A refinement-based development of a distributed signalling system

Paulius Stankaitis¹, Alexei Iliasov⁴, Tsutomu Kobayashi², Yamine Aït-Ameur³,
Fuyuki Ishikawa², Alexander Romanovsky¹

¹ School of Computing, Newcastle University, Newcastle upon Tyne, United Kingdom

² National Institute of Informatics, Tokyo, Japan

³ INPT-ENSEEIH, 2 Rue Charles Camichel, Toulouse, France

⁴ The Formal Route Limited, 32A Woodhouse Grove, E12 6SR, London, United Kingdom

Abstract. The decentralised railway signalling systems have a potential to increase capacity, availability and reduce maintenance costs of railway networks. However, given the safety-critical nature of railway signalling and the complexity of novel distributed signalling solutions, their safety should be guaranteed by using thorough system validation methods. To achieve such a high-level of safety assurance of these complex signalling systems, scenario-based testing methods are far from being sufficient despite that they are still widely used in the industry. Formal verification is an alternative approach which provides a rigorous approach to verifying complex systems and has been successfully used in the railway domain. Despite the successes, little work has been done in applying formal methods for distributed railway systems. In our research we are working towards a multifaceted formal development methodology of complex railway signalling systems. The methodology is based on the Event-B modelling language which provides an expressive modelling language, a stepwise development and a proof-based model verification. In this paper, we present the application of the methodology for the development and verification of a distributed protocol for reservation of railway sections. The main challenge of this work is developing a distributed protocol which ensures safety and liveness of the distributed railway system when message delays are allowed in the model.

Keywords: Distributed railway signalling, Distributed resource allocation, Event-B method

1. Introduction

Railway signalling is a safety-critical system whose responsibility is to guarantee a safe and efficient operation of railway networks. In the last few decades there have been proposals to utilise distributed system concepts in the railway signalling as a way of increasing railway network capacity and reducing maintenance costs (e.g. [HP00, WK12]). These emerging distributed railway signalling concepts propose to use a radio-based communication technology and novel distributed signalling protocols for decentralising contemporaneous signalling systems.

Because of their complex concurrent behaviour, distributed systems are notoriously difficult to validate and this could curtail the development, and deployment of novel distributed signalling solutions. To achieve a high-level of safety assurance of novel distributed signalling systems popular railway system validation methods, such as scenario-based testing methods, would be far from being sufficient. In recent years, there has been a push by industries with a strong focus on distributed systems to incorporate formal methods into their system development processes [HHK+17, New14]. The railway domain has proved to be a fruitful area for applying various formal methods, but considerably less has been done in applying them for distributed railway systems by industry and academia [BBFM99, ED06]. Therefore, the long-term aim of our research is to lower the barriers of applying formal methods for the development of complex railway signalling systems, including distributed [SIK+20], heterogeneous [SI17] and hybrid [SDS+19] railway systems.

In order to manage the modelling and verification complexity of complex railway systems we are working towards a multifaceted formal development methodology. The methodology primarily relies on three concepts: a refinement-based model development, communication modelling patterns [SIA+19] and a proof-based model verification. The methodology is built upon the Event-B method [Abr13] which provides an expressive modelling language, flexible refinement mechanism and proof-based verification. It is also paramount that the methodology should support quantitative evaluation; as stated by Fantechi and Haxthausen [FH18], signalling solutions will only be adopted in practice if system availability is demonstrated. Therefore, in our proposed multifaceted methodology we also integrate a stochastic simulator for a quantitative railway system analysis.

In this paper, we present our work on a refinement-based formal development of a novel distributed signalling protocol. The paper attempts to demonstrate the advantages of the proposed methodology and refinement-based formal development, particularly, in the early stages of the system development (system prototyping) where specifications and requirements are incomplete. The main objective of the developed protocol was to guarantee a safe allocation of railway sections while ensuring liveness of the distributed signalling system. In our work, the liveness property of the system had to be guaranteed by the protocol in a model where message delays are permitted. The authors of related researches did not consider liveness and fairness properties which, as discussed in [FH18], can directly affect system availability. The consideration of message delays has also introduced modelling and deductive verification challenges as subtle deadlock scenarios were discovered.

Related work. In [FH18] the authors formalise the railway interlocking problem as a distributed mutual exclusion problem and discuss the related literature on distributed interlocking (e.g. [HP00, FHN17, WK12]). In principle all railway models share similar high-level safety, liveness and fairness requirements, as summarised on page 2 in [FH18]. One difference between our work and the studies overviewed in [FH18] is the interlocking engineering concept and the system model (e.g. allowed message delays). Another difference is the formal consideration of liveness and fairness requirements. In our work we not only prove the safety properties of the protocol, but also ensure systems liveness, fairness and analyse performance. A similar distributed signalling concept is presented as a case study in [Pro16]. The authors verified their system design via a simulation approach and only considered scenarios with up to two trains. In our verification approach we prove the distributed signalling system mathematically and hence guarantee its safety for any number of trains. In the study by Morley [Mor96] the author formally proved a distributed protocol, which is used in the real-world railway signalling systems to reserve a route, which is jointly controlled by adjacent signalling systems. Even though the distributed signalling concepts of our works are different, the effects of message delays to the safety were considered in both works. Our developed distributed signalling protocol is also based on serialisability and similar to ones defined from transactions processing [EGLT76, BSR80, GR92] in centralised and distributed database systems.

Paper structure. The rest of the paper is organised in the following way. Section 2 revisits a previously proposed multifaceted methodology and provides background information on the Event-B formal specification language. In Sect. 3 we elicit high-level distributed system requirements and specifications as well as provide a semi-formal description of the distributed resource allocation protocol. In Sect. 4 we present the formal protocol model development with the Event-B method. In the last section, Sect. 5, we summarise our work and discuss future work directions.

Fig. 1. Multifaceted modelling and verification framework (with focus of the paper highlighted)

$\text{event_label} \triangleq [\text{specification}_{id}]$
ANY
 parameters
WHEN
 guards
THEN
 actions

INVARIANTS
 $\text{invariant_label} \quad \text{invariant} \quad [\text{requirement}_{id}]$

Fig. 2. An example of annotated Event-B model event and invariant in the Rodin platform

2. Multifaceted formal development methodology

The generally accepted approach to any formal system development is starting a development process by first informally describing system specifications and requirements. Therefore, this is the first step (Step 1 in Fig. 1) in our formal development process. At the moment we do not suggest or provide a systematic approach to defining requirements and specifications of a distributed system, however, we require labelling each requirement with a unique label. By labelling requirement and specifications we intend to ensure the traceability between informal and formal methodology artefacts. The traceability aspect guarantees the completeness of the model, meaning, that all informal specifications and requirements have been captured by the formal model. On the formal Event-B artefact, we require to annotate events and invariants of the Event-B model with labelled references to a specific informal artefact (see Fig. 2).

The following step (Step 2) in our methodology is a development and verification of a functional formal Event-B model. The purpose of formally modelling a distributed system is to have a formal artefact, which can be animated, analysed and formally verified. For the development and verification of functional system models we selected the Event-B [Abr13] specification language, which has been successfully used for a formal development of various distributed protocols (e.g. [CM06, HKBA09, ILTR11]). The Event-B method provides an expressive modelling language, flexible refinement mechanism and is also proof-driven, meaning that model correctness is demonstrated by generating and discharging proof obligations with available automated theorem provers [DFGV14, ISAYR16]. The method is also supported by tools such as ProB [LB03] which makes it possible to animate a model and validate it via model-checking. ProB can be particularly useful in early modelling stages where it could be too onerous to deductively verify a model. On the other hand, the Event-B method does not have an adequate probabilistic reasoning support, which was essential to verifying our distributed railway section reservation protocol. Therefore, it was decided to integrate a well-known stochastic model checker, PRISM [HKNP06], into the framework, so the probabilistic properties of a system can be verified.

The last step (Step 3) in the proposed formal development process is analysing performance of the developed distributed system. For analysing system's performance we have implemented a high-fidelity protocol simulator which could help to evaluate a protocol under normal or *stressed* conditions. Because of the paper focus on Steps 1 and 2, in the following sections we provide more detail on the functional Event-B model development and the Event-B method.

Fig. 3. Event-B model structure

Table 1. Event-B proof obligations

Theorems	$A(s, c) \Rightarrow T_c(s, c)$
Invariant Preservation	$A(s, c) \wedge I(s, c, v) \wedge G(s, c, v, x) \wedge BA(s, c, v, x, v') \Rightarrow I(s, c, v')$
Event Feasibility	$A(s, c) \wedge I(s, c, v) \wedge G(s, c, v, x) \Rightarrow \exists v' \cdot BA(s, c, v, x, v')$
Variant Progress	$A(s, c) \wedge I(s, c, v) \wedge G(s, c, v, x) \wedge BA(s, c, v, x, v') \Rightarrow V(s, c, v') < V(s, c, v)$

2.1. Event-B

The Event-B mathematical language [Abr13] used in the system development and analysis is an evolution of the classical B method [Abr96] and Action Systems [Bac90]. The formal specification language offers a fairly high-level mathematical language based on a first-order logic and Zermelo-Fraenkel set theory as well as an economical yet expressive modelling notation. The formalism belongs to a family of state-based modelling languages where a state of a discrete system is simply a collection of variables and constants whereas the transition is a guarded variable transformation.

A cornerstone of the Event-B method is the step-wise development that facilitates a gradual design of a system implementation through a number of correctness preserving refinement steps. The model development starts with a creation of a very abstract specification and the model is completed when all requirements and specifications are covered. The Event-B model is made of two key components-machines and contexts which respectively describe dynamic and static parts of the system (see Fig. 3). The context contains modeller declared constants c and associated axioms $A(s, c)$ which can be made visible in machines. The dynamic part of the model contains variables v which are constrained by invariants $I(s, c, v)$ and initialised by an action. The state variables are then transformed by actions which are part of events and the modeler may use predicate guards $G(s, c, v, x)$ to denote when event is triggered.

The Event-B method is a proof driven specification language where model correctness is demonstrated by generating and discharging proof obligations-theorems in the first-order logic. Table 1 shows a few of more important proof obligations of the Event-B language. The model is considered to be correct when all proof obligations are discharged.

Rodin [The06] is an open source Eclipse-based integrated development environment (IDE) for Event-B model development. The Rodin is a core set of plug-ins for project management, formal development, syntactic analysis, proof assistance and proof-based verification.

