ )  $\vee$ 
                        ( $moved' = \text{FALSE} \wedge mcars' = mcars \wedge$ 
                           $current' = current$ )
                        END
```

Gate operations

predicate $mcars' = mcars \wedge current' = current$ in

$$(moved' = \text{TRUE} \wedge mcars' = mcars + 1 \wedge current' = \text{empty}) \vee \\ (moved' = \text{FALSE} \wedge mcars' = mcars \wedge current' = current)$$

is necessary to indicate that *mcars* and *current* remain unchanged in the second branch of the post-condition. This is only required when a disjunction is used and not all variables are assigned new values in the disjunction branches

Operating the gate

to open the gate and let a car through it, the following has to happen:

- ▶ a driver must press the gate button (operation *Button*)
- ▶ the controller must activate the gate (operation *OpenGate*)

in our model, the main development models both driver's and controller's behaviour

First refinement machine

The first refinement imports the gate module interface. Prefix **entry** is used to avoid name clashes (with another gate added later on).

When a prefixed interface is imported, all its constants and sets appear prefixed in the importing context. This is not always convenient. We use type instantiation to replace the type of an imported module by a typing expression known in the importing context. We also define a property (an axiom) that equates a prefixed and unprefixed versions of constant *empty*.

```
USES entry : ParkingGate
TYPES
  entry_CAR  $\mapsto$  CAR
PROPERTIES
  entry_empty = empty
```

First refinement machine

two new variables are defined in the refinement machine. They help to link the states of the controller and the entry gate.

- ▶ *incar* - the id of an entering car
- ▶ *inmoved* - a flag indicating whether a car has passed through the (open) entry gate

INVARIANT

incar \in *CAR*

inmoved \in *BOOL*

Import invariant

it is necessary to provide an invariant relating the states of an imported module and the importing machine (**import invariant**)

without this, a module import does not make much sense as an overall model would be composed of two independently evolving systems

Import invariants

when there is no car at the gate, the gate car counter has the same value as the controller counter:

$$inmoved = \text{FALSE} \implies entered = entry_mcars$$

when a car is passing through the entrance gate, only the gate counter has been incremented:

$$inmoved = \text{TRUE} \implies entered + 1 = entry_mcars$$

Import invariants

when a car is passing through the gate there must be no other car at the gate:

$$inmoved = \text{TRUE} \implies entry_current = empty$$

when a car is coming through the entrance gate there is certainly free space in the parking lot:

$$inmoved = \text{TRUE} \wedge entry_current \neq empty \implies entered - left < LOT_SIZE$$

Model events

a driver presses the gate button at the entrance gate (new event):

```
UserPressButton  =  WHEN
                    entered – left < LOT_SIZE
                    entry_current = empty
                    inmoved = FALSE
                    THEN
                    incar := entry_Button
                    END
```

here **entry_Button** is a call of the *Button* operation from the *entry* module.

Model events

the parking lot controller orders the gate to open (new event):


```
CtrlOpenGate  =  WHEN
                  entry_current  $\neq$  empty  $\wedge$  inmoved = FALSE
                THEN
                  inmoved := entry_OpenGate
                END
```


Model events

finally, the *enter* event is refined to reflect the model changes:

