



HAL
open science

Moving from Event-B to probabilistic Event-B

Mohamed Amine Aouadhi, Benoit Delahaye, Arnaud Lanoix

► **To cite this version:**

Mohamed Amine Aouadhi, Benoit Delahaye, Arnaud Lanoix. Moving from Event-B to probabilistic Event-B. [Research Report] LINA-University of Nantes. 2016. ⟨hal-01316610v2⟩

HAL Id: hal-01316610

<https://hal.science/hal-01316610v2>

Submitted on 24 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Moving from Event-B to Probabilistic Event-B

Mohamed Amine Aouadhi

Université de Nantes / LINA

Nantes, France

mohamed-amine.aouadhi@univ-nantes.fr

Benoît Delahaye

Université de Nantes / LINA

Nantes, France

benoit.delahaye@univ-nantes.fr

Arnaud Lanoix

Université de Nantes / LINA

Nantes, France

arnaud.lanoix@univ-nantes.fr

Abstract

We propose a fully probabilistic extension of Event-B where all the non-deterministic choices are replaced with probabilities. We present the syntax and the semantics of this extension and introduce novel and adapted proof obligations for proving the correctness of probabilistic Event-B models. As a preliminary step towards handling refinement of probabilistic Event-B models, we propose sufficient conditions for the almost-certain convergence of a set of events and express them in terms of proof obligations. We illustrate our work by presenting a case study specified in both standard and probabilistic Event-B.

1 Introduction

As systems become more and more complex, with randomised algorithms [18], probabilistic protocols [3] or failing components, it is necessary to add new modelling features in order to take into account complex system properties such as reliability [23], responsiveness [7, 22], continuous evolution, energy consumption etc.

In this way, several research works have focused on the extension of Event-B to allow the expression of probabilistic information in Event-B models. Event-B [1] is a formal method used for discrete systems modelling. It is equipped with *Rodin* [2], an open toolset for modelling and proving systems. The development process in Event-B is based on refinement: systems are typically developed progressively using an ordered sequence of models, where each model contains more details than its predecessor.

In [16], Morgan *et al.* have summarised the difficulties of embedding probabilities into Event-B. This seminal paper suggests that probabilities need to be introduced as a refinement of *non-determinism*. In Event-B, non-determinism occurs in several places such as the choice between enabled events in a given state, the choice of the parameter values in a given event, and the choice of the value given to a variable through some non-deterministic assignments. The ideal probabilistic extension of Event-B should therefore allow using probabilities in all these places. To the best of our knowledge, the existing works on extending Event-B with probabilities have mostly focused on refining non-deterministic assignments into probabilistic assignments. In [8], Hallerstede and Hoang propose to focus on a qualitative aspect of probability. They propose to refine non-deterministic assignments into *qualitative* probabilistic assignments where the actual probability values are not specified, and adapt the Event-B semantics and proof obligations to this new setting. In [24], Yilmaz study the refinement of qualitative probabilistic Event-B models and propose a tool support inside Rodin. Other works [19, 21, 20] have extended this approach by refining non-deterministic assignments into *quantitative* probabilistic assignments where, unlike in [8], the actual probability values are specified. This new proposition is then exploited in order to assess several system properties such as reliability and responsiveness.

Unfortunately, other sources of non-determinism than assignments have been left untouched in these works. In [16], the authors argue that probabilistic choice between events or parameter values can be achieved by transformations of the models that embed these choices inside probabilistic assignments. While this is unarguably true, such transformations are not trivial and greatly impede the understanding of Event-B models. Moreover, these transformations would need to be included in the refinement chain when designers need it, which would certainly be counter-intuitive to engineers.

Instead, we propose a different approach in which probabilistic choices can be introduced as a refinement of any potential non-deterministic choice, be it between enabled events, parameter values or assignments. Our long-term goal is to produce a probabilistic extension of Event-B where probabilistic events/parameters/assignments can be introduced natively either as standalone modelling artifacts or as a refinement of their non-deterministic counterparts. This long-term goal is clearly ambitious and will require several years of study to be achieved.

As a first step towards this long-term objective, we consider a slightly simplified modelling process where the engineer introduces probabilities in the last refinement step of a model, when the system is already sufficiently detailed. For now, we also restrict ourselves to purely probabilistic systems: when probabilities are introduced in the model, they replace *all* non-deterministic choices. We therefore propose a *fully probabilistic* extension of Event-B where *all* non-deterministic choices are replaced with probabilistic ones. As for standard Event-B models, the consistency of probabilistic Event-B models is expressed in terms of proof obligations. We therefore introduce new proof obligations dedicated to the consistency of probabilistic Event-B models and explain how standard Event-B proof obligations can be adapted to the probabilistic setting. In order to prove the correctness of our approach, we show that the semantics of a probabilistic Event-B model is a (potentially infinite-state) discrete time Markov chain.