Fig. 4. Protocol modelling patterns

2.2. Developing functional protocol models in Event-B

A formal functional Event-B model can have a multitude of uses, but the main application is for formally proving properties about the distributed system. The completed distributed system's model in Step 2 should cover all requirements and specifications, and would be considered correct when all generated proof obligations are proved.

The model development approach we utilise (visualised in Fig. 4) is a rather standard and starts with the abstract model $m_{ab.}$, which formally specifies the objective of the distributed protocol. In fact, distributed aspects of the system are ignored at this model level and the abstract model considers a centralised configuration. The following group of refinement steps ($m_{0..n}$) introduce more details about the model by primarily modelling communication aspects. For protocol modelling we propose to use backward unfolding style where the next refinement step introduces the preceding protocol step. The communication modelling patterns are primarily made of patterns for defining communicating actors of a system, introducing new messages into a model through context models (and machine variables) and machine event patterns for message exchanges. The previously developed communication modelling patterns and a generic model refinement plan are described in more detail in [SIA+19].

A key aspect of our methodology is the scenario validation and analysis. Particularly, in early protocol development stages, it might be too onerous to verify a model only to discover design mistakes. To facilitate design exploration we apply animation and model-checking enabled by ProB. Nonetheless, the final (concrete) model should be proved by adding invariants to the model and proving generated proof obligations with available automated theorem provers.

Fig. 5. An example of a distributed railway system and its abstract model: sub-route mapping to resources $SR_0 \rightarrow r_0, \dots, SR_3 \rightarrow r_3$ and trains to agents mapping $T_0 \rightarrow a_0, T_1 \rightarrow a_1$. The bold arrows illustrate resources locked by agents while dashed arrows show resources which will need to be locked by an agent. For example, if train T_1 wants to travel to sub-route SR_2 it must simultaneously lock sub-routes SR_1 and SR_2 .

Table 2. High-level distributed signalling system specifications

SYS_1	The distributed railway system is made of agents and resources (respectively represent trains and railway subsections).
SYS_2	Agents are only allowed to communicate with resources and not other agents.
SYS_3	Agents have an objective, which is a set of resources agents have to reserve before proceeding with the next objective.
SYS_4	Agents have a memory in which sent messages are stored.
SYS_5	Resources have a memory where agents allocation order can be stored.
SYS_6	Resources have a promised pointer (ppt) of the memory, which indicates a currently available resource allocation order.
SYS_7	Resources have a read pointer (rpt) of the memory, which specifies a memory slot (with an associated agent) that currently uses a resource.
ENV_1	Message delays are permitted in the distributed signalling system model.

3. High-level distributed signalling system model

We abstract the railway model and instead of trains, routes, and switches our system model consists of agents and resources (resources controllers). The system model permits message exchanges only between agents and resources, and messages can be delayed. Each resource controller has an associated queue-like memory, where the order of agent allocation can be stored. A resource also has a promise (ppt) and read pointers (rpt), which respectively indicate the currently available slot in the queue and the reserved slot (with an associated agent) that currently uses the resource (abstract railway model visualised in Fig. 5). An agent has an objective, which is a collection of resources an agent will attempt to reserve (all at the same time) before using and eventually releasing them. In Fig. 6 we visualise a simple resource locking protocol and illustrate how ppt and rpt variables are updated. In Table 2 high-level distributed system specifications are elicited in the format, which is suggested by the proposed methodology.

The main objective of the protocol is to enable safe and deadlock-free distributed atomic reservation of collection of resources. Where by a safe resource reservation we mean that no two different agents have reserved the same resource at the same time. The protocol must also guarantee that each agent eventually gets all requested resources-partial request satisfaction is not permitted. These are fundamental requirements of railway signalling systems which ensure a safe train separation by utilising a so-called fixed-block signalling [FH18].

Table 3. High-level system safety and liveness requirements

SAF ₁	A resource will not be allocated to different agents at the same time.
SAF ₂	An agent will not use a resource until all requested resources are allocated.
LIV ₁	An agent must be eventually allocated requested set of resources.
LIV ₂	Resource allocation must be guaranteed in the presence of message delays.

Fig. 6. A scenario (left to right) depicting a simple resource locking protocol with promised (ppt) and read (rpt) variable value changes.

The main high-level safety and liveness requirements of the distributed system are expressed in Table 3. The following section attempts to justify the need for an adequate distributed protocol by discussing problematic distributed resource allocation scenarios.

3.1. Problematic distributed resource allocation scenarios

Let us consider Scenarios 1-2 (visualised in Fig. 7) to see how requirement LIV₁ could not be guaranteed (while ensuring SAF₂) without an adequate distributed resource allocation protocol.

Scenario 1 In this scenario, agents a_0 and a_1 are attempting to reserve the same set of resources $\{r_0, r_1\}$. Agents start by firstly sending request messages to both resources. Once a resource receives a request message, it replies with the current value of the promised pointer $ppt(r_k)$ and then increments the $ppt(r_k)$. For instance, in this scenario, resource r_0 firstly received a request message from agent a_0 and thus replied with the value $ppt(r_0) = 0$, which was then followed by a message to a_1 with an incremented $ppt(r_0)$ value of 1. In Fig. 7, we denote a_n^* as the $ppt(r_k)$ value sent to a_n . Request messages at resource r_1 have been received and replied in the opposite order.

In this preliminary protocol, once an agent receives promised pointer values from all requested resources, it sends messages to requested resources to lock them at the promised queue-slot. In this scenario, agent a_0 was promised queue-slots $\{(r_0, 0), (r_1, 1)\}$ while a_1 queue-slots $\{(r_0, 1), (r_1, 0)\}$. If agents would lock these exact queue-slots, resource r_0 would allow a_0 to use it first, while r_1 would concurrently allow a_1 . The distributed system would deadlock and fail to satisfy LIV₂ requirement as both agents would wait for the second *use* message to ensure SAF₂.

In order to prevent the cross-blocking type of deadlocks, an agent should repeatedly re-request the same set of resources (and not lock them) until all received promised queue slot values are the same. We define a process of an agent attempting to receive the same promised queue slots as an agent forming a distributed lane (dl).

Definition 1 A distributed lane dl of an agent a is a set of pairs $dl(a) = \{(r_k, s), \dots, (r_{k+m}, s)\}$ where $\{r_k, \dots, r_{k+m}\}$ are all resources requested by an agent a and s is a natural number representing a slot value of the queue promised by all requested resources.

Fig. 7. Problematic scenarios: Scenario 1 (left) and Scenario 2 (right)

Important to note, that this solution relies on the assumption, that there is a non-zero probability of distinct messages arriving at the same destination in different orders, even if they are simultaneously sent by different sources.

The modified situation is depicted in Scenario 1, where, after agents $\{a_0, a_1\}$ initially receiving $\{(r_0, 0), (r_1, 1)\}$ and $\{(r_0, 1), (r_1, 0)\}$ slots, mutually re-request resources again. This time they receive $\{(r_0, 2), (r_1, 2)\}$ and $\{(r_0, 3), (r_1, 3)\}$ slots, and are able to form distributed lanes $dl_0(a_0)$ and $dl_1(a_1)$.

Scenario 2 However, simply re-requesting the same resources might result in a different problem. In Scenario 2, agent a_1 has requested and has been allocated a single resource r_1 which in turn modified $ppt(r_1)$ to 1 while $ppt(r_0)$ remained 0. If another agent a_0 attempts to reserve resources $\{r_0, r_1\}$, it will never receive the same promised pointer value from both resources per (re)-requesting round, and hence, will not be able to lock them.

3.2. Semi-formal protocol description

In order to address the issues described in the previous section, we developed a two-stage protocol, where the $stage_1$ of the distributed protocol specifies how an agent forms a distributed lane and $stage_2$ of the protocol addresses deadlock scenarios, which can occur after agents form distributed lanes. In the following paragraphs we semi-formally describe both stages of the protocol, which then will be formally modelled with the Event-B specification language (Algorithm 1—agent's algorithm for $stage_1$, Algorithm 2—resource's algorithm for both stages and Algorithm 3—agent's algorithm for $stage_2$).

Stage₁ An agent, which intends to reserve a set of resources starts by sending request messages to resources. The messages are sent to those resources which are part of the agent's current objective. In the provided pseudocode excerpt, we first denote relations $sent_requests$ and $objective$ where they are mappings from agents to resource collections (ln. 1-3 Algorithm 1). The messages request are sent by an agent a_n to a resource r_k ($r_k \in objective[a_n]$) until $sent_requests[a_n] = objective[a_n]$ (images are equal) (ln. 4-7 Algorithm 1). When a resource r_k receives a request message from an agent a_n it responds with a reply message which contains the current promised pointer value of resource $ppt(r_k)$ to that agent and increments the promised pointer (ln. 6-8 Algorithm 2). After sending all request messages an agent waits until reply messages are received from requested resources (ln. 8 Algorithm 1) and then makes a decision.

If all received promised pointer values are the same (a distributed lane can be formed) an agent will complete $stage_1$ by sending write messages which contain the negotiated index to all requested resources (ln. 20-24 Algorithm 1). But if one of the received promised pointer values is different an agent will start a renegotiation cycle (ln. 9-19 Algorithm 1). An agent will now send a srequest messages which contain a desired slot index to resources.

Algorithm 1 Agent stage₁ communication algorithm

```

1: variables sent_requests, received_replies, sent_srequests, sent_write typeof AGT  $\leftrightarrow$  RES init  $\emptyset$ 
2: variables objective typeof AGT  $\leftrightarrow$  RES init objective : $\in$  AGT  $\leftrightarrow$  RES
3: variables replies typeof AGT  $\leftrightarrow$   $\mathbb{N}$  init  $\emptyset$ 
4: while sent_requests[an]  $\neq$  objective[an] do                                 $\triangleright$  requesting resources which belong to the objective
5:   request(an)  $\rightarrow$  rk
6:   sent_requests := sent_requests  $\cup$  {(an, rk)}
7: end
8: wait until received_replies[an] = objective[an]
9: while | replies[an] |  $\neq$  1 do                                 $\triangleright$  enter while loop if all received indices are not the same
10:   sent_srequests[an] :=  $\emptyset$ 
11:   received_replies[an] :=  $\emptyset$ 
12:   replies[an] :=  $\emptyset$ 
13:   while sent_srequests[an]  $\neq$  objective[an] do
14:     m := max(replies[an]) + 1
15:     srequest(an, m)  $\rightarrow$  rk                                 $\triangleright$  send a special request message with a desired slot index m
16:     sent_srequests := sent_srequests  $\cup$  {(an, rk)}
17:   end
18:   wait until received_replies[an] = objective[an]
19: end
20: while sent_write[an]  $\neq$  objective[an] do
21:   m := max(replies[an])
22:   write(an, m)  $\rightarrow$  rk
23:   sent_write := sent_write  $\cup$  {(an, rk)}
24: end

```

A desired index is computed by taking the maximum of all received promised pointer values and adding a constant (one is sufficient) - ln. 14 Algorithm 1. A resource will reply to srequest message with the higher value of the current ppt(r_k) or received srequest message value and will update the promised pointer (ln. 9-11 Algorithm 2). After sending all srequest messages, an agent will wait for reply messages (ln. 18 Algorithm 1) and will restart the loop if received slot indices are not the same (ln. 9 Algorithm 1).

