

Formal distributed protocol development for reservation of railway sections

Paulius Stankaitis¹, Alexei Iliasov¹, Tsutomu Kobayashi², Yamine Aït-Ameur³,
Fuyuki Ishikawa², and Alexander Romanovsky¹

¹ School of Computing, Newcastle University, Newcastle upon Tyne, UK

² National Institute of Informatics, Tokyo, Japan

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

Corresponding author: p.stankaitis@newcastle.ac.uk

Abstract. The decentralisation of railway signalling systems has the potential to increase railway network capacity, availability and reduce maintenance costs. 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. In this paper, we present a rigorous formal development and verification of a distributed protocol for reservation of railway sections, which we believe could deliver benefits of a decentralised signalling while ensuring safety and liveness properties. For the formal distributed protocol development and verification, we devised a multifaceted framework, which aims to reduce modelling and verification effort, while still providing complementary techniques to study protocol from all relevant perspectives.

Keywords: formal verification · distributed resource allocation · performance analysis · Event-B · PRISM model checker · railway signalling

1 Introduction

Railway signalling is a safety-critical system whose responsibility is to guarantee a safe and efficient operation of railway networks. In recent decades there have been proposals to utilize distributed system concepts (e.g. [13, 24]) in railway signalling as a way to increase railway network capacity and reduce maintenance costs. These emerging distributed railway signalling concepts propose using a radio-based communication technology to decentralise 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.

In recent years there has been a push (e.g. [12, 22]) by the industry with a strong focus on distributed systems to incorporate formal methods into their system development processes to improve system assurance and time-to-market.

[†]A single signalling computer may be responsible for controlling tens of routes (case studies [18, 20]) whereas novel distributed systems would reduce that number.

Yet, despite that for years the railway domain has proved to be a fruitful area for applying various formal methods [3, 7], considerably less has been done in applying them for distributed railway systems by industry and academia. Therefore, the long-term aim of our research is to lower the effort the barriers to applying formal methods in developing correct-by-construction distributed signalling systems.

In order to manage the modelling and verification complexity of distributed protocols we are working towards an integrated multifaceted methodology, which is based on three concepts: stepwise refinement, communication modelling patterns and validation through proofs. In spite of advancements in proof automation, it might be too onerous to mathematically prove the model in early development stages. Therefore, it is also desirable that the framework should support model animation and scenario validation. It is also paramount that the framework should support quantitative evaluation; as stated by Fantechi and Haxthausen [10], distributed signalling solutions will only be adopted in practice if system availability is demonstrated. The authors (as discussed in [10]) of related researches did not consider liveness and fairness properties, which directly affect system availability. In our proposed multifaceted methodology we integrate stochastic simulators for quantitative analysis.

In this paper, we present a research, which uses the proposed methodology to formally develop and verify a distributed railway signalling protocol, which would deliver decentralised signalling benefits, while meeting high safety requirements. The developed distributed signalling protocol is based on serialisability and is inspired by protocols used in transactions processing [4, 8, 11] in centralised and distributed database systems. The main objective of our protocol is to guarantee mutual exclusion of railway sections while ensuring systems liveness. In a nutshell, our key contributions are the formally proved distributed railway section allocation protocol inspired by past protocols for database systems and the formalisation of the multifaceted verification framework.

Related work. In Fantechi and Haxthausen [10] the authors formalise the railway interlocking problem as a distributed mutual exclusion problem and discuss the related literature on distributed interlocking (e.g. [9, 13, 24]). In principle all railway models share similar high-level safety, liveness and fairness requirements, as summarised on page 2 in [10]. One difference between our work and the studies overviewed in [10] 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 [1]. 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 paper by Morley [21] 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.

The rest of the paper is organised as follows. Section 2 outlines the motivation for developing the protocol, semi-formally describes its functionality, elicits the requirements and introduces its specifications and the properties to be proved. Section 3 further discusses the integrated methodology we are proposing. The following section briefly discusses formal model development and also provides technical details on property verification and performance analysis. In the last section we summarise our work and discuss future work directions.

2 Distributed Resource Allocation Model and Protocol

The distributed railway signalling can increase networks capacity (as trains could run closer), improve systems agility to delays and possibly reduce repair costs. On the other hand, an increased system complexity and a safety-critical (SIL4) nature requires the highest level of safety assurance. In order to apply formal methods one must clearly state system requirements and specifications. In the following subsections we describe an abstract model of the distributed railway system and its requirements as well as the `stage1` of the distributed protocol, which guarantees the safety and liveness of the distributed system.