```
enter  =  WHEN
          inmoved = TRUE
        THEN
          entered := entered + 1
          inmoved := FALSE
        END
```


Development structure



all proof obligations are discharged automatically (18 total)

Second refinement

the second refinement is very similar: we add another gate - an exit gate. the same module is imported with a new prefix to obtain two separate modules modelling two gates.



Second refinement machine

two new variables are defined :

- ▶ *outcar* - the id of the leaving car
- ▶ *outmoved* - the flag indicating whether a leaving car has passed through the (open) exit gate

INVARIANT

outcar \in *CAR*

outmoved \in *BOOL*


Import invariant

in addition to the conditions relating the variables of the exit gate module with the variables of the main machine (the controller) we are also able to specify a link between the states of the two gates

when the gates are closed, the number of cars entered through the entry gate minus the number of cars left via the exit gate may not be less than zero and is not greater than the parking lot capacity:

$$\begin{array}{l} inmoved = \text{FALSE} \wedge outmoved = \text{FALSE} \implies \\ entry_mcars - exit_mcars \in 0 \dots LOT_SIZE \end{array}$$

Development structure




one interactive proof (17 total)

Third refinement

the third refinement step is concerned with keeping the record of car stays; this step introduces the notion of time

the definition of time will be used more than once and thus it is convenient to place in an interface



Clock interface

the clock interface models the progress of time; the following is taken as the definition of time:

- ▶ time value increase is monotonic
- ▶ time changes in discrete increments when it is observed

this reflects our modelling approach to time; at an implementation stage it may have to be mapped onto a differing concept of time progress

Clock constants and variables

- ▶ *from* - the lowest time value (constant)
- ▶ *to* - the highest time value (constant)
- ▶ *delta* - the smallest observable time increment (constant)
- ▶ *prev* - the last reading of the clock (variable)

AXIOMS

$to \in \mathbb{N} \wedge from \in \mathbb{N}$
 $to - from \in delta$
 $delta > 0$

INVARIANT

$prev \in from \dots to$

Clock operations

the time progress is observed and the current time value is returned:

```
t ← currentTime  =  PRE
                    TRUE
                    POST
                     $prev' \in from \dots to$ 
                     $prev' \geq prev + delta \wedge prev' = to \wedge t' = prev'$ 
                    END
```

the clock is reset:

```
clockReset  =  PRE
               TRUE
               POST
                $prev' = from$ 
               END
```

Third refinement machine

constant definitions:

- ▶ *TOD* - the time-of-the-day type
- ▶ *DAY_START* - day start time value
- ▶ *DAY_END* - day end time value

AXIOMS

$$TOD = DAY_START \dots DAY_END$$
$$DAY_START \in \mathbb{N}$$
$$DAY_END \in \mathbb{N}$$
$$DAY_END > DAY_START$$

Third refinement machine

new variables:

- ▶ *register* - function recording the time when a car enters the parking lot
- ▶ *cartime* - for a given car gives the accumulated stay time since the midnight

INVARIANT

$register : CAR \mapsto TOD$
 $cartime : CAR \rightarrow \mathbb{N}$

INITIALISATION

$register := \emptyset$
 $cartime := CAR \times \{0\}$

the time spent in the parking lot since the midnight is no greater than the latest registration time:

$$\forall x \cdot x \in \text{dom}(\text{register}) \implies \text{register}(x) - \text{DAY_START} \geq \text{cartime}(x)$$

Clock module import

the clock module is imported without a prefix; the time limits are set to correspond to the *TOD* data type:

```
USES Clock
  PROPERTIES
    from = DAY_START
    to = DAY_END
```

Import invariants

all the car registration timestamps have the time value not exceeding the current time:

$$\forall x \cdot x \in \text{dom}(\text{register}) \implies \text{register}(x) \leq \text{prev}$$

the time a car has spent in the park since the midnight is not more than the time elapsed since the midnight:

$$\forall x \cdot x \in \text{CAR} \implies \text{cartime}(x) \leq \text{prev} - \text{DAY_START}$$

Model events

a record is made of the time when a car enters the car park:

```
enter  =  EXTENDS enter
        BEGIN
            register(incar) := currentTime
        END
```

currentTime is an operation call returning (and also advancing) the current time

Model events

when a car leaves, the registration record is removed and the total stay time is updated:

```
CtrlOpenGateL  =  EXTENDS CtrlOpenGateL
                  WHEN
                     $outcar \in dom(register)$ 
                  THEN
                     $cartime(outcar) := cartime(outcar) +$ 
                       $(currentTime - register(outcar))$ 
                     $register := outcar \triangleleft register$ 
                  END
```

here $currentTime - register(outcar)$ is the length of the current stay of the leaving car $outcar$

Model events

according to the requirements, upon midnight, the car stay times and registration timestamps are reset:

```
RegisterReset  =  WHEN
                  prev = DAY_END
                THEN
                  clockReset
                  register := dom(register) × {DAY_START}
                  cartime := CAR × {0}
                END
```

operation **clockReset** sets the clock reading *prev* to the start of a day time value *DAY_START*


Model events

to make sure that clock and register resets happen even when there are no cars entering or leaving the parking lot, the system must actively observe time

ObserveTime	=	BEGIN
		currentTime
		END

this event forces the progress of time even if no other time-related activity takes place

Development structure




five interactive proofs (30 total)

Fourth refinement

in this step, before a car may leave, the car driver must pay the amount determined by the length of stay since the midnight

the functionality of a device collecting payment is decoupled from the controller logic and is placed in a separate module



Payment machine constants and variables

- ▶ *payPerTimeUnit* - the cost of unit of time in the parking lot (constant)
- ▶ *maxPay* - the limit on the amount paid in a single transaction (constant)
- ▶ *balance* - the outstanding balance to be paid (variable)

AXIOMS

payPerTimeUnit $\in \mathbb{N}$

maxPay $\in \text{nat}$

INVARIANT

balance $\in \mathbb{N}$

Payment machine operations

the payment machine is configured by supplying the accumulated length of stay as a parameter; the operation computes the balance to be paid:

```
Configure  =  ANY stay PRE
               stay ∈ ℕ
               balance = 0
            POST
               balance' = min(staypayPerTimeUnit, maxPay)
            END
```

the payment is taken from a driver; the amount paid is at least as large as the outstanding balance (i.e., a driver may overpay but not underpay):

```
Pay  =  PRE
        balance ≠ 0
      POST
        p' ∈ ℕ ∧ p' ≥ balance ∧ balance' = 0
      END
```

Fourth refinement machine

there are two new variables used to constrain the ordering of concrete events:

- ▶ *confPay* - a flag indicating the configuration phase of payment collection
- ▶ *paid* - a flag indicating that payment has been collected

INVARIANT

confPay \in *BOOL*

cpayed \in *BOOL*

Import invariants

when there is no car at the exit gate the pay machine balance is zero:

$$exit_current = empty \vee confPay = \text{TRUE} \implies pm_balance = 0$$

during payment configuration there is always a car at the exit gate:

$$confPay = \text{TRUE} \implies exit_current \neq empty$$

Model events

the controller configures the payment machine by calling the *Configure* operation with the accumulated stay time:


```
CtrlPay  =  WHEN
            confPay = TRUE
        THEN
            pm_void := pm_Configure(cartime(outcar))
            confPay := FALSE
        END
```

Model events

a driver pays if there is an outstanding balance to be paid (always the case with the current payment machine and clock interfaces):

```
UserPay  =  WHEN
              ...  $\wedge pm\_balance > 0$ 
            THEN
              payed := TRUE
              pm_Pay
            END
UserNoPay =  WHEN
              ...  $\wedge pm\_balance = 0$ 
            THEN
              payed := TRUE
            END
```

Development structure



all proof obligations are discharged automatically (34 total)

Implementing Modules

The development relies on three modules that are so far defined only by their interfaces. To complete the development, we will construct developments corresponding to these interfaces. One exception is the Clock interface that represents a simple time theory and cannot be usefully detailed in a module body

Implementing Modules

A machine providing the realisation of an interface is said to *implement* the interface. This is recorded by adding the interfaces into the **IMPLEMENTS** section of a machine. The fact that a machine provides a correct implementation of interfaces is established by a number of static checks and a set of proof obligations. The latter appear automatically in the list of machine proof obligations. The implementation relation is maintained during machine refinement (subject to some syntactic constraints) and thus the bulk of the module implementation activity is the normal Event B refinement process.

Event Group

The first step of implementing an interface is to provide at least one event for each interface operation. In general, an operation is realised by a set of events (an event **group**). Some events play a special role of operation termination events and are called **final** events. A final event returns the control to a caller. It must satisfy the operation post-conditions but there is no need to prove the convergence of a final event.

Implementing **ParkingGate**: abstract machine

To simplify proofs, the initial implementation is a simple machine with few events mirroring the interface operations. The machine retains interface variables *current* and *mcars* and also defines the operation return variables *Button_carid* and *OpenGate_moved*.

The names of the operation return variables are fixed for the first machine of a module implementation. In further refinements they may be replaced or removed using data refinement.

Implementing **ParkingGate**: abstract machine

The *button* event implements operation *Button* in a single atomic step. The fact that it is associated with operation *Button* is stated by GROUP Button. Being the only event in its operation group it is also a FINAL event.