It is important to note that the stage 1 protocol solution to the described deadlock scenarios has a stochastic nature and one needs to guarantee that a desirable state is probabilistically reachable. In Table 4 we elicit additional safety and liveness requirements for the stage₁ of the protocol, which will need to be proved in the formal model. Stage₂ Once an agent completes sending all write messages it will wait for all pready messages from resources (ln. 4 Algorithm 3). A pready message is sent by a resource firstly if it has received a write message and no other agent is consuming that resource at that moment-resource is not locked (ln. 12–15 Algorithm 2). Secondly, a pready message will only be sent to an agent if it is distributed lane is the new minimum. In our protocol resources read pointer always take a new minimum value in the queue, once an agent sends a release message and allocation is removed from the queue.

Table 4. Low-level protocol stage₁ safety and liveness requirements

SAF ₃	An agent will not send write (form a distributed lane) messages until all receive promised pointer values are identical.
SAF ₄	Agents with overlapping resource objectives will negotiate distributed lanes with different index.
LIV ₃	An agent will eventually negotiate a distributed lane.

Algorithm 2 Resource communication algorithm

```

1: variables ppt typeof AGT  $\rightarrow \mathbb{N}$  init RES  $\rightarrow \{0\}$ 
2: variables rpt typeof AGT  $\rightarrow \mathbb{N}$  init  $\emptyset$ 
3: variables rlock typeof RES  $\rightarrow$  AGT init  $\emptyset$ 
4: variables mem typeof RES  $\rightarrow (\mathbb{N} \rightarrow \text{AGT})$  init RES  $\leftrightarrow \{\emptyset\}$ 
5: switch received_message do
6:   case request( $a_n$ ) ▷ a resource replies to a request message with a slot index ppt( $r_k$ )
7:     reply(ppt( $r_k$ ),  $r_k$ )  $\rightarrow a_n$ 
8:     ppt( $r_k$ ) := ppt( $r_k$ ) + 1
9:   case srequest( $a_n$ ,  $n$ ) ▷ a resource replies to a special request message with a slot index ppt( $r_k$ ) or  $n$ 
10:    reply(max(ppt( $r_k$ ),  $n$ ),  $r_k$ )  $\rightarrow a_n$ 
11:    ppt( $r_k$ ) := max(ppt( $r_k$ ),  $n$ ) + 1
12:   case write( $a_n$ ,  $m$ ) ▷ a resource replies to a write message with a pready message if
13:     if rlock( $r_k$ ) =  $\emptyset \wedge m = \text{rpt}(r_k)$  ▷ a resource is not locked and read pointer is at slot  $m$ 
14:       mem( $r_k$ ,  $m$ ) :=  $a_n$ 
15:       pready( $r_k$ )  $\rightarrow a_n$ 
16:   case lock( $a_n$ )
17:     if rlock( $r_k$ ) =  $\emptyset$  ▷ a resource replies with a ready message if a resource is not locked
18:       rlock( $r_k$ ) :=  $a_n$  ▷ a resource is locked and ready message is sent
19:       response( $r_k$ , READY)  $\rightarrow a_n$ 
20:     if rlock( $r_k$ )  $\neq \emptyset$ 
21:       response( $r_k$ , DENY)  $\rightarrow a_n$  ▷ resource replies with a deny message if a resource is locked
22:   case release( $a_n$ ,  $m$ ) ▷ a resource will unlock itself and remove an agent from memory slot  $m$ 
23:     mem( $r_k$ ,  $m$ ) :=  $\emptyset$ 
24:     rlock( $r_k$ ) :=  $\emptyset$ 
25:     if mem( $r_k$ )  $\neq \emptyset$  ▷ if memory is not empty a resource will
26:       rpt( $r_k$ ) := min(dom(mem( $r_k$ ))) ▷ update a read pointer to the next reserved slot
27:        $a_n := \text{mem}(r_k)(\text{rpt}(r_k))$  ▷ and send pready message to that agent
28:       pready( $r_k$ )  $\rightarrow a_n$ 
29:

```

When an agent receives all pready messages it will send lock messages to requested resources (ln. 5–8 Algorithm 3). If a resource is unlocked upon receiving lock message it will reply with ready message and lock itself, meaning, that it will stop sending pready messages (even to agent with smaller distributed lanes) until a release message is received from that agent (ln. 17–19 Algorithm 2). However, if, a resource is locked upon receiving lock message, it replies with deny message (ln. 20–21 Algorithm 2). If for every lock message an agent received a ready message, it can proceed to use resources and eventually release them. If, at least, one message is a deny message an agent will send release messages to resource (ln. 11–14 Algorithm 3) which sent ready messages to unlock them and will wait for pready messages again to repeat the process.

Algorithm 3 Agent stage₂ communication algorithm

```

1: variables received_response, received_pready, sent_lock, sent_release, consumed, released typeof
   AGT  $\leftrightarrow$  RES init  $\emptyset$ 
2: variables response typeof RES  $\leftrightarrow$  (AGT  $\leftrightarrow$  {DENY, READY}) init  $\emptyset$ 
3: while received_response[an]  $\neq$  objective[an] do
4:   wait until received_pready[an] = objective[an] ▷ wait until all pre-ready messages are received
5:   while sent_lock[an]  $\neq$  objective[an] do ▷ attempt to lock all requested resources
6:     lock(an)  $\rightarrow$  rk
7:     sent_lock := sent_lock  $\cup$  {(an, rk)}
8:   end
9:   wait until received_response[an] = objective[an] ▷ wait until all ready messages are received
10:  if DENY  $\in$  ran(response)[an] do
11:    while sent_release[an]  $\neq$  response-1[(an, DENY)] do ▷ resources which sent DENY message
12:      release(an)  $\rightarrow$  rk
13:      sent_release := sent_release  $\cup$  {(an, rk)}
14:    end
15:    received_response[an] :=  $\emptyset$ 
16:    received_pready[an] :=  $\emptyset$ 
17:    sent_lock[an] :=  $\emptyset$ 
18:    sent_release[an] :=  $\emptyset$ 
19:  end
20: end
21: while consumed[an]  $\neq$  objective[an] do ▷ all locked resources (stage2 completed) are consumed
22:   consumed := consumed  $\cup$  {(an, rk)}
23: end
24: while released[an]  $\neq$  objective[an] do ▷ all consumed resources are eventually released
25:   release(an)  $\rightarrow$  rk
26:   released := released  $\cup$  {rk, an}
27: end

```

4. Formal protocol development with Event-B

We apply the Event-B formalism to develop a high-fidelity functional model and prove the protocol functional correctness requirements. We follow the modelling process presented in Section 2. Important to note that the protocol model was redeveloped multiple times as various deadlock scenarios were found with ProB animator and model-checker. In following sections we present the final (verified) model.

4.1. A refinement strategy of the distributed protocol

The model development approach we utilise is model refinement, which starts with the abstract model that formally specifies the objective of the protocol. In fact distributed aspects of the system are ignored at this model level and the abstract model considers a centralised configuration. The abstract resource allocation protocol model was captured by two machines (m_0 and m_1). The former model essentially summarises the high-level objective of the protocol which is agents safely capturing and releasing collection of resources (objectives). This abstract model contains individual events for capturing and releasing objectives. The next refinement step introduces resources into the model and decomposes two previously introduced events according to the loop modelling pattern defined in [SIA+19].

The following group of refinement steps introduces more details about the model by primarily modelling aspects of communication. For protocol modelling, we propose to use the backward unfolding style where the next refinement step introduces the preceding protocol step. The abstract models were firstly refined with stage₂ segment of the protocol. In the refinement, m_2 , we introduced lock, response and release messages and associated events into the model. In this step we also demonstrated that the protocol stage₂ ensures safe distributed resource reservation by proving an invariant. The invariant states that no two agents will be both at resource consuming stage if both requested intersecting collections of resources. The following refinement, m_3 , is the bridge between protocol stages stage₁ and stage₂ and introduces two new messages write and pready into the model. In the final refinement step (m_4) we model stage₁ of the distributed protocol which is responsible for creating distributed lanes. The remaining messages request, reply, srequest and associated events are introduced together with the distributed lane data structure. In this refinement, we prove that distributed lanes are correctly formed.

Abstract Model Context The abstract model context introduces agents, resources and objectives into the model as finite carrier sets. The context also contains the enumerated set for the agent's status.

Message Context All messages in the protocol were defined in individual context models according to the communication pattern presented in [SIA+19].

Abstract Model The initial machine (abstract model) summarises the purpose of the protocol. The high-level objective of the protocol is to facilitate the reservation of collection of resources (objectives) by agents. The abstract model captures an agent reserving and eventually releasing an objective.

Refinement 1 (Abstract ext.) In this refinement we introduce resources into the model and now an agent tries to fulfil the objective by locking individual resources (and releasing). To model the iterative process of locking/releasing resources we use loop modelling pattern.

Refinement 2 Because of the backward unfolding style the model is then refined with stage₂ part of the protocol. In the refinement lock, response and release messages are introduced. With this refinement step we also demonstrate that the protocol's stage₂ ensures safe distributed resource reservation by proving an invariant. The invariant states that no two agents will be both at resource consuming stage if both requested intersecting collections of resources.

Refinement 3 This refinement can be considered as a bridge between the protocol's stages stage₁ and stage₂. Here, two new messages—write and pready—are introduced into the model.