2.1 High-level distributed system model and requirements

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 agents allocation order 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. 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.

SAF₁ | A resource will not be allocated to different agents at the same time.

SAF₂ | An agent will not use a resource until all requested resources are allocated.

LIV₁ | An agent must be eventually allocated requested set of resources.

LIV₂ | Resource allocation must be guaranteed in the presence of message delays.

Requirements 1: High-level systems safety and liveness requirements

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. The main high-level safety and liveness requirements of the distributed system are expressed in Requirements 1.

The following section attempts to justify the need for an adequate distributed protocol by discussing problematic distributed resource allocation scenarios.

2.2 Problematic distributed resource allocation scenarios

Let us consider Scenarios 1-2 (visualised in Figure 1) to see how requirement LIV_1 could not be guaranteed (while ensuring SAF_2) 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 Figure, 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, after 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 LIV_2 requirement as both agents would wait for the second *use* message to ensure SAF_2 .

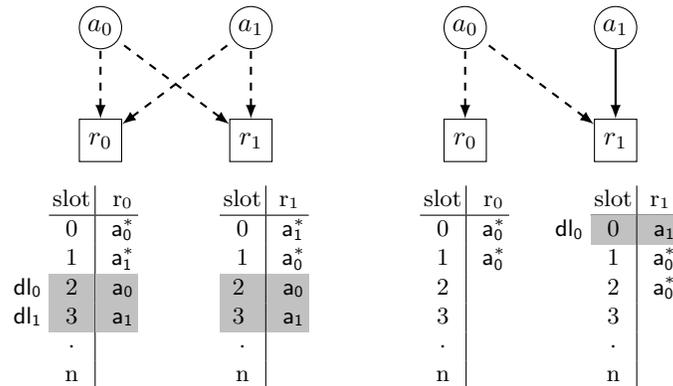


Figure 1: Problematic scenarios: Scenario 1 (left) and Scenario 2 (right)

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).

A distributed lane of agent a_n is $dl(a_n) = \{(r_k, s), (r_{k+1}, s), \dots, (r_{k+m}, s)\}$, where $\{r_k, r_{k+1}, \dots, r_{k+m}\}$ are all resources requested by agent a_n and s is the queue slot value promised by all requested resources. 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 values from both resources, and hence, will not be able to lock them.

To address the two issue described above, we developed a two-stage protocol, where the **stage₁** of the distributed protocol specifies how an agent forms a distributed lane. **Stage₂** of the protocol, which is out of this paper scope, addresses other deadlock scenarios, which can occur after agents form distributed lanes. In the following subsection we semi-formally describe the **stage₁** of the protocol.

2.3 Semi-formal description of the 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 agents 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 ($r_k \in \text{objective}[a_n]$) until $\text{sent_requests}[a_n] = \text{objective}[a_n]$ (images are equal). 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. 2-4 Algorithm 2). After sending all **request** messages an agent waits until **reply** messages are received from requested resources and then makes a decision.

When all received promised pointer values are the same (a distributed lane can be formed) an agent completes the **stage₁** by sending **write**, to all requested resources, messages which contain the negotiated index (ln. 14-17 Algorithm 1). But if one of the received promised pointer values is different an agent will start a renegotiation cycle (ln 5-13 Algorithm 1). By sending a **srequest** messages which

Algorithm 1 Agent stage₁ communication algorithm

```

1: while sent_requests[an] ≠ objective[an] do
2:   request(an) → rk   sending request message from agent an to resource rk
3: end
4: wait until received_replies[an] = objective[an]
5: while |replies[an] ≠ 1 do   cardinality of agents received slot indices
6:   m' = max(replies[an]) + 1
7:   sent_srequests[an]' = ∅   reset sent special request messages buffer
8:   received_replies[an]' = ∅   reset received reply messages buffer
9:   while sent_srequests[an] ≠ objective[an] do
10:    srequest(an, m) → rk
11:   end
12:   wait until received_replies[an] = objective[an]
13: end
14: while sent_write[an] ≠ objective[an] do
15:   m' = max(replies[an])
16:   write(an, m) → rk
17: end The end of stage1 of the protocol.

```

contain a desired slot index to resources. A desired index is computed by taking the maximum of all received promised pointer values and adding a constant (one is sufficient) - ln. 6 Algorithm 1. A resource will reply to `srequest` message with the higher value of the current `ppt(rk)` or received `srequest` message value and will update the promised pointer (ln. 5-7 Algorithm 2). After sending all `srequest` messages, an agent waits for `reply` messages and then restarts the loop if received slot indices are not the same.