```
MACHINE iParkingGate IMPLEMENTS
VARIABLES current mcars Button_carid OpenGate_moved
EVENTS
  button  =  FINAL GROUP Button
              WHEN
                current = empty
              THEN
                current :∈ CAR \ {empty}
                Button_carid := current
              END
```


Implementing **ParkingGate**: abstract machine

The machine declares two more events, both realising the *OpenGate* operation. The events are final and each one handles one of the cases of the *OpenGate* operation post-condition.

```
gate_succ  =  FINAL GROUP OpenGate
              WHEN
                  current  $\neq$  empty
              THEN
                  OpenGate_moved := TRUE
                  mcars := mcars + 1
                  current := empty
              END
gate_nocar  =  FINAL GROUP OpenGate
              WHEN
                  current  $\neq$  empty
              THEN
                  OpenGate_moved := FALSE
              END
```

Development structure



one interactive proof (5 total)

Implementing **ParkingGate**: first refinement

new variables:

- ▶ *gate* - the gate state: open or closed
- ▶ *sensor* - the state of the car sensor placed; the sensor is placed on the parking lot of a gate
- ▶ *stage* - the current step of the gate operation

INVARIANT

gate \in *GATE*

sensor \in *BOOL*

stage $\in 0 \dots 3$

stage = 1 \implies *gate* = *OPEN*

stage = 2 \implies *gate* = *CLOSED*

Implementing **ParkingGate**: first refinement

The refined implementation of the *OpenGate* operation includes events for opening and closing the gate.


```
open_gate  =  GROUP OpenGate
              WHEN
                  gate = CLOSED  $\wedge$  stage = 0
              THEN
                  gate := OPEN
                  stage := 1
              END
close_gate  =  GROUP OpenGate
              WHEN
                  stage = 2
              THEN
                  gate := CLOSED
                  stage := 3
              END
```

Implementing **ParkingGate**: first refinement

the gate detects whether a car has passed through the gate while the gate was open:

```
readSensor  =  GROUP OpenGate
              WHEN
                stage = 1
              THEN
                sensor :∈ BOOL
                stage :∈ 1, 2
              END
```

Development structure



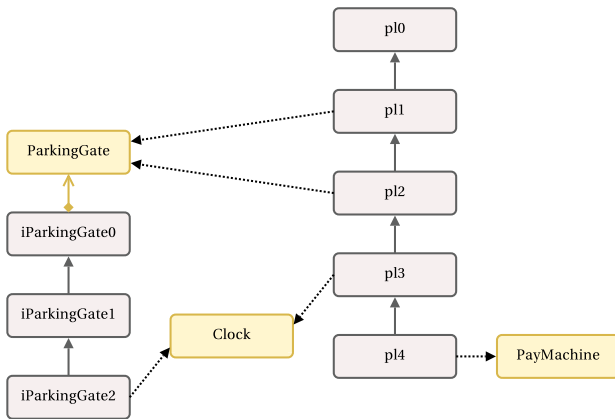
all proof obligations are discharged automatically (26 total)

Implementing **ParkingGate**: second refinement

To prove the convergence of anticipated event *readSensor*, the car sensor waits for a car for a given time interval. The time model is imported from the *Clock* interface.

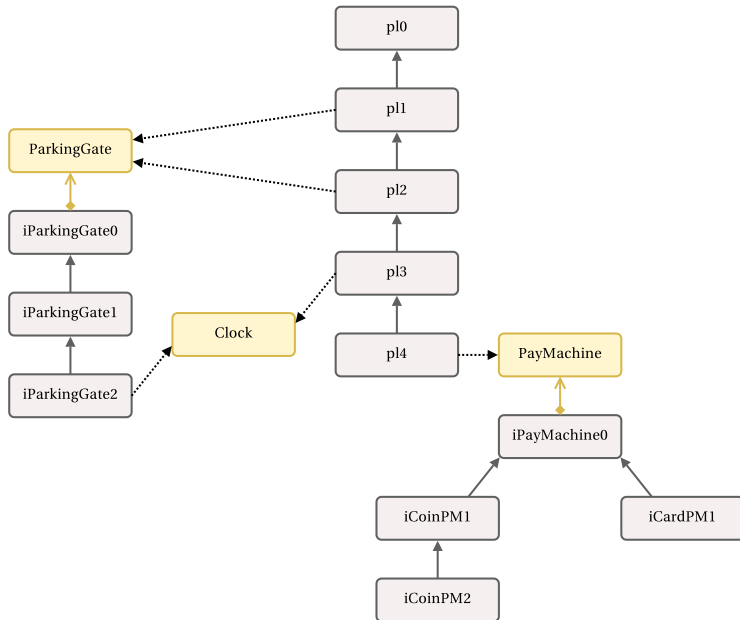
```
readSensor  =  GROUP OpenGate
              WHEN
                stage = 1
                prev < delay
              THEN
                sensor, stage :| (sensor' = TRUE ∧ stage' = 2) ∨
                                (sensor' = FALSE ∧ stage' = 1)
                time := currentTime
              END
```

Development structure




two interactive proofs (8 total)

The overall development structure



Decomposition Patterns

The purpose of the patterns is to facilitate a specific form of model decomposition where the core functionality of an abstraction is distributed among two or more separate modules connected by a relatively primitive middleware model. Typically, the modules correspond to the software or hardware being developed while the middleware model should match an existing coordination infrastructure.



Component template


A generic component description is given by an interface. The interface variables are seen by the middleware component and can be used to formulate a gluing invariant.

The integration is achieved by the means of two operations: one for informing the component about a new incoming message and another confirming the processing of an outgoing message.

The production of new messages is modelled as an independent thread of control within a component. At the interface level the thread manifests itself by adding new messaging into output message queue.


Component template variables

Interface defines message sequences recording all the incoming and outgoing messages during the component lifetime. The output message history also plays the role of output message queue.



In the above and further, light green are messages already processed by the middleware and are either delivered or in transmission. Dark green slots correspond to fresh messages not yet processed by the middleware. Blue slots are messages originated at another component and delivered by the middleware.

Component template variables



$ih \in 1 \dots il \rightarrow MSG$

input history sequence

$oh \in 1 \dots ol \rightarrow MSG$

output history sequence

$ol \geq r \wedge olr \leq QUEUE_LENGTH$

output queue capacity

Component template operations

The *receive* operation reacts on an incoming. At the interface level the observed effect is a new message in the input history.

```
receive  =  ANY  $m$  PRE  
            $m \in MSG$   
           POST  
              $ih' = ih \cup \{il + 1 \mapsto m\}$   
              $il' = il + 1$   
           END
```

An actual model would define further pre- and post-conditions.

Component template operations

Operation *deliver* marks a message in the output queue as processed. It also frees one slot in the output message queue.

$m \leftarrow \text{deliver}$	=	PRE
		$ol > r$
		POST
		$m' = oh(r + 1)$
		$r' = r + 1$
		END

When the operation is called it returns a message to be delivered.

Component template process


A component has a thread of control that allows it accomplish some tasks independently of middleware and other components. One observed effect of the thread execution is the generation of new messages in the output queue.