Refinement 4 The final refinement step captures the stage₁ of the distributed protocol which is responsible for creating distributed lanes. The remaining messages request, reply, srequest and associated events are introduced together with the distributed lane data structure. In this refinement, we prove that distributed lanes are correctly formed (req. SAF₃₋₄).

```

CONTEXT context_abstract
SETS
  AGT, RES, OBJ  ▷ agents, resources and objectives
CONSTANTS
  objr           ▷ a constant function which maps objectives to a set of resources (axm7)
AXIOMS
  axm1..3      finite(AGT, RES, OBJ)
  axm4..6       $\emptyset \subset \text{RES, AGT, OBJ}$ 
  axm7          $\text{objr} \in \text{OBJ} \rightarrow \mathbb{P} 1(\text{RES})$ 
  axm8          $\exists o \cdot o \in \text{dom}(\text{objr}) \Rightarrow \text{card}(\text{objr}(o)) \geq 1$ 

```

Listing 1. A context of the abstract resource allocation model

```

CONTEXT context_agent_state
SETS
  AST           ▷ agents, resources and objectives
CONSTANTS
  REQUEST, WRITE, RENEGOTIATE, CONFIRM, CONSUME, RELEASE, CONFIRMW,
  LOCK, CONFIRMC, UNLOCK, CONFIRMP
AXIOMS
  axm1         partition(AST, {REQUEST}, {CONFIRMW}, {WRITE}, {RENEGOTIATE},
  {CONFIRM}, {CONFIRMP}, {LOCK}, {UNLOCK}, {CONFIRMC}, {CONSUME}, {RELEASE})

```

Listing 2. A context model defining program counter values of an agent

4.2. Protocol Event-B model: abstract context

The formal protocol modelling was started by defining static model information such as carrier sets, constants and axioms. First of all we create a context for the abstract model which contains three finite size carrier sets representing agents (AGT), resources (RES) and objectives (OBJ) as shown in Listing 1. The latter carrier set is used as an extractor operator for groups of resources in a constant function (objr).

We also introduce an enumerated set (AST) to denote agent status or in other words agent's program counter values in a separate context model context_agent_state (Listing 2). For the abstract model only (CONSUME) and (RELEASE) elements are needed whereas remaining elements will be introduced in the following subsections.

4.3. Protocol Event-B model: machine m_0

In modelling the distributed resource allocation protocol we follow a standard Event-B modelling approach where the abstract model summarises the protocol with a centralised view of the system. As previously discussed, the objective of the distributed protocol is to enable safe resource locking. This can be abstracted as agents consuming and releasing objectives (visualised in Fig. 8). In the abstract model we want to prevent agents consuming identical objectives.

To begin with, we introduce two variables for storing consumed objectives (cons) and agents status (pct0) as shown in Listing 3. The agent consume event updates the variable cons with a new pair (act₁ in Listing 4) if an objective has not been consumed (guard grd₁) and an agent is not consuming any other objectives (guard grd₂). In addition the event updates agent's program counter-variable which helps to track the steps of the agent and discharge proof obligations (guard act₂).

The second event (Listing 5) models the release of an objective by removing a pair which belonged to the cons variable and updating the program counter. The correctness of the abstract model can be verified by proving invariant (inv_{saf} in Listing 3) which asks to prove that an objective is only consumed by a single agent. The Rodin platform automatically generates proof obligations, like one below, which requires to prove that an invariant inv_{saf} is preserved by the agent_consume event.

Fig. 8. A visualisation of the abstract resource allocation model (m_0): model consists of agents which consume free objectives and then release them

Guards (and other invariants, axioms) of the `agent_consume` event become hypothesis of the proof obligation and if they are *strong* enough make it is possible for available automatic theorem provers to prove them automatically. A similar proof obligation was generated and automatically proved for the `agent_release` event.

$$\begin{array}{ll}
 cons^{-1} \in OBJ \Rightarrow AGT & (inv_{saf}) \\
 obj \in OBJ \setminus ran(cons) & (event : agent_consume : grd_1) \\
 ag \in AGT \setminus dom(cons) & (event : agent_consume : grd_2) \\
 pct0(ag) = CONSUME & (event : agent_consume : grd_3) \\
 \vdash & \\
 (cons \triangleleft \{ ag \mapsto ob \})^{-1} \in OBJ \Rightarrow AGT &
 \end{array}$$

4.4. Protocol Event-B model: machine m_1

The refinement step m_1 expands on the previous refinement by introducing resources into the model. Instead of simply consuming an objective, an agent captures resources until an objective is fulfilled (visualised in Fig. 9). Captured resources are stored in newly created variable `capt` whereas objective an agent will try to complete in the function `objt` (shown in Listing 6).

In contrary to capturing a single objective, an agent might need to consume multiple resources in order to fulfill its objective. For the iterative process, we previously introduced a two event pattern which we instantiate in this refinement step for capturing and releasing events (see Section IV.D in [SIA+19]). The loop body event `agent_consume_b` (Listing 7) takes any agent with previously initialised objective and assign a new resource `rs` to the variable `capt` (action `act1` in Listing 7). The agent must be at `CONSUME` state (guard `grd5`), the resources must be within agent's objective (guard `grd2`) and not yet be captured by any agent (guard `grd3`).

MACHINE m_0
SEES context_abstract
VARIABLES
 cons, pct0
INVARIANTS

inv_1 $cons \in AGT \Rightarrow OBJ$ $\triangleright$ variable which maps agents to consumed objectives
 inv_2 $pct0 \in AGT \Rightarrow AST$ $\triangleright$ variable which maps agents to its current program counter status
 inv_{saf} $cons^{-1} \in OBJ \Rightarrow AGT$ $\triangleright$ safety invariant which requires an objective to be consumed by only a single agent

Listing 3. The machine m_0 of the distributed resource allocation model

```

EVENT agent_consume  $\hat{=}$ 
ANY
  ag, ob  $\triangleright$  agent (ag) and objective (ob)
WHERE
  grd1 ob  $\in$  OBJ  $\setminus$  ran(cons)  $\triangleright$  take an objective that has not been consumed
  grd2 ag  $\in$  AGT  $\setminus$  dom(cons)  $\triangleright$  take an agent that has not consumed an objective
  grd3 pct0(ag) = CONSUME  $\triangleright$  program counter of the agent (ag) must be CONSUME
THEN
  act1 cons(ag) := ob  $\triangleright$  assign (consume) an objective (ob) to an agent (ag)
  act2 pct0(ag) := RELEASE  $\triangleright$  update program counter
END

```

Listing 4. Events modelling an agent consuming an objective

```

EVENT agent_release  $\hat{=}$ 
ANY
  ag, ob
WHERE
  grd1 ag  $\in$  dom(cons)  $\triangleright$  take an agent that has consumed an objective
  grd2 ob  $\in$  cons(ag)  $\triangleright$  objective consumed by agent ag
  grd3 pct0(ag) = RELEASE
THEN
  act1 cons := cons  $\setminus$  {ag  $\mapsto$  ob}  $\triangleright$  remove agent and its consumed objective
  act2 pct0(ag) := CONSUME  $\triangleright$  update program counter
END

```

Listing 5. Events modelling an agent releasing resources an objective

Fig. 9. A visualisation of the (m_1) refinement model: agents are capturing free resources to fulfil an objective (a resource can belong to different objectives)


```

MACHINE m1
REFINES m0
VARIABLES
  capt, objt, pct1
INVARIANTS
  inv1  capt ∈ AGT → ℙ(RES)  ▷ a variable mapping agents to captured resources
  inv2  objt ∈ AGT → OBJ      ▷ variable mapping an agent to an objective it tries to fulfil
  ⋮
  invsaf  ∀ a1, a2 · a1 ∈ dom(capt) ∧ a2 ∈ dom(capt) ∧ a1 ≠ a2 ⇒ capt(a1) ∩ capt(a2) = ∅

```

Listing 6. Variables and invariants of the refinement step m₁

```

EVENT agent_consume_b ≡
ANY
  ag, rs  ▷ agent (ag) and resource (rs)
WHERE
  grd1  ag ∈ dom(capt)  ▷ take some agent (ag) from domain of variable captured
  grd2  rs ∈ objr(objt(ag))  ▷ take a resource (rs) which is part of ag objective (objt(ag))
  grd3  rs ∉ union(ran(capt))  ▷ resource must not be captured by other agent
  grd4  union(ran({ag} ◁ capt)) ∩ objr(objt(ag)) = ∅  ▷ no resources from the objective can be captured
  grd5  pct1(ag) = CONSUME  ▷ program counter of agent must be at CONSUME
THEN
  act1  capt(ag) := capt(ag) ∪ {rs}  ▷ add resource (rs) to agent's captured set
END

```

Listing 7. The event modelling an agent consuming a free resource: loop body event

The agent *ag* will also capture a resource *rs* only if none of the resources from its objective have been captured by other agents (guard *grd₄*). The guard *grd₄* prevents a deadlock scenario in which agents with overlapping objectives are unable to complete their objective as remaining resources are captured by the other overlapping agent.

The loop completion event *agent_consume_c* (Listing 8) would be triggered as soon as the objective has been fulfilled (guard *grd₁* in Listing 8) and program counter would be updated to new state—RELEASE (action *act₁* in Listing 8). Similarly, in this refinement we transform *agent_release* event according based on the event pattern presented. To show correctness of the extended model, we prove an invariant (*inv_{saf}* in Listing 6) which states that no two agents can have the same resource captured.

```

EVENT agent_consume_c ≡
ANY
  ag
WHERE
  grd1  capt(ag) = objr(objt(ag))  ▷ guard which checks wheter resource obective has been fulfilled
  grd2  pct1(ag) = CONSUME
THEN
  act2  pct0(ab) := RELEASE  ▷ update agent's program counter to RELEASE
END

```

Listing 8. The event modelling an agent consuming a free resource: loop completion event

Fig. 10. Sequence diagram of stage₂ of the protocol at refinement level m_2

4.5. Protocol Event-B model: machine m_2

In this refinement step, we start to consider communication as we introduce protocol (described in Section 3) into the model. Because of backward unfolding style we introduce stage₂ part of the protocol first (sequence diagram shown in Fig. 10). At this stage of the protocol three types of messages are sent: lock and release by an agent and response message by a resource. In this refinement step, we apply communication modelling patterns described in Section IV.C of the paper [SIA+19].