Algorithm 2 Resource stage₁ communication algorithm

```

1: switch received_message do
2:   case request(an)
3:     reply(ppt(rk), rk) → an
4:     ppt(rk)' = ppt(rk) + 1
5:   case srequest(an, n)
6:     reply(max(ppt(rk), n), rk) → an
7:     ppt(rk)' = max(ppt(rk), n) + 1

```

SAF₃ | An agent will not send `write` (form a distributed lane) messages until all receive promised pointer values are identical.

SAF₄ | Agents with overlapping resource objectives will negotiate distributed lanes with different index.

LIV₃ | An agent will eventually negotiate a distributed lane.

Requirements 2: Low-level protocol stage₁ safety and liveness requirements

It is important to note that the stage₁ protocol solution to the described deadlock scenarios has a stochastic nature and one needs to guarantee that a desirable state is probabilistically reachable. In Requirements 2 we summarise requirements for the stage₁ of the protocol.

After an agent completes **stage₁** and thus negotiates a distributed lane it will start protocol **stage₂** to prevent other deadlock scenarios. Predominantly because of papers verification focus towards properties from **stage₁** (all complimentary verification/analysis techniques used) we provide protocol **stage₂** description in the online appendix[§].

3 Multifaceted Modelling and Verification Framework

As stated before, the long-term objectives of our research are to reduce modelling and verification effort of distributed systems and to have a multifaceted framework to study protocols from all relevant perspectives. In the introduction, we defined key formal concepts the framework should rely on and in the following section we discussed protocol requirements we need to guarantee.

The following subsections proposes an engineering process with different formal techniques each of which is efficient to handle parts of above requirements and help to manage modelling and verification complexity.

3.1 Formalised multifaceted verification framework

For any adequate formal system development, system requirements should be clearly stated, and so, this is the first step (Step 1 in Fig. 2) in the modelling process. Currently, we do not suggest or provide a specific structural approach for defining distributed system requirements. The next step (Step 2) in the process is developing and verifying a pivotal formal model. The purpose of formally modelling a distributed system is to have a formal artefact, which can be animated, analysed and formally verified.

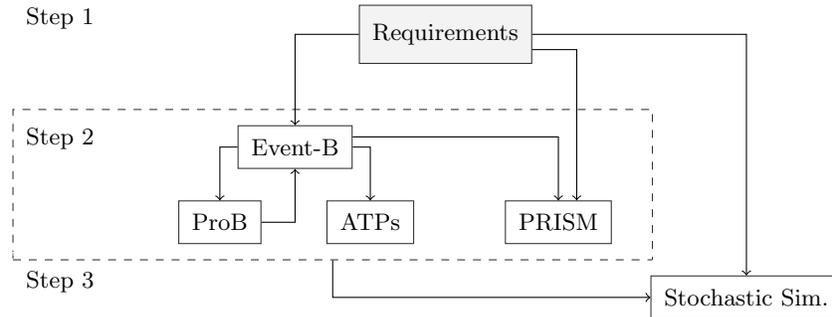


Figure 2: Multifaceted modelling and verification framework

For the development and verification of pivotal functional system models we selected the Event-B [2] specification language, which has previously been

[§]A complete protocol description and formal models can be found at <http://stankaitis.uk/2019/02/>.

successfully used for modelling and verification of various distributed protocols [5, 15, 16]. The Event-B method provides an expressive modelling language, flexible refinement mechanism and is also proof driven, meaning model correctness is demonstrated by generating and discharging proof obligations with available automated theorem provers [6, 17]. The method is supported by tools such as ProB [19] which enable animating and model-checking a model. On the other hand, the Event-B method does not have an adequate probabilistic reasoning support, which, for example, was essential for verifying the distributed railway section reservation protocol. Therefore, it was decided to integrate the well-known PRISM [14] stochastic model checker into the framework, so stochastic system's properties can be verified.

The last step (Step 3) in the proposed engineering process is analysing a developed distributed system's performance. For that, we have implemented a high-fidelity protocol simulator which could help to evaluate protocols under normal or *stressed* conditions. Following subsections provide more detail on how each of the formal techniques would be used in the development and verification of a distributed protocol.