```
proc    =    GUARANTEE
            $ol' \geq ol \wedge ol' r \leq \text{QUEUE\_LENGTH}$ 
            $oh = 1 \dots ol \triangleleft oh'$ 
           END
```

The process guarantee states that the process may change the output queue by adding new messages to the queue tail until the capacity limit is reached.

Synchronous Coordination

The templates defines a middleware transmitting a message between two components in one atomic step. The middleware has no memory (hence no model variables of its own) and acts as a wiring logic relating two modules.




It is convenient to represent bi-directional channels as a pair of uni-directional ones.

Synchronous Coordination

When two components are linked by a synchronous template the following properties are maintained:


$$left_ih = 1 \dots right_r \triangleleft right_oh$$

$$left_il = right_r$$




Synchronous Coordination

Event *copy_left_right* transfers one message from component *left* to component *right* in an atomic step.



Asynchronous Coordination


With the asynchronous communication a middleware is equipped with some memory used to temporarily save messages travelling across two components. This permits one to relax an assumption that a message is received the moment it is sent.



For each channel there are three message classes: messages sent and received, messages in transmission and messages yet to be transmitted.

Asynchronous Coordination

A buffer stores the messages currently in transmission. Its content is a projection of the messages generated by a sender.



$$\forall i \cdot i \in lr_read \dots left_r - 1 \implies$$
$$lr_buffer(i \bmod LR_BUFFER_SIZE + 1) = left_oh(i + 1)$$

Asynchronous Coordination


Buffer is a FIFO realised as a circular buffer (should have used an infinite sequence to simplify the proofs!).

lr_buffer is a buffer is attached to the *left* to *right* channel. Its write position is $lr_write \bmod LR_BUFFER_SIZE + 1$ and read position is $lr_read \bmod LR_BUFFER_SIZE + 1$. To maintain a lossless buffer we require that $lr_read \leq lr_write$.

The number of messages received on a channel equals the number of messages read from the channel buffer: $right_il = lr_read$ and $left_il = rl_read$


Asynchronous Coordination

Initially, a message appears in the output queue (and also the output history of a component). Such message is generated by an internal process of a component.



Asynchronous Coordination


The middleware copies the message in a channel buffer



right_ih (empty)


Asynchronous Coordination

The message is delivered to the receiver




Asynchronous Coordination

The sender generates another message




Asynchronous Coordination

The message is copied by the middleware into the buffer



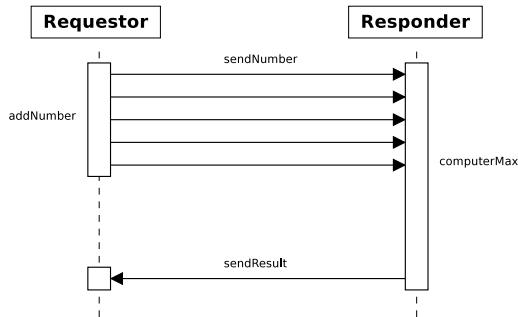
Asynchronous Coordination

The sender generates another message



Example: MaxOf5

As an example we consider the specification of a simple protocol following the typical request/response pattern.



The requester sends five distinct numbers to the responder which, as reply, sends a message with the maximum of the numbers.

Example: MaxOf5

The complexity in the example is not in the nature or the properties of the protocol but rather in proving the refinement step leading from an abstract protocol specification to a decomposed (distributed/concurrent) model.

The development in this example is built as around the synchronous communication template.

Abstract Model

- ▶ *set* - set of generated numbers (requests)
- ▶ *result* - the computed result variable

$$\begin{array}{l} \textit{set} \subseteq \mathbb{N} \\ \textit{finite}(\textit{set}) \\ \textit{card}(\textit{set}) \leq 5 \\ \textit{result} \in \mathbb{Z} \end{array}$$

Abstract Model

generates a new number until there are total 5 numbers in the set

```
addNumber = ANY  $n$  WHERE  
             $n \in \mathbb{N}$   
             $n \notin set$   
             $card(set) < 5$   
        THEN  
             $set := set \cup \{n\}$   
        END
```


Abstract Model

computes the maximum of the generated number set and saves the result