MACHINE m_2

REFINES m_1

VARIABLES

rdpt, lck, lcke, rsp, rel, pct2

INVARIANTS

inv₁ rdpt ∈ RES ⇒ AGT ▷ variable representing locks of resources (resource locked by an agent)

inv₂ lck ⊆ LCK ▷ lock message channel variable

inv₃ lcke ∈ AGT ⇒ $\mathbb{P}(\text{RES})$

inv₄ rsp ⊆ RSP ▷ response message channel variable

inv₅ rel ⊆ REL ▷ release message channel variable

inv₆ pct2 ∈ AGT ⇒ AST

Listing 9: Variables and invariants of the refinement step m_2

First of all, machine m_2 is extended with three context models—separately defining each message type according to the generic message context pattern (response message context model shown in Listing 12). Then, by following machine communication modelling pattern we create variables lck, rsp and rel which represent communication channels of each message. In addition an agent's variable lcke is created to store already sent lck messages (shown in Listing 9).

In the stage₂ an agent tries to lock resources associated with negotiated distributed lane by sending lock messages. To model that, we apply loop modelling patterns and create two new events: agent_lock_b and agent_lock_c (Listings 10 and 11). Since, preceding protocol messages are modelled in the next refinements, we model lock message as a initiating message at this stage but later convert it to a reply event type. The first event is the body of a message sending loop which sends a new lock message (action act₁ in Listing 10) if a message lc has not been sent before (guard grd₁) and destination (resource) is within agent's objective (guard grd₃). The program counter of the agent must be in LOCK protocol phase (guard grd₄). An agent also saves a local copy of the sent message with the event action act₂.

The second event in the loop pattern is a loop completion event agent_lock_c which detects the end of the loop and updates the program counter (Listings 11). For the lock message sending event—an agent must detect when all messages have been sent or in other words the objective has been fulfilled (guard grd₂).

```

EVENT agent_lock_b  $\hat{=}$ 
ANY
  lc  $\triangleright$  lc (lock) message which will be sent
WHERE
  grd1 lc  $\in$  LCK  $\setminus$  lck  $\triangleright$  only a new lock message will be sent (lc message not already in the lck channel)
  grd2 lckd(lc)  $\notin$  lcke(lcks(lc))
  grd3 lckd(lc)  $\in$  objr(objt(lcks(lc)))  $\triangleright$  a lock message is sent to a resource which belongs to agent's objective
  grd4 pct2(lcks(lc)) = LOCK  $\triangleright$  lock message can only be sent by an agent with LOCK program status
THEN
  act1 lck := lck  $\cup$  {lc}  $\triangleright$  a new lock message is added to lock message channel
  act2 lcke(lcks(lc)) := lcke(lcks(lc))  $\cup$  lckd(lc)  $\triangleright$  message sending receipt is saved locally by an agent
END

```

Listing 10. The event modelling an agent locking resources: loop body event

```

EVENT agent_lock_c  $\hat{=}$ 
ANY
  ag
WHERE
  grd1 ag  $\in$  dom(lcke)
  grd2 lcke(ag) = objr(objt(ag))  $\triangleright$  lock messages are sent to all resources which belong to agent's objective
  grd3 pct2(ag) = LOCK
THEN
  act1 pct2(ab) := CONFIRM_C  $\triangleright$  once all lock messages are sent agent's program counter is updated
END

```

Listing 11. The event modelling an agent locking resources: loop completion event

This event simply updates the program counter to CONFIRM_C state with the action (action act₁). According to the protocol, after sending all lock messages an agent will wait for replies before it proceeds to consuming resources.

In order to model a resource's response to the lock message first we created a read pointer variable rdpt to work as a resource lock for an agent. The resource lock is released once a resource receives a release message from an agent who locked the resource (see Listing 9). A resource sending a response message is a reply type message therefore we used a reply event modelling pattern (Listings 13 and 14). A resource sends response message when lock message has been received-guard grd₁ in both events. The following two guards (grd₄ and grd₅) define the source and destination of the new message which are respectively destination and source of received lock message.

```

CONTEXT
resource_response_message
SETS
  RSP, CNF
CONSTANTS
  rps, rspd, rspv, READY, DENY
AXIOMS
  axm1 rps  $\in$  RSP  $\rightarrow$  RES  $\triangleright$  source constant function of the response message
  axm2 rspd  $\in$  RSP  $\rightarrow$  AGT  $\triangleright$  destination constant function of the response message
  axm3 rspv  $\in$  RSP  $\rightarrow$  CNF  $\triangleright$  value constant function of the response message
  axm4 partition(CNF, {READY}, {DENY})  $\triangleright$  response message value can be READY or DENY
  axm5 finite(RSP)
  axm6  $\forall s, d, v \cdot s \in \text{RES} \wedge d \in \text{AGT} \wedge v \in \text{CNF} \Rightarrow \exists m \cdot \text{rps}(m) = s \wedge \text{rspd}(m) = d \wedge \text{rspv}(m) = v$ 

```

Listing 12: The context model of the resource response type message

Lastly, the response message also carries a value and guards grd_6 define the value in both events. The message carries READY value if the resource is not locked by other agent (grd_3 in Listing 13) otherwise a DENY message is sent (grd_3 in Listing 14). In addition to sending a message, the `resource_response_ready` event also removes answered message and if READY message was sent resource locks itself for that agent with action act_3 .

```

EVENT resource_response_ready  $\hat{=}$ 
ANY
  lc, rs  $\triangleright$  lock (lc) and response (rs messages)
WHERE
   $\text{grd}_1$   $lc \in \text{lck}$   $\triangleright$  a response READY message will be only sent if a lock message was sent
   $\text{grd}_2$   $rs \in \text{RSP} \setminus \text{rsp}$   $\triangleright$  only a previously not sent response READY message will be sent
   $\text{grd}_3$   $\text{rsps}(rs) \notin \text{dom}(\text{rdpt})$   $\triangleright$  a response READY message is sent by a resource only if it is not locked
   $\text{grd}_4$   $\text{rsps}(rs) = \text{lckd}(lc)$   $\triangleright$  source of the response message is destination of received lock message
   $\text{grd}_5$   $\text{rspd}(rs) = \text{lcks}(lc)$   $\triangleright$  destination of the response message is source of received lock message
   $\text{grd}_6$   $\text{rspn}(rs) = \text{READY}$   $\triangleright$  response message contains READY value
THEN
   $\text{act}_1$   $\text{rps} := \text{rsp} \cup \{rs\}$   $\triangleright$  response message is added (sent) to response channel
   $\text{act}_2$   $\text{lck} := \text{lck} \setminus \{lc\}$   $\triangleright$  received and replied lock message is removed from the channel
   $\text{act}_3$   $\text{rdpt} := \text{rdpt} \leftarrow \{\text{rsps}(rs) \mapsto \text{rspd}(rs)\}$   $\triangleright$  resource which sent READY response message is locked for the agent
END

```

Listing 13: The event modelling a resource sending a response message: READY type

```

EVENT resource_response_deny  $\hat{=}$ 
ANY
  lc, rs
WHERE
   $\text{grd}_1$   $lc \in \text{lck}$   $\triangleright$  a response DENY message will be only sent if a lock message was sent
   $\text{grd}_2$   $rs \in \text{RSP} \setminus \text{rsp}$ 
   $\text{grd}_3$   $\text{rsps}(rs) \in \text{dom}(\text{rdpt})$   $\triangleright$  a response DENY message is sent by a resource if it is locked
   $\text{grd}_4$   $\text{rsps}(rs) = \text{lckd}(lc)$ 
   $\text{grd}_5$   $\text{rspd}(rs) = \text{lcks}(lc)$ 
   $\text{grd}_6$   $\text{rspn}(rs) = \text{DENY}$   $\triangleright$  response message contains DENY value
THEN
   $\text{act}_1$   $\text{rps} := \text{rsp} \cup \{rs\}$   $\text{rps} := \text{rsp} \cup \{rs\}$   $\triangleright$  response message is added (sent) to response channel
   $\text{act}_2$   $\text{lck} := \text{lck} \setminus \{lc\}$   $\text{lck} := \text{lck} \setminus \{lc\}$   $\triangleright$  received and replied lock message is removed from the channel
END

```

Listing 14: The event modelling a resource sending a response message: DENY type

Once an agent sends all lock messages it must receive all response messages before it can progress. In the model, we created a new event `agent_decide` (Listing 15) which checks the values of received response messages and selects the new program counter value. In principle this event is a loop completion event but in the model we decided to model *decision* events separately to reduce events complexity. The main difference from the basic loop completion event is that depending on conditions different program counter values can be selected. In `agent_decide` event if all messages contained READY value (guard grd_2) then the event will update program counter to CONSUME. But if a DENY message (guard grd_3) was received program counter is changed to UNLOCK state. In the latter scenario an agent must then send release messages to resources, which sent READY messages and locked their resources, so protocol can progress.