3.2 Step 2: Developing functional pivot 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 propose is a rather standard and starts with the abstract model 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 abstract model is then iteratively refined by introducing more details about the distributed protocol, primarily by modelling communication aspects. To reduce modelling effort we previously developed communication modelling patterns and described a generic model refinement plan in [23]. 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.

3.3 Step 2: Proving stochastic properties with PRISM

As the distributed signalling protocol had a stochastic nature it was important to formally demonstrate that a satisfying state could be reached. Probabilistic or liveness properties are hard to formalise and prove in the Event-B method.

Therefore, it was decided to prove progress of the protocol outside of Event-B by redeveloping part of the model (`stage1`) in the PRISM model checker.

The drawback of using PRISM model checker, if a bounded problem abstraction cannot be found, the verification is limited to bounded models. As we could not find protocol's `stage1` abstraction, we created a skeleton model, which then could be instantiated to model specific scenarios of `stage1` with n agents, m resources and other initial conditions. Additionally, we developed a model generator, which can automatically instantiate the skeleton model to capture a random scenario and run probabilistic verification conditions.

3.4 Step 3: Analysing system's performance

With Event-B and PRISM we aim to demonstrate that the protocol addresses the formulated requirements but it is necessary in our application domain to understand how the protocol is going to perform under various conditions if it were deployed in a real system. To conduct such a simulation we have implemented a high fidelity protocol simulator that can be populated with any number of resources and agents while realising any conceivable agents' goal formation and message delivery policies.

The simulator is parametrised with a function of probability of picking a certain message out of a pool of available messages. The probability function is itself parametrised by message source, destination, timestamp and type. The simulation would help to answer how fast, in terms of vital steps such as messages sent, a protocol's `stage1` can be completed and how the performance is affected by messages delays. With function D we can simulate slow agents and resources, fair, arbitrary and unfair delivery policies, agents that operate much faster than others and so on.

4 Formal Protocol Modelling, Verification and Analysis

In this section we present the application of previously introduced modeling and verification framework for developing distributed railway signalling protocol. In Section 2 we defined protocol's requirements (Step 1), thus following subsections focuses on formal methodology aspects.

4.1 Step 2. Formal protocol model development in 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 3.2. Important to note that the protocol model was redeveloped multiple times as various deadlock scenarios were found with ProB animator and model-checker. Below, we overview the final (verified) model.

Modelling was started by creating an abstract model context which contains constants, given sets and uninterpreted functions. In the abstract context, we introduced three (finite) sets, to respectively represent agents (`agt`), resources (`res`) and objectives (`obj`). The context also contains an objective function which

is a mapping from objectives to a collection of resources ($\text{ob} \in \text{obj} \rightarrow \mathbb{P}(\text{res})$) and an enumerated set for agents status counter.

The dynamic protocol parts, such as messages exchanges, are modelled as variables and events computing next variable states and contained in a *machine*. According to the proposed model development process, the initial machine (abstract) should summarise the objective of protocol, which is an agent completing an objective (locking all necessary resources). To capture that, the abstract protocol machine contains two events, respectively modelling an agent locking and then releasing a free objective ($\text{ob} \in \text{obj}$). The abstract model is refined by mostly modelling communication aspects of the distributed signalling protocol and for that we use a backward unfolding style where the next refinement step introduces preceding protocol step. Below, we overview the refinement chain and properties we proved at that modelling stage.

Refinement 1 (Abstract ext.) In this refinement we introduce resources into the model and now an agent tries to fulfill the objective by locking resources. Previous two events (lock/release) are now decomposed to two for each and capture iterative locking and releasing of resources.

Refinement 2 The abstract models are firstly refined with stage_2 part of the protocol. In the refinement, r_2 , we introduced `lock`, `response` and `release` messages and associated events into the model. In this step we also demonstrated that the protocol stage_2 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 Model r_3 , is the bridge between protocol stages stage_1 and stage_2 and introduces two new messages `write` and `pready` into the model.

Refinement 4 The final refinement step - r_4 - models stage_1 of the distributed protocol which is responsible for creating distributed lanes. 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. SAF_{3-4}).

4.2 Step 2: Proving functional correctness properties in Event-B

As shown in Section 2.2 (Scenarios 1 - 2) high-level system's requirements can only be met if an agent invariably and correctly forms a distributed lane. The probabilistic lane forming eventuality (LIV_3) is discussed separately while in the following paragraphs we focus on the proof regarding requirements SAF_{3-4} .