As explained in [16], ensuring the refinement of Event-B models where probabilistic choice is not reduced to assignments is a difficult problem. While we do not solve this problem in its entirety, we take a preliminary step towards this goal by providing sufficient conditions, expressed in terms of proof obligations, for the almost-certain convergence of a set of events. Convergence is a required property in standard Event-B for proving refinement steps as soon as new events are introduced in the model. Almost-certain convergence has already been studied in [8] and [11], in the context of non-deterministic models with probabilistic assignments, but we show that the proof obligations developed in this context are not sufficient for models where probabilistic choice also appears in the choice of events and parameter valuations.

Finally, we illustrate our work on a classical case study: the emergency brake system. In particular, we show that some of the requirements provided in this case study cannot be taken into account using standard Event-B while their specification using probabilistic Event-B is intuitive, in particular when probabilities can be taken into account for the choice between enabled events.

All the results we present in this paper are being implemented in a prototype plugin for Rodin, which we briefly present at the end of the paper.

Related work. As said above, [16] is a seminal paper that identifies the challenges when considering a probabilistic extension for Event-B: introducing probabilities in the three places where non-determinism appears in standard Event-B (between enabled events, on the parameter values choice and on non-deterministic assignments) is the major difficulty to preserve the practicality of Event-B, and in particular the refinement development framework. Existing works such as the book of Morgan and McIver [17] partially answer these challenges. In particular, [17] introduces a probabilistic refinement calculus where the refinement of probabilistic guarded commands (\equiv assignments) is worked out. The PhD work of Hoang [15, 13, 10] adapts these results to the classical B-Method and the underlying guarded substitution language. Following these works, Hallerstede and Hoang [8] have proposed a first probabilistic extension of Event-B, where probabilities are introduced as a refinement of non-determinism in non-deterministic assignments. In this first extension, Hallerstede and Hoang focus on a *qualitative* aspect of probabilities and adapt the Event-B semantics and proof obligations to this new setting. *Quantitative* probabilistic assignments are then introduced in [19, 21, 20]. Almost-certain convergence of a set of events is studied in [8] and [11] in this context of non-deterministic models with probabilistic assignments. Event-B refinement of such models is studied in [24].

Unfortunately, while the refinement method and underlying theory introduced in these works is undoubtedly a good starting point for addressing refinement of fully probabilistic Event-B models, none of these works can be directly adapted to our setting as they do not consider probabilistic choice between enabled events or in the choice of parameter values. We show in Section 6 that the proof obligations dedicated to almost-certain convergence for non-deterministic systems with probabilistic assignments cannot be adapted to our setting and that additional proof obligations are required. The inherent complexity of introducing probabilistic choice between events and for parameter valuations is such that the introduction of a functional refinement procedure in our setting is still out of our reach, although we are making progress as Section 6 shows.

Outline. The paper is structured as follows. Section 2 presents an overview of the Event-B method and of our running case study. In Section 3, we introduce the syntax of fully probabilistic Event-B and illustrate our approach on the running case study. Section 4 presents new and modified proof obligations for the consistency of probabilistic Event-B models. The semantics of a fully probabilistic Event-B model is described in Section 5 and

Section 6 treats the almost-certain convergence of fully probabilistic Event-B models. Finally, Section 7 concludes and presents hints for future work. For space reasons, the full proofs of our results as well as additional material and examples are presented in a separate appendix, to be consulted at the discretion of the reviewers.

2 Event-B

We first present the basic elements of the Event-B method and then introduce our running case study.

2.1 Preliminaries

Event-B [1] is a formal method used for the development of complex systems. Systems are described in Event-B by means of models. For the sake of simplicity, we assume in the rest of the paper that an Event-B model is expressed by a tuple $M = (\bar{v}, I(\bar{v}), V(\bar{v}), \text{Evts}, \text{Init})$ where $\bar{v} = \{v_1 \dots v_n\}$ is a set of variables, $I(\bar{v})$ is an invariant, $V(\bar{v})$ is an (optional) variant used for proving the convergence of the model, Evts is a set of events and $\text{Init} \in \text{Evts}$ is an initialisation event. The invariant $I(\bar{v})$ is a conjunction of predicates over the variables of the system specifying properties that must always hold.