For selecting correct message from the channel we use messages constant functions and a channel variable— $[\text{rsp} \cap \text{rspd}^{-1}[\{\text{ag}\}]]$ selects all `rsp` messages which were sent to agent `ag`. Inserting the result to the `rsps` constant function we get sources (resources) of these messages. This message extraction guard pattern is used widely when relevant messages need to be selected. Events of the unlocking phase are modelled using reply communication pattern and hence not covered in this subsection. Once relevant resources are unlocked agent tries to lock them again in this refinement step.

```

EVENT agent_decide  $\hat{=}$ 
ANY
  ag, pc  $\triangleright$  agent (ag) and program counter (pc)
WHERE
  grd1   $\text{rsps}[\text{rsp} \cap \text{rspd}^{-1}[\{\text{ag}\}]] = \text{lcke}(\text{ag}) \triangleright$  all response messages have to be received
  grd2   $\text{rspn}[\text{rsp} \cap \text{rspd}^{-1}[\{\text{ag}\}]] = \{\text{READY}\} \Rightarrow \text{pc} = \text{CONSUME} \triangleright$  if all response messages are READY
  grd3   $\text{DENY} \in \text{rspn}[\text{rsp} \cap \text{rspd}^{-1}[\{\text{ag}\}]] \Rightarrow \text{pc} = \text{UNLOCK} \triangleright$  if at least one response message is DENY
  grd4   $\text{pct2}(\text{ag}) = \text{CONFIRM}$ 
THEN
  act1   $\text{pct2}(\text{ag}) := \text{pc} \triangleright$  update agent's program counter based on guards 2 – 3
END

```

Listing 15. The event modelling agents decision between restarting stage₂ or consuming resources

INVARIANTS

$$\text{SAF}_1 \quad \forall a_1, a_2 \cdot \text{pct2}(a_1) = \text{CONSUME} \wedge \text{pct2}(a_2) = \text{CONSUME} \wedge a_1 \neq a_2 \Rightarrow \text{objr}(\text{objt}(a_1)) \cap \text{objr}(\text{objt}(a_2)) = \emptyset$$

Listing 16. The safety invariant for prohibiting mutual resource locking by different agents

In this refinement step, which specifies stage₂ of the distributed protocol, we were required to prove invariant expressed in Listing 16. The safety invariant relates to the mutual exclusion safety protocol property SAF₁ specified in Fig. 3. The property requires proving that if any two agents are in the CONSUME state their objectives (set of resources) must not be overlapping. The Rodin platform has generated five proof obligations based from the SAF₁ invariant. Four out of five proof obligations were proved automatically, however, the safety invariant relating to SAF₁ property and agent_decide event had to be proved interactively. The interactively discharged proof obligation is shown below, where the guards and invariant form the hypothesis of the proof obligation while the invariant modified by the agent_decide event action is goal of the sequent.

$$\begin{aligned}
 &\text{rsps}[\text{rsp} \cap \text{rspd}^{-1}[\{\text{ag}\}]] = \text{lcke}(\text{ag}) && (\text{grd}_1) \\
 &\text{rspn}[\text{rsp} \cap \text{rspd}^{-1}[\{\text{ag}\}]] = \{\text{READY}\} \Rightarrow \text{pc} = \text{CONSUME} && (\text{grd}_2) \\
 &\forall a_1, a_2 \cdot (\text{pct2})(a_1) = \text{CONSUME} \wedge \text{pct2}(a_2) = \text{CONSUME} \wedge a_1 \neq a_2 \Rightarrow && (\text{inv}_{\text{saf}}) \\
 &\text{objr}(\text{objt}(a_1)) \cap \text{objr}(\text{objt}(a_2)) = \emptyset \\
 &\vdash \\
 &\forall a_1, a_2 \cdot (\text{pct2} \triangleleft \{ \text{ag} \mapsto \text{pc} \})(a_1) = \text{CONSUME} \wedge (\text{pct2} \triangleleft \{ \text{ag} \mapsto \text{pc} \})(a_2) = \text{CONSUME} \wedge a_1 \neq a_2 \Rightarrow \\
 &\text{objr}(\text{objt}(a_1)) \cap \text{objr}(\text{objt}(a_2)) = \emptyset
 \end{aligned}$$

4.6. Protocol Event-B model: machine m₃

In this refinement step we continue to unfold protocol backwards and introduce two new messages write and pready. In fact this refinement step can be thought as an intermediate step gluing protocol stages stage₁ and stage₂ (sequence diagram shown in Fig. 11). Events for capturing write message sending were developed using the initiating message pattern. Since they were structurally identical to the previous refinement events agent_lock_b and agent_lock_c events we do not discuss them here. The more interesting challenge of this refinement step was correctly capturing the pready message. A variable of type $\text{mem} \in \text{RES} \rightarrow \mathbb{P}(\text{AGT})$ was created to abstract queue like resource memories (see Listing 17) which are introduced in the next refinement step.

A resource sends pready message to inform an agent of its availability for consumption and that it can be locked now. There are two different cases when this message should be sent and instead of constructing a single event to cover all scenarios it was decided to create more trivial events for each of the cases.

Fig. 11. Sequence diagram of stage₁ of the protocol at refinement level m₃**MACHINE** m₃**REFINES** m₂**VARIABLES**

wrt, prd, mem, wrte, pct3

INVARIANTS

inv_1 wrt ⊆ WRT ▷ write message channel variable

inv_2 prd ⊆ PRD ▷ pready message channel variable

inv_3 mem ∈ RES → P(AGT) ▷ resource memory variable storing agents which have sent write message

Listing 17. Variables and invariants of the machine model m₃

In the model we created a reply type `resource_write_preready` event which creates a new message and removes the answered message from the `wrt` channel if a resource is not locked. A pready message will be sent if a write message has been received and resource is not locked by an agent.

The second event `resource_release_pready` was necessary to cover a scenario where pready message was sent to an agent in response to its write message but in the end the agent was not able to lock all resources. This would mean that write messages have been removed and `resource_write_preready` guards would never be satisfied for that agent. Yet an agent would be still waiting to receive pready messages eventually. Therefore we introduced initiating message `resource_release_pready` event which sends a new pready message to any agent which is interested in that resource if a resource is not locked.

Fig. 12. Sequence diagram of stage₁ of the protocol at refinement level m_4

4.7. Protocol Event-B model: machine m_4

The final refinement step m_4 introduces the remaining part of the protocol—stage₁. This stage starts with an agent requesting a set of resources and finishes with agent forming a distributed lane by sending a write messages to the same resources. In between, an agent might need to send a special request message srequest if a distributed lane was not negotiated in the first attempt. To allow distributed lane forming—a resource memory introduced in the previous refinement now must be refined to queue-like data structure with corresponding variables (e.g. read pointer, promised pointer).

MACHINE m_4

REFINES m_3

VARIABLES

req, rqst, rqs, rqt, rep...

INVARIANTS

inv_1 req $\subseteq$ REQ $\triangleright$ request message channel

inv_2 rqst $\in$ AGT $\rightarrow \mathbb{P}(\text{RES})$

inv_3 rqs $\subseteq$ REQ $\triangleright$ special request message channel

inv_4 rqt $\in$ AGT $\rightarrow \mathbb{P}(\text{RES})$

inv_5 rep $\subseteq$ REP $\triangleright$ reply message channel

inv_6 ppt $\in$ RES $\rightarrow \mathbb{N}$ $\triangleright$ promised pointer variable

inv_7 rpt $\in$ RES $\rightarrow \mathbb{N}$ $\triangleright$ read pointer variable

inv_8 lan $\in$ RES $\rightarrow (\mathbb{N} \rightarrow \text{AGT})$ $\triangleright$ variable where distributed lanes are stored

inv_9 pct4 $\in$ AGT $\rightarrow \text{AST}$

Listing 18: Variables and invariants of the machine model m_4

To begin with we start with introducing communication channel variables for request, reply and srequest messages. In addition two local variables rqst and rqt were created to store messages on the agent side (shown in Listing 18). Following the loop and initiating message modelling patterns we introduced two events agent_request_b and agent_request_c for requesting resources (Listings 19 and 20).


```

EVENT agent_request_b  $\hat{=}$ 
ANY
  rq
WHERE
  grd1  rq  $\in$  REQ  $\setminus$  req  $\triangleright$  send a new request message
  grd2  reqd(rq)  $\notin$  rqst(reqs(rq))
  grd3  reqd(rq)  $\in$  objr(objt(reqs(rq)))  $\triangleright$  request message destination (resource) must be part of the objective
  grd4  pct4(reqs(rq)) = REQUEST
THEN
  act1  req := req  $\cup$  {rq}  $\triangleright$  send a request message
  act2  rqts(reqs(rq)) := rqts(reqs(rq))  $\cup$  {reqd(rq)}
END

```

Listing 19: The event modelling an agent requesting a resource: loop body event

After all request messages have been sent event agent_request_c changes agents program counter to CONFIRMW state (act₁ in Listing 20) which informs an agent to wait until all reply messages have been received.

```

EVENT agent_request_c  $\hat{=}$ 
ANY
  ag
WHERE
  grd1  ag  $\in$  dom(rqst)
  grd2  rqst(ag) = objr(objt(ag))  $\triangleright$  check if all request messages have been sent
  grd3  pct4(ag) = REQUEST
THEN
  act1  pct4(ag) := CONFIRMW  $\triangleright$  update program counter value
END

```

Listing 20: The event modelling an agent requesting a resource: loop completion event

On the resource side, a reply type event resource_reply_general (Listing 21) was created which simply sends a new reply message (action act₁) containing the current promised pointer value (guard grd₅), removes request message (act₂) and increments the promised pointer (act₃). Once an agent has sent all request and respectively received all reply messages a decision must be made whether to form a distributed lane or renegotiate new indexes.

```

resource_reply_general  $\hat{=}$ 
ANY
  rp, rq  $\triangleright$  reply message (rp) and request message (rq)
WHERE
  grd1  rq  $\in$  req
  grd2  rp  $\in$  REP  $\setminus$  rep
  grd3  repd(rp) = reqs(rq)
  grd4  reps(rp) = reqd(rq)
  grd5  repn(rp) = ppt(reqs(rq))
THEN
  act1  rep := rep  $\cup$  {rp}  $\triangleright$  send a reply message
  act2  req := req  $\setminus$  {rq}  $\triangleright$  remove request message which was replied
  act3  ppt(reps(rp)) := ppt(reps(rp)) + 1  $\triangleright$  update (increment) resource's promised pointer value
END

```

Listing 21: The event modelling a resource replying to a EVENTrequest message

The distributed lane will be formed if all reply messages contained the same promised pointer value otherwise an agent must renegotiate a distributed lane. In the model, we created a new *decision* agent_confirm_write_renegotiate event which changes the program counter according to the named conditions (Listing 22). To check whether all reply messages contained the same index it is sufficient to check the set size (cardinality) of the

$\text{repn}[\text{rep} \cap \text{repd}^{-1}[\text{ag}]]$ values set. The guard grd_1 must have both local variables because this event can be enabled after general request or srequest message sending.