SAF_3 is required to ensure that agent's resource objectives are not satisfied or satisfied on full. The model addresses this via event *guards* restricting 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. For illustrative purposes we focus on details of verifying a slightly more interesting case of SAF_4 and assume that SAF_3 is proven.

Requirement SAF_4 addresses potential cross-blocking deadlocks or resource double locking due to distributed lane overriding. The strategy is to prove the

requirement is to show that agents that are interested in at least one common resource (related) always form distributed lanes with differing indices. We start by assuming that agents only form distributed lanes if all received indices are the same (proved as SAF_3). Then, if a resource (or resources) shared between any two related agents send unique promised pointer values to these agents, these indices will be distributed lane *deciders* as all other indices from different resources must be the same to form a distributed lane. Hence, to prove SAF_4 it is enough to show that each resource replies to a request or special request message with a unique promised pointer value.

```

resource_reply_general  $\hat{=}$ 
ANY
    rq, rp
WHERE
    grd1  rq  $\in$  req  take a sent request message
    grd2  rp  $\in$  REQ \ rep  create a new reply message
    grd3  repd(rp) = reqs(rq)  destination of reply message is source of request message
    grd4  reps(rp) = reqd(rq)  source of reply message is destination of request message
    grd5  repn(rp) = ppt(reps(rp))  reply message contains promised pointer
THEN
    act1  rep := rep  $\cup$  {rp}  add new message to reply channel
    act2  req := req \ {rq}  remove request message from request channel
    act3  ppt(res) := ppt(res) + 1  increment promised pointer
    act4  hisppt(res) := hisppt(res)  $\Leftarrow$  {(hiswr(res))  $\mapsto$  ppt(res)}
    act5  hiswr(res) := hiswr(res) + 1
END
    
```

Figure 3: Event-B model excerpt of a resource sending a reply message

To prove that all resources replies to a request or special request message with a unique promised pointer value, we firstly introduced a history variable his_{ppt} of type $\text{his}_{\text{ppt}} \in (\text{res} \rightarrow (\mathbb{N} \leftrightarrow \mathbb{N}))$ into our model. The main idea behind the history variable was to chronologically store the promised pointer values sent by a resource. We also introduced a time-stamp variable his_{wr} of the type $\text{his}_{\text{wr}} \in \text{res} \rightarrow \mathbb{N}$ to chronologically order the promised pointer values stored in the history variable.

After introducing history variables, we modified events `resource_reply_general` and `resource_reply_special`, which in the protocol update the promised pointer variables, by adding two new actions (see Figure 3). The first action `act4` updates the history variable with the promised pointer value ($\text{ppt}(\text{res})$) that was sent to the agent at the time stamp ($\text{his}_{\text{wr}}(\text{res})$). The second action, `act5`, simply increments resource's res time-stamp ($\text{his}_{\text{wr}}(\text{res})$) variable.

$$\text{inv_saf_4} \quad \forall r, n_1, n_2 \cdot r \in \text{RES} \wedge n_1, n_2 \in \text{dom}(\text{his}_{\text{ppt}}(r)) \wedge n_1 < n_2 \Rightarrow \text{his}_{\text{ppt}}(r)(n_1) < \text{his}_{\text{ppt}}(r)(n_2)$$

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 (**inv_saf_4**) 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.

$$\begin{aligned} \mathbf{inv_his_ppt} \quad \forall \mathbf{res}. \quad & (\mathbf{his_wr}(\mathbf{res}) = 0 \wedge \mathbf{his_ppt}(\mathbf{res}) = \emptyset) \\ & \vee (\text{dom}(\mathbf{his_ppt}(\mathbf{res})) = 0 .. \mathbf{his_wr}(\mathbf{res}) - 1 \\ & \wedge \mathbf{his_ppt}(\mathbf{res})(\mathbf{his_wr}(\mathbf{res}) - 1) = \mathbf{ppt}(\mathbf{res}) - 1) \end{aligned}$$

To prove that **resource_reply_{general, special}** preserve **inv_saf_4**, the following properties play the key role: (1) the domain of **his_ppt** (i.e., ‘indices’ of **his_ppt**) is $\{0, \dots, \mathbf{his_wr} - 1\}$, (2) $\mathbf{his_ppt}(\mathbf{his_wr} - 1) < \mathbf{his_ppt}(\mathbf{his_wr})$. Property (2) holds because **his_ppt(his_wr)** 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 (**inv_his_ppt**) and proved they are preserved by the events which helped to prove **inv_saf_4**.

Proof statistics. In Table 1 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 [23] use and SMT-based verification support [6, 17].