Events. An event has the following form:

$\text{event } e_i \hat{=} \text{any } \bar{t} \text{ where } G_i(\bar{t}, \bar{v}) \text{ then } S_i(\bar{t}, \bar{v}) \text{ end}$
--

where e_i is the name of the event, $\bar{t} = \{t_1 \dots t_n\}$ represents the set of parameters of the event, $G_i(\bar{t}, \bar{v})$ is the guard of the event and $S_i(\bar{t}, \bar{v})$ is the action of the event. An event is *enabled* in a given valuation of the variables (also called a configuration) if and only if there exists a parameter valuation such that its guard $G_i(\bar{t}, \bar{v})$ is satisfied in this context. Parameters and guards are optional. The action $S_i(\bar{t}, \bar{v})$ of an event may contain several assignments that are executed in parallel. Assignments can be written in several forms in standard Event-B, but they can always be reduced to what we call *Predicate assignments* in the following. In order to simplify the writing of (Probabilistic) Event-B programs, we only distinguish the following three forms of assignments:

- *Deterministic assignment:* $x := E(\bar{t}, \bar{v})$ means that the expression $E(\bar{t}, \bar{v})$ is assigned to the variable x .
- *Predicate assignment:* $x :| Q_x(\bar{t}, \bar{v}, x')$ means that the variable x is assigned a new value x' such that the predicate $Q_x(\bar{t}, \bar{v}, x')$ is satisfied.
- *Enumerated assignment:* $x \in \{E_1(\bar{t}, \bar{v}) \dots E_n(\bar{t}, \bar{v})\}$ means that the variable x is assigned a new value taken from the set $\{E_1(\bar{t}, \bar{v}) \dots E_n(\bar{t}, \bar{v})\}$.

Before-after predicate. The formal semantics of an assignment is described by means of a before-after predicate (BAP) $Q_x(\bar{t}, \bar{v}, x')$, which describes the relationship between the values of the variable before (x) and after (x') the execution of an assignment.

- The BAP of a deterministic assignment is $x' = E(\bar{t}, \bar{v})$.
- The BAP of a predicate assignment is $Q_x(\bar{t}, \bar{v}, x')$.
- The BAP of an enumerated assignment is $x' \in \{E_1(\bar{t}, \bar{v}) \dots E_n(\bar{t}, \bar{v})\}$.

Recall that the action $S_i(\bar{t}, \bar{v})$ of a given event may contain several assignments that are executed in parallel. Assume that $v_1 \dots v_i$ are the variables assigned in $S_i(\bar{t}, \bar{v})$ – variables $v_{i+1} \dots v_n$ are thus not modified – and let $Q_{v_1}(\bar{t}, \bar{v}, v'_1) \dots Q_{v_i}(\bar{t}, \bar{v}, v'_i)$ be their corresponding BAP. Then the BAP $S_i(\bar{t}, \bar{v}, \bar{v}')$ of the event action $S_i(\bar{t}, \bar{v})$ is:

$S_i(\bar{t}, \bar{v}, \bar{v}') \hat{=} Q_{v_1}(\bar{t}, \bar{v}, v'_1) \wedge \dots \wedge Q_{v_i}(\bar{t}, \bar{v}, v'_i) \wedge (v'_{i+1} = v_{i+1}) \wedge \dots \wedge (v'_n = v_n)$
--

Proof obligations. The consistency of a standard Event-B model is characterised by *proof obligations* (POs) which must be discharged. These POs allow to prove that the model is sound with respect to some behavioural semantics. Formal definitions of all the standard Event-B POs are given in [1]. In the following, we only recall the most important of them: (event/INV) for *invariant preservation*, which states that the invariant still holds after the execution of each event in the Event-B model M . Given an event e_i with guard $G_i(\bar{t}, \bar{v})$ and action $S_i(\bar{t}, \bar{v})$, this PO is expressed as follows:

$I(\bar{v}) \wedge G_i(\bar{t}, \bar{v}) \wedge S_i(\bar{t}, \bar{v}, \bar{v}') \vdash I(\bar{v}')$	(event/INV)
---	-------------




Figure 1: Event-B model of the emergency brake system

2.2 The Emergency Brake System

We now introduce our running example, based on a simplified scenario of the emergency brake system in charge of manoeuvring the brake of a vehicle.

Specification. To command the brake, a pedal is provided to the driver: when the pedal is switched to “down”, the brake must be applied; when the pedal is switched to “up”, the brake must be released. Some requirements constrain the model:

- R1. Pedal failure: when the driver tries to switch “down” the pedal, it may stay in the same position;
- R2. Risk of pedal failure: the risk of pedal failure is set to 10%;
- R3. Brake failure: the brake may not be applied although the pedal has been switched down;
- R4. Maximum brake wear: the brake cannot be applied more than a fixed number of times;
- R5. Brake wear: due to brake wear, the risk of brake failure increases each time the brake is applied.