EVENT $\text{agent_confirm_write_renegotiate} \triangleq$

ANY

ag, pc

WHERE

$\text{grd}_1 \quad \text{ag} \in \text{repd}[\text{rep}]$

$\text{grd}_2 \quad \text{card}(\text{rqst}(\text{ag}) \cup \text{rqts}(\text{ag})) = \text{card}(\text{rep} \cap \text{repd}^{-1}[\{\text{ag}\}]) \quad \triangleright \text{all sent request messages have been replied}$

$\text{grd}_3 \quad \text{reps}[\text{rep} \cap \text{repd}^{-1}[\{\text{ag}\}]] = \text{objr}(\text{objt}(\text{ag}))$

$\text{grd}_4 \quad \text{pct4}(\text{ag}) = \text{CONFIRMW}$

$\text{grd}_5 \quad \text{card}(\text{repn}[\text{rep} \cap \text{repd}^{-1}[\{\text{ag}\}]] > 1 \Rightarrow \text{pc} = \text{RENEGOTIATE} \quad \triangleright \text{renegotiate resources}$

$\text{grd}_6 \quad \text{card}(\text{repn}[\text{rep} \cap \text{repd}^{-1}[\{\text{ag}\}]] = 1 \Rightarrow \text{pc} = \text{WRITE} \quad \triangleright \text{form a distributed lane if all reply values identical}$

THEN

$\text{act}_1 \quad \text{pct4}(\text{ag}) := \text{pc}$

END

Listing 22: Agent decision that decides whether to restart stage_1 or start stage_2

If an agent was not able to negotiate a distributed lane it must send a new, special message—srequest. This time instead of requesting an arbitrary index a new srequest message contains a desired index which is computed by adding a non-zero constant to the highest reply index. To send srequest messages we created two reply type events: $\text{agent_renegotiate_b}$ and $\text{agent_renegotiate_c}$.

As a standard reply type event $\text{agent_renegotiate_b}$ creates a new message and removes the answered reply message from the channel. Whereas $\text{agent_renegotiate_c}$ event updates the program counter again to CONFIRMW state - this cycle repeats until a distributed lane is negotiated.

Lastly in this refinement we introduce distributed lane data structure and accordingly update relevant events. Two variables ppt and rpt are created to respectively represent promise pointer and read pointer of a resource. For each resource the promised pointer is initialised to zero and only two reply events $\text{resource_reply_general}$ and $\text{resource_reply_special}$ (Listing 23) modify a promised pointer variable. The first event simply increments the current $\text{ppt}(\text{r})$ value after reply message has been sent whereas the latter event updates the $\text{ppt}(\text{r})$ by computing action act_3 , where maximum function parameters are the current $\text{ppt}(\text{r})$ value and received srequest message value.

EVENT $\text{resource_reply_special} \triangleq$

ANY

$\text{rp}, \text{rq} \quad \triangleright \text{reply message (rp) and special request message (rq)}$

WHERE

$\text{grd}_1 \quad \text{rq} \in \text{rqs}$

$\text{grd}_2 \quad \text{rp} \in \text{REP} \setminus \text{rep}$

$\text{grd}_3 \quad \text{repd}(\text{rp}) = \text{rqss}(\text{rq})$

$\text{grd}_4 \quad \text{reps}(\text{rp}) = \text{rqsd}(\text{rq})$

$\text{grd}_5 \quad \text{repn}(\text{rp}) = \max(\{\text{ppt}(\text{reps}(\text{rp})), \text{rqsn}(\text{rq})\}) \quad \triangleright \text{reply message value is a larger value between resources ppt value and value from received special request message}$

THEN

$\text{act}_1 \quad \text{rep} := \text{rep} \cup \{\text{rp}\}$

$\text{act}_2 \quad \text{rqs} := \text{rqs} \setminus \{\text{rq}\}$

$\text{act}_3 \quad \text{ppt}(\text{reps}(\text{rp})) := \max(\{\text{ppt}(\text{reps}(\text{rp})), \text{rqsn}(\text{rq})\}) + 1 \quad \triangleright \text{update resource's promised pointer}$

END

Listing 23: The event modelling a resource replying to a srequest message

Read pointers are updated by two events which send pready messages. In contrary to the previous variable the read pointer $\text{rpt}(\text{r})$ is always set to the minimum value of the request pool. This is necessary as some agent might negotiate a distributed lane with lower index than others but its write messages are delayed (or even lost) so the protocol would halt. Allowing agents with higher distributed lane indexes but sooner write message arriving to consume resources introduces fault-tolerance into the protocol. Guards of these two events were also strengthened to only send pready messages to newly incoming agents if their index is the lowest in the request pool.

INVARIANTS

$$\text{SAF}_3 \quad \forall ag \cdot \text{pct4}(ag) = \text{WRITE} \wedge \text{reps}[\text{rep} \cap \text{repd}^{-1}[\{ag\}]] = \text{objr}(\text{objt}(ag)) \Rightarrow \\ \text{card}(\text{repn}[\text{rep} \cap \text{repd}^{-1}[ag]]) = 1$$

Listing 24. Event-B safety invariant ensuring that SAF_3 requirement is preserved in the model**MACHINE** m_4 **VARIABLES**

$\text{req}, \dots, \text{his}_{\text{ppt}}, \text{his}_{\text{wr}}$

INVARIANTS

$$\begin{aligned} & \vdots \\ & \text{his}_{\text{ppt}} \in \text{RES} \rightarrow (\mathbb{N} \rightarrow \mathbb{N}) \\ & \text{his}_{\text{wr}} \in \text{RES} \rightarrow \mathbb{N} \\ & \text{inv}_{\text{saf4}} \quad \forall r, n_1, n_2 \cdot r \in \text{RES} \wedge n_1, n_2 \in \text{dom}(\text{his}_{\text{ppt}}(r)) \wedge n_1 < n_2 \Rightarrow \text{his}_{\text{ppt}}(r)(n_1) < \text{his}_{\text{ppt}}(r)(n_2) \\ & \text{inv}_{\text{his}_{\text{ppt}}} \quad \forall \text{res} \cdot (\text{his}_{\text{wr}}(\text{res}) = 0 \wedge \text{his}_{\text{ppt}}(\text{res}) = \emptyset) \vee \\ & \quad (\text{dom}(\text{his}_{\text{ppt}}(\text{res})) = 0 \dots \text{his}_{\text{wr}}(\text{res}) - 1 \wedge \text{his}_{\text{ppt}}(\text{res})(\text{his}_{\text{wr}}(\text{res}) - 1) = \text{ppt}(\text{res}) - 1) \end{aligned}$$

Listing 25. Additional variables and invariants added for provid safety property SAF_4

Proving SAF_{3-4} Safety Properties. As shown in Section 3.2 (Scenarios 1 - 2) high-level system requirements can only be met if an agent invariably and correctly forms a distributed lane. The probabilistic lane forming eventuality (LIV_3) is discussed in [SIK+20] while in the following paragraphs we describe proving requirements SAF_{3-4} in the Event-B model.

First of all, requirement SAF_3 (see Table 4) relates to SAF_2 (see Table 3) and is needed to ensure that that an agent is not allocated only a part of resources it has requested. In the model we ensure that SAF_3 requirement is satisfied by strengthening event guards which restrict enabling states of the event that generates an outgoing write message. To cross-check this implementation we add an invariant that directly shows that SAF_3 is maintained in the model (Listing 24).

In the following paragraphs we provide details on proving a slightly more interesting case of SAF_4 which assumes that SAF_3 property is proven. Requirement SAF_4 addresses potential cross-blocking deadlock scenarios described in Section 3.1 or a resource double locking due to distributed lane overriding. The strategy to proving the SAF_4 requirement is to formally prove that agents with intersecting objectives (at least one common resource) always form distributed lanes with differing indices. Based on a distributed lane definition (Definition 1 in Section 3.1) we could formally express the safety invariant as expression (1) which states if two different agents (a_1, a_2) have formed a distributed lane (dl) over an intersecting set of resources ($\text{dom}(\text{dl}(a_m)) = \{r_0, r_1, \dots, r_n\}$) then slot values s of their distributed lanes must be different.

$$\begin{aligned} \forall a_1, a_2 \cdot a_1, a_2 \in \text{AGENTS} \wedge a_1 \neq a_2 \wedge \text{dom}(\text{dl}(a_1)) \cap \text{dom}(\text{dl}(a_2)) \neq \emptyset \\ \Rightarrow \text{ran}(\text{dl}(a_1)) \cap \text{ran}(\text{dl}(a_2)) \neq \emptyset \end{aligned} \quad (1)$$

Instead of directly including predicate (1) into the Event-B model and completing a complex interactive SAF_4 property proof, it was realised that property (1) can be expressed (and proved) as a simpler verification condition. Let's first assume that agents only form distributed lanes if all received indices are the same (proved as SAF_3). If the same resource is requested by different agents and they are replied with different promised pointer values, these promised pointer values (slots) will be distributed lane *deciders* as all other slots from remaining resources must be the same in order to form a distributed lane. Therefore, to prove that model preserves SAF_4 requirement it is enough to show that a resource always replies to a request or special request message with a unique promised pointer value. We can prove that a resource always replies with a unique promised pointer value by introducing a history variable into our model.

To prove this property we first extend our model with a variable his_{ppt} and an entry variable his_{wr} or *time-stamp* as shown in Listing 25. The former variable stores each resources chronological promised pointer updates and latter works as a write pointer for history variable.

```

EVENT resource_reply_general  $\hat{=}$ 
ANY
  rq, rp, res
WHERE
  grd1  rq  $\in$  req
   $\vdots$ 
  grd6  res = reps(rp)
THEN
  act1  rep := rep  $\cup$  {rp}  $\triangleright$  send a reply message
   $\vdots$ 
  act4  hisppt(res) := hisppt(res)  $\triangleleft$  {hiswr(res)  $\mapsto$  ppt(res)}  $\triangleright$  update history ppt variable
  act5  hiswr(res) := hiswr(res) + 1  $\triangleright$  update history write variable
END

```

Listing 26. The event modelling a resource replying to a request message with history variable extension

Table 5. Event-B protocol model proof statistics

Model	No. of POs	Aut. Discharged	Int. Discharged
context c_0	0	0	0
context mes.	9	9	0
machine m_0	12	12	0
machine m_1	23	21	2
machine m_2	59	43	16
machine m_3	43	32	11
machine m_4	103	57	46
Total	249	174	75

After introducing history variables, we updated two events `resource_reply_{general, special}` which update promised pointer variables (after sending reply messages) by adding two new actions `act4` and `act5` as shown in Listing 26 (identical extension for `resource_reply_special` event).

Action `act4` updates a history variable for a resource `res` with the current write stamp and promised pointer (`ppt(res)`) value sent. The next action `act5` simply updates the resource’s write stamp. We can then add the main invariant to prove (`invsaf4` in Listing 25) which states that if we take any two entries `n1`, `n2` of the history variable for the same resource where one is larger, then that larger entry should have larger promised pointer value.