Model	No. of POs	Aut. Discharged	Int. Discharged
context c₀	0	0	0
context mes.	9	9	0
machine m₀	12	12	0
machine m₁	23	21	2
machine m₂	59	43	16
machine m₃	43	32	11
machine m₄	103	57	46
Total	249	174	75

Table 1: Event-B protocol model proof statistics

4.3 Step 2: Proving liveness (req. **LIV₃**) with PRISM

In this subsection, we discuss stochastic model checking results with which we intend to prove level that **LIV₃** requirement is preserved. In particular, we focus on showing that **LIV₃** requirement is ensured in Scenario 2 (Section 2.2).

In order to demonstrate that **LIV₃** requirement holds in Scenario 2 (Section 2.2) we used **stage₁** protocol’s skeleton PRISM model to replicate Scenario 2. In this experiment we were interested in observing the effects a promised pointer offset has on an probability of agent forming a distributed lane while

the upper limit of the promised pointer is increased[†] (n in Scenario 2). Early experiments showed that verification would not scale well (several hours for a single data-point) if we would increase the number of resources and agents above two resources and three agents (each agent trying to reserve both resources) so we kept these parameters constant.

For each scenario, we would run a quantitative property: $P = ? [F \text{ dist}_0 > -1]$ which asks what is the probability of an agent negotiating a distributed lane until the upper promised pointer limit is reached. The three curves (red, green and violet) in Figure 3 show the effect a promised pointer offset has on negotiation probability as queue depth is increased. Results suggest that increasing the offset reduces the probability of negotiating a distributed lane as queue depth is increased, but the probability still approaches one as the number of rounds is increased.

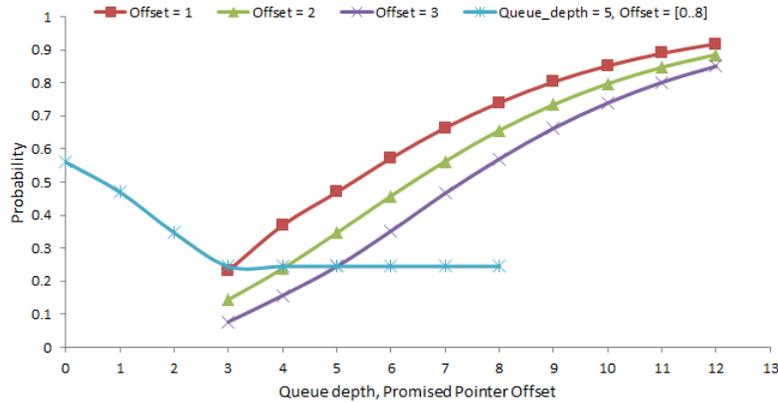


Figure 4: Scenario 2 with varied resource promised pointer offset and queue depth.

To further see the effects of the offset, we considered a different experiment where the same quantitative property would be run when the number of possible renegotiations value is kept constant and offset is increased (light blue plot). Results indicate that offset has only effect until a specific threshold and after that the probability of agent negotiating a distributed lane is not affected by the offset. These results suggest that the situation in Scenario 2 does not violate LIV_3 requirement as distributed lanes can be negotiated.

4.4 Step 3: Analysing performance

The goal of this part is to study the protocol performance under various stress conditions and thus provide assurances of its applicability in real life situations. To build simulation, we simply capture protocol’s stage_1 behaviour using a program. We are also able to obtain bounds on the number of messages required to

[†]Instead, of ppt upper limit we decided plot the probability against the queue depth, ($\text{offset} - n$) as it directly shows how many times an agent can renegotiate resources.

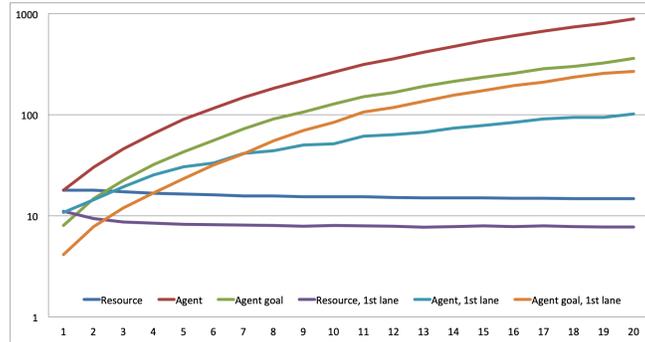


Figure 5: Time to form all or first lanes, logarithmic scale.

form lanes in different setups. This can be directly translated into real-life time bounds on the basis of point to point transmission times.