```
computeMax = WHEN
               card(set) = 5
               result = 1
            THEN
               result := max(set)
            END
```

Concrete Model: Import Invariants

own variables: **none**

set is replaced by the contents of the output history of the generator component:

$$set = ran(g_oh)$$

result is replaced by the first entry of the processor component output history:

$$\begin{array}{l} p_ol > 0 \implies result = p_oh(1) \\ p_ol = 0 \implies result = 1 \end{array}$$

Concrete Model: Import Invariants

Connecting the communication histories of the components

$$\begin{aligned}p_{ih} &= 1 \dots g_r \triangleleft g_{oh} \\ p_{il} &= g_r \\ g_{ih} &= 1 \dots p_r \triangleleft p_{oh} \\ g_{il} &= p_r\end{aligned}$$

Protocol property: when the responder has received five messages the requestor has nothing more left to send

$$\text{card}(\text{ran}(p_{ih})) = 5 \implies g_r = g_{ol}$$

Concrete Model (Middleware)

an unbuffered channel connecting requestor and responder (one way only) this event takes a message from the output history and places it in the input history of another component

```
copy_gp  =  WHEN
             $g_{ol} > g_r$ 
        THEN
            p_receive(g_deliver)
        END
```

Concrete Model (Middleware)

an unbuffered channel connecting responder and requestor (one way only)

```
copy_pg  =  WHEN
             $p_{ol} > p_r$ 
        THEN
            g_receive(p_deliver)
        END
```

Module process definitions

an unbuffered channel connecting responder and requestor (one way only)

```
generate = WHEN
             $n \in \mathbb{N} \setminus \text{ran}(oh)$ 
             $\text{card}(oh) < 5$ 
        THEN
             $oh(ol + 1) := n$ 
             $ol := ol + 1$ 
        END
```

```
compute = WHEN
             $\text{card}(\text{ran}(ih)) = 5$ 
             $ol = 0$ 
        THEN
             $oh(ol + 1) := \max(\text{ran}(ih))$ 
             $ol := 1$ 
        END
```

Summary and Future Work

- ▶ working on several other patterns: circular buffer, lossy buffer
- ▶ maxof5 examples for all the patterns
- ▶ longer-term: distributed AOCS

Experience

- ▶ AOCS mode consistency: NCL/Abo (1400 POs)
- ▶ AOCS messaging model: (900 POs)
- ▶ SAP protocol model: (350 POs)

Tool evolution (until 1 Oct):

- ▶ 4 releases
- ▶ 22 bug reports (all fixed)
- ▶ 12 feature requests (10 implemented)
- ▶ >200 installations (45 for the latest version)

Experience

Initial modelling attempts by industrial partners were unsuccessful

- ▶ not using an implementation machine
- ▶ introducing modules too early
- ▶ too complex pre- and post-conditions
- ▶ misunderstanding of the purpose of operations

Half day training sessions were sufficient for further independent work

Some common mistakes

Modularisation should not be used to compose models

- ▶ composition is a (degenerate) form of decomposition refinement where no abstraction is given for the new functionality introduced by a module
- ▶ results in a fragmented model
- ▶ not a part of the refinement methodology
- ▶ it is a poor specification technique: inflates a model without generating any interesting (related to the modelled problem) proof obligations

Some common mistakes

It is mandatory to supply an import invariant

- ▶ import invariant is a form of a gluing invariant connecting the exported variables of a module with the variables of the importing machine
- ▶ only abstract and preserved variables of the parent machine may be used in import invariant
- ▶ the good technique is to try to replace some variables of parent machine with the exported variables
- ▶ this allows to refine actions of the abstract parent machine with operation calls

Semi-automatic decomposition

A tool that mechanises a specific decomposition approach using in the modelling of communication protocols.

- ▶ define number of modules and their names
- ▶ define variable distribution
- ▶ define operations