To prove that `resource_reply_{general, special}` preserve `invsaf4`, the following properties play the key role: (1) the domain of `hisppt` (i.e., ‘indices’ of `hisppt`) is $\{0, \dots, \text{his}_{wr} - 1\}$, (2) `hisppt(hiswr - 1) < hisppt(hiswr)`. Property (2) holds because `hisppt(hiswr)` is the maximum of promised pointer (`ppt`) and special request slot number and promised pointer is incremented as `resource_reply_{general, special}` occurs. We also specified these properties as an invariant (`invhisppt`) and proved they are preserved by the events which helped to prove `invsaf4`.

Proof statistics. In Table 5 we provide an overall proof statistics of the Event-B protocol model which may be used as a metric for models complexity. The majority of the generated proof obligations were automatically discharged with available solvers and even a large fraction of interactive proofs required minimum number of steps. We believe that a high proof automation was due to modelling patterns [SIA+19] use and the SMT-based verification support [ISAYR16, DFGV14].

5. Conclusions and future work

In this paper, we demonstrated how a previously proposed formal development methodology of distributed systems can be used to developing a distributed signalling system protocol. Starting only with high-level system requirements, we developed an early formal protocol prototype model, which with the help of ProB model checker was refined as subtle deadlock scenarios were discovered. Because of the stepwise modelling principle and our refinement strategy, re-modelling efforts were significantly reduced as only a part of the complete functional Event-B model needed to be reworked. For example, issues discovered in the stage₁ of the distributed protocol only affected final refinement step m_4 and so machine models $m_{0..3}$ and associated (completed) proofs were preserved. The stepwise distributed protocol development, as also shown before [CM06, HKBA09, ILTR11], together with adequate tools [ISAYR16, DFGV14] helped to achieve fairly high verification automation. The final protocol model was formally proved to guarantee a safe and deadlock free distributed signalling operation. We demonstrated this by expressing and deductively proving invariants which stated that a distributed lane would be correctly and eventually formed by an agent. An alternative deadlock-freedom proving strategy would have been proving a deadlock freedom proof obligation (page 5 DLK [Hoa14]) which was considered in our work. The resulting proof obligation results in a complex proof obligation to be discharged in the Rodin platform.

In this work, we have also demonstrated that the previously developed communication modelling patterns [SIA+19] can model a variety of message sending/receiving scenarios, and therefore, provide a more systematic communication modelling method. On the other hand, the proposed formal development methodology does not provide a systematic process to refine the generic communication-based signalling model. Formalising refinement steps would fully utilise the benefits of the Event-B stepwise modelling and further reduce modelling and verification effort.

We believe that the future work of defining refinement patterns can be built around some initial refinement ideas in our paper [SIA+19]. Furthermore, the modelling effort with our methodology could be significantly reduced by providing an automatic refinement of communication models from a high-level protocol specification, for example, sequence diagrams, which are a widely used method to specify protocols in the railway industry. Another important direction we would like to explore in the future work is protocol's viability to be deployed in practice. At the moment the distributed resource allocation protocol solves railway resource allocation problem at a high abstraction level and in the future work the protocol model should be refined to consider other railway signalling elements such as switches, signals and routes as in [HP00]. The refined version of the protocol would also need to solve another type of deadlock scenario which would occur once railway routes are introduced into the model. Specifically, the head-to-head deadlock scenario, which can occur when two trains cannot progress as each train is preventing the other one from locking the next set of resources (example discussed in Section 7.1.4 of [GH21]). In the future work we could address this issue by extending the protocol to specify how objectives are formed and selected. Finally, in the future work, we would like to prototype a physical distributed signalling system with the extended distributed resource allocation protocol could be implemented by using, for example the CLEARSY Safety Platform [LDFO20] as suggested in paper [Pro16].

Acknowledgements

This work is supported by an iCASE studentship and funded by EPSRC/UK and Siemens Rail Automation (grant number EP/P510580/1). This work was partially supported by the EPSRC STRATA platform grant (EP/N023641/1).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

- [Abr96] Abrial J R (1996) The B-book: assigning programs to meanings. Cambridge University Press, New York
- [Abr13] Abrial JR (2013) Modeling in Event-B: system and software engineering. Cambridge University Press, New York
- [Bac90] Back RJR (1990) Refinement calculus, part II: parallel and reactive programs. In: de Bakker JW, de Roever WP, Rozenberg G (eds) Stepwise refinement of distributed systems models, formalisms, correctness. Springer, Berlin, pp 67–93
- [BBFM99] Behm P, Benoit P, Faivre A, Meynadier JM (1999) Météor: a successful application of B in a large project. In Wing JM, Woodcock K, Davies J (eds) FM'99—formal methods. Springer, Berlin, pp 369–387
- [BSR80] Bernstein PA, Shipman DW, Rothnie JB Jr (1980) Concurrency control in a system for distributed databases (SDD-1). ACM Trans Database Syst 5(1):18–51
- [CM06] Cansell D, Méry D (2006) Formal and incremental construction of distributed algorithms: on the distributed reference counting algorithm. Theor Comput Sci 364(3):318–337
- [DFGV14] Déharbe D, Fontaine P, Guyot Y, Voisin L (2014) Integrating SMT Solvers in Rodin. Sci Comput Program 94(P2):130–143
- [ED06] Essamé D, Dollé D (2006) B in large-scale projects: the Canarsie line CBTC experience. In: Julliand J, Kouchnarenko O (eds) B 2007: formal specification and development in B. Springer, Berlin, pp 252–254
- [EGLT76] Eswaran Kapali P, Gray Jim, Lorie Raymond A, Traiger Irving L (1976) The Notions of Consistency and Predicate Locks in a Database System. Commun. ACM, 19(11):624–633
- [FH18] Fantechi A, Haxthausen AE (2018) Safety interlocking as a distributed mutual exclusion problem. In Howar F, Barnat J (eds) Formal methods for industrial critical systems. Springer, Cham, pp 52–66
- [FHN17] Fantechi A, Haxthausen AE, Nielsen MBR (2017) Model checking geographically distributed interlocking systems using UMC. In: 25th Euromicro international conference on parallel, distributed and network-based processing (PDP), pp 278–286
- [GH21] Geisler S, Haxthausen AE (2021) Stepwise development and model checking of a distributed interlocking system using RAISE. Formal Aspects Comput 33:87–125
- [GR92] Gray J, Reuter A (1992) Transaction processing: concepts and techniques, 1st edn. Morgan Kaufmann Publishers Inc., San Francisco
- [HHK+17] Hawblitzel C, Howell J, Kapritsos M, Lorch JR, Parno B, Roberts ML, Setty S, Zill B (2017) IronFleet: proving safety and liveness of practical distributed systems. Commun ACM 60(7):83–92
- [HKBA09] Hoang TS, Kuruma H, Basin D, Abrial JR (2009) Developing topology discovery in Event-B. Sci Comput Program 74(11):879–899
- [HKNP06] Hinton A, Kwiatkowska M, Norman G, Parker D (2006) PRISM: a tool for automatic verification of probabilistic systems. In: Hermanns H, Palsberg J (eds) Tools and algorithms for the construction and analysis of systems. Springer, Berlin, pp 441–444
- [Hoa14] Hoang TS (2014) Reasoning about almost-certain convergence properties using Event-B. Sci Comput Program 81:108–121
- [HP00] Haxthausen AE, Peleska J (2000) Formal development and verification of a distributed railway control system. IEEE Trans Softw Eng 26(8):687–701
- [ILTR11] Iliasov A, Laibinis L, Troubitsyna E, Romanovsky A (2011) Formal derivation of a distributed program in Event B. In: Qin S, Qiu Z (eds) Formal methods and software engineering. Springer, Berlin, pp 420–436
- [ISAYR16] Iliasov A, Stankaitis P, Adjepon-Yamoah D, Romanovsky A (2016) Rodin platform Why3 plug-in. In: Proceedings of the 5th international conference on abstract state machines, alloy, B, TLA, VDM, and Z, ABZ 2016. Springer, Berlin, pp 275–281
- [Pro16] INTO-CPS Project (2016) Deliverable D1.2-Case studies 2. Available at <https://into-cps.org/publications/>
- [LB03] Leuschel M, Butler M (2003) ProB: a model checker for B. In: Araki K, Gnesi S, Mandrioli D (eds) FME 2003: formal methods. Springer, Berlin, pp 855–874
- [LDFO20] Lecomte T, Déharbe D, Fournier P, Oliveira M (2020) The CLEARSY safety platform: 5 years of research, development and deployment. Sci Comput Program 199:102524
- [Mor96] Morley MJ (1996) Safety assurance in interlocking design. PhD thesis, University of Edinburgh. College of Science and Engineering. School of Informatics
- [New14] Newcombe C (2014) Why Amazon Chose TLA+. In: Ait Ameer Y, Schewe KD (eds) Abstract state machines, alloy, B, TLA, VDM, and Z. Springer, Berlin, pp 25–39
- [SDS+19] Stankaitis P, Dupont G, Singh NK, Ait-Ameer Y, Iliasov A, Romanovsky A (2019) Modelling hybrid train speed controller using proof and refinement. In: 2019 24th International conference on engineering of complex computer systems (ICECCS), pp 107–113
- [SI17] Stankaitis P, Iliasov A (2017) Theories, techniques and tools for engineering heterogeneous railway networks. In: Fantechi A, Lecomte T, Romanovsky A (eds) Reliability, safety, and security of railway systems. Modelling, analysis, verification, and certification. Springer, Cham, pp 241–250
- [SIA+19] Stankaitis P, Iliasov A, Ait-Ameer Y, Kobayashi T, Ishikawa F, Romanovsky A (2019) A refinement based method for developing distributed protocols. In: IEEE 19th international symposium on high assurance systems engineering (HASE), pp 90–97
- [SIK+20] Stankaitis P, Iliasov A, Kobayashi T, Ait-Ameer Y, Ishikawa F, Romanovsky A (2020) Formal distributed protocol development for reservation of railway sections. In: Raschke A, Méry D, Houdek F (eds) Rigorous state-based methods. Springer, Cham, pp 203–219
- [The06] The RODIN platform (2006). Available at https://sourceforge.net/projects/rodin-b-sharp/files/Core_Rodin_Platform/
- [WK12] Whitwam F, Kanner A (2012) Control of automatic guided vehicles without wayside Interlocking. Patent US 20120323411 A1

Received 21 January 2021

Accepted in revised form 30 September 2021 by Alessandro Fantechi, Anne Haxthausen and Jim Woodcock

Published online 24 November 2021