Simulation construction. Simulation is setup as a collection of actors of two types - agents and resources - and an orchestration component observing and recording message passing among the actors. A message is said to be in transit as soon as it is created by an actor. Every act of message receipt (and receipt only) advances the simulation (world) clock by one unit. Hence, any number of computations leading to message creation can occur in parallel but message delivery is sequential. To model delays we define a function that probabilistically picks a message to be delivered among all the messages currently in transit. A special message, called skip, is circulated to simulate idle passage of time. This message is resent immediately upon receipt by an implicit idle actor.

Let \mathbb{M} be set of all messages that can be generated by agents and resources. Also, let $\text{skip} \notin \mathbb{M}$ denote the skip message and $\mathbb{M}' = \mathbb{M} \cup \{\text{skip}\}$. By its structure, set \mathbb{M}' is countable (each message identified by unique integer) and one can define a measure space over \mathbb{M}' . Let D signify the probability that some message $m \in M \subseteq \mathbb{M}'$ from message pool M is selected for reception. We shall define D via the current message pool, the attributes of m such its source, destination, time stamp and protocol stage, and the world time: $D = D(M, m, t) = D(M, m.s, m.d, m.c, m.o, t)$. Here M is the set of available message, $m.s$ and $m.d$ are the message source and destination agent or resource, $m.c$ is the message type (e.g., WRITE), $m.o$ is the message timestamp (the point of its creation) and t is the world clock. Defining differing probabilities D we are able to address most scenarios of interest.

Uniform distribution. With $D(M, m, t) = \text{card}(M)^{-1}$ the simulator picks a message from M using a uniform distribution. It is an artificial setting as the time in transit bears no influence over the probability of arrival. Counter-intuitively, the said probability may decrease with the passage of time when new messages are created quicker than they are delivered. The skip message has equal probability with the rest so the system "speeds up" when M is large. The plots in Fig. 5 shows how the protocol performance changes when the number of resources (Resource line), agents (Agent lines), and resources an agent attempts

to acquire (Agent goal) increase. We plot separately time to form all lanes and any first lane. The values plotted are averaged over 10000 runs.

5 Conclusions and Future Work

In this paper we proposed a multifaceted framework with which we aim to reduce modelling and verification of distributed (railway signalling) systems. The framework was applied in the development of the novel distributed signalling protocol. Starting only with high-level system requirements we developed an early formal protocol prototype which with the help of ProB was refined as subtle deadlock scenarios were discovered. This in part is the advantage of a stepwise development supported by Event-B as complex distributed models can be *decomposed* into smaller problems and errors found earlier. The stepwise distributed protocol development as also shown before [5, 15, 16] together with adequate tools [6, 17] helped to achieve fairly high verification automation. On the other hand, protocol verification was complicated by the need of stochastic reasoning and not adequate Event-B support for reasoning about probabilistic properties. The current solution relied on a model redevelopment in stochastic model checker PRISM which did not scale well for verification of larger scenarios. As a future direction it is essential to address this problem by most likely improving stochastic reasoning in Event-B. In the future we would also like to a much closer tool integration and support an automatic translation to PRISM and the stochastic simulator.

References

1. INTO-CPS Project. Case Studies 2, Deliverable D1.2. Available at: http://projects.au.dk/fileadmin/D1.2a_Case_Studies.pdf. Tech. rep. (11 2016)
2. Abrial, J.R.: Modeling in Event-B: System and Software Engineering. Cambridge University Press, New York, NY, USA (2013)
3. Behm, P., Benoit, P., Faivre, A., Meynadier, J.M.: Météor: A Successful Application of B in a Large Project. In: Wing, J.M., Woodcock, J., Davies, J. (eds.) FM'99 — Formal Methods. pp. 369–387. Springer, Berlin, Heidelberg (1999)
4. Bernstein, P.A., Shipman, D.W., Rothnie, Jr., J.B.: Concurrency Control in a System for Distributed Databases (SDD-1). ACM Trans. Database Syst. **5**(1), 18–51 (Mar 1980)
5. Cansell, D., Méry, D.: Formal and Incremental Construction of Distributed Algorithms: On the Distributed Reference Counting Algorithm. Theor. Comput. Sci. **364**(3), 318–337 (Nov 2006)
6. Déharbe, D., Fontaine, P., Guyot, Y., Voisin, L.: Integrating SMT Solvers in Rodin. Sci. Comput. Program. **94**(P2), 130–143 (Oct 2014)
7. Essamé, D., Dollé, D.: B in Large-Scale Projects: The Canarsie Line CBTC Experience. In: Julliand, J., Kouchnarenko, O. (eds.) B 2007: Formal Specification and Development in B. pp. 252–254. Springer Berlin Heidelberg (2006)
8. Eswaran, K.P., Gray, J., Lorie, R.A., Traiger, I.L.: The Notions of Consistency and Predicate Locks in a Database System. Commun. ACM **19**(11), 624–633 (1976)

9. Fantechi, A., Haxthausen, A.E., Nielsen, M.B.R.: Model Checking Geographically Distributed Interlocking Systems Using UMC. In: 25th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP). pp. 278–286 (March 2017)
10. Fantechi, A., Haxthausen, A.E.: Safety Interlocking as a Distributed Mutual Exclusion Problem. In: Howar, F., Barnat, J. (eds.) Formal Methods for Industrial Critical Systems. pp. 52–66. Springer International Publishing, Cham (2018)
11. Gray, J., Reuter, A.: Transaction Processing: Concepts and Techniques. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edn. (1992)
12. Hawblitzel, C., Howell, J., Kapritsos, M., Lorch, J.R., Parno, B., Roberts, M.L., Setty, S., Zill, B.: IronFleet: Proving Safety and Liveness of Practical Distributed Systems. *Commun. ACM* **60**(7), 83–92 (Jun 2017)
13. Haxthausen, A.E., Peleska, J.: Formal development and verification of a distributed railway control system. *IEEE Transactions on Software Engineering* **26**(8), 687–701 (Aug 2000)
14. Hinton, A., Kwiatkowska, M., Norman, G., Parker, D.: 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. pp. 441–444. Springer Berlin Heidelberg (2006)
15. Hoang, T.S., Kuruma, H., Basin, D., Abrial, J.R.: Developing Topology Discovery in Event-B. *Science of Computer Programming* **74**(11), 879 – 899 (2009)
16. Iliasov, A., Laibinis, L., Troubitsyna, E., Romanovsky, A.: Formal Derivation of a Distributed Program in Event B. In: Qin, S., Qiu, Z. (eds.) Formal Methods and Software Engineering. pp. 420–436. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
17. Iliasov, A., Stankaitis, P., Adjepon-Yamoah, D., Romanovsky, A.: Rodin Platform Why3 Plug-In. In: Proceedings of the 5th International Conference on Abstract State Machines, Alloy, B, TLA, VDM, and Z. pp. 275–281. ABZ 2016, Springer-Verlag, Berlin, Heidelberg (2016)
18. Iliasov, A., Taylor, D., Laibinis, L., Romanovsky, A.: Formal Verification of Signalling Programs with SafeCap. In: Computer Safety, Reliability, and Security. pp. 91–106. Springer Berlin Heidelberg (2018)
19. Leuschel, M., Butler, M.: ProB: A Model Checker for B. In: Araki, K., Gnesi, S., Mandrioli, D. (eds.) FME 2003: Formal Methods. pp. 855–874. Springer Berlin Heidelberg (2003)
20. Limbrée, C., Cappart, Q., Pecheur, C., Tonetta, S.: Verification of Railway Interlocking - Compositional Approach with OCRA. In: Lecomte, T., Pinger, R., Romanovsky, A. (eds.) Reliability, Safety, and Security of Railway Systems. Modelling, Analysis, Verification, and Certification. pp. 134–149. Springer International Publishing, Cham (2016)
21. Morley, M.: Safety assurance in interlocking design. Ph.D. thesis, University of Edinburgh. College of Science and Engineering. School of Informatics. (1996)
22. Newcombe, C.: Why Amazon Chose TLA+. In: Ait Ameer, Y., Schewe, K.D. (eds.) Abstract State Machines, Alloy, B, TLA, VDM, and Z. pp. 25–39. Springer Berlin Heidelberg (2014)
23. Stankaitis, P., Iliasov, A., Ait-Ameer, Y., Kobayashi, T., Ishikawa, F., Romanovsky, A.: A Refinement Based Method for Developing Distributed Protocols. In: IEEE 19th International Symposium on High Assurance Systems Engineering (HASE). pp. 90–97 (Jan 2019)
24. Whitwam, F., Kanner, A.: Control of Automatic Guided Vehicles Without Wayside Interlocking. Patent US 20120323411 A1 (2012)