Event-B model. The model StdEmergencyBrake given in Fig. 1 presents an Event-B specification of the emergency brake system. The state of the system is described by means of three variables: pedal models the driver command, brake represents the state of the emergency brake (applied or released) and wear counts the number of times the brake is applied. The constant MAX_WEAR represents the maximum number of times the brake can be applied.

The event PushPedal models the driver command, i.e, switching the pedal to down. For taking into account the possible pedal failure mentioned in R1, we use an enumerated non-deterministic assignment $pedal := \{down, up\}$ to express that the pedal is switched to down (the attempted behaviour) or remains in the up position (failure). Using standard Event-B, we cannot take into account the quantitative risk of failure expressed in R2. The event ApplyBrake models the brake application, i.e. the variable brake is assigned the value applied (and the variable wear is increased). The event ApplyBrakeFailure models failure during the brake application: the value of variable brake remains released. When $wear < MAX_WEAR$, the events ApplyBrake and ApplyBrakeFailure are enabled at the same time (when $pedal = down \wedge brake = released$), the subsequent non-determinism between these two events reflects requirement R3. On the other hand, when $wear = MAX_WEAR$, ApplyBrakeFailure is enabled while the guard of ApplyBrake is not satisfied. Therefore, the brake necessarily fails as soon as $wear = MAX_WEAR$, which means that the brake event cannot be triggered more than MAX_WEAR times (the maximum brake wear) as expressed by R4. Requirement R5 cannot be modelled in standard Event-B.

3 Introducing Probabilities in Event-B

The typical way of defining a probabilistic Event-B model from a classical Event-B model M is to go through M and replace all occurrences of non-deterministic choices with probabilistic choices. In Event-B, non-determinism can appear in three places: the choice of the enabled event to be executed, the choice of the parameter value to

be taken and the choice of the value to be assigned to a given variable in a non-deterministic assignment. In the following, we go through these three sources of non-determinism and explain how to turn them into probabilistic choices.

3.1 Introducing probabilistic choices

In standard Event-B, when several events are enabled in a given configuration, the event to be executed is chosen non-deterministically. In order to resolve this non-deterministic choice, we propose to equip each probabilistic event with a *weight*. In configurations where several probabilistic events are enabled, the probability of choosing one of them will therefore be computed as the ratio of its weight against the total value of the weights of all enabled events in this state. Using weights instead of actual probability values is convenient as the set of enabled events evolves with the configuration of the system. Moreover, for the sake of expressivity, we propose to express the weight $W_i(\bar{v})$ of a probabilistic event e_i as an expression over the variables \bar{v} of the probabilistic Event-B model. The probability of executing a given event can therefore evolve as the system progresses. A probabilistic event is therefore allowed to be executed only if *i*) its guards is fulfilled and *ii*) its weight is strictly positive.

In standard Event-B, events can be equipped with parameters. In each configuration where this is possible, a valuation of the parameters is chosen such that the guard $G_i(\bar{t}, \bar{v})$ of the event is satisfied. When there are several such parameter valuations, one of them is selected non-deterministically. We therefore propose to replace this non-deterministic choice by a uniform choice over all parameter valuations ensuring that the guard of the event is satisfied. The uniform distribution is a default choice but our results can be extended to any other discrete distribution.

Recall that non-deterministic assignments in Event-B are expressed in two forms: predicate non-deterministic assignments and enumerated non-deterministic assignments. We propose to replace predicate non-deterministic assignments by *predicate probabilistic assignments* written

$$x := \oplus Q_x(\bar{t}, \bar{v}, x')$$

Instead of choosing non-deterministically among the values of x' such that the predicate $Q_x(\bar{t}, \bar{v}, x')$ is true as in standard predicate non-deterministic assignments, we propose to choose this new value using an uniform distribution. For simplicity reasons, we enforce that this uniform distribution must be discrete, and therefore that the set of values x' such that $Q_x(\bar{t}, \bar{v}, x')$ is true must always be finite. As above, the uniform distribution we propose by default could be replaced by any other discrete distribution.

We propose to replace enumerated non-deterministic assignments by *enumerated probabilistic assignments* written

$$x := E_1(\bar{t}, \bar{v}) @_{p_1} \oplus \dots \oplus E_m(\bar{t}, \bar{v}) @_{p_m}$$

In this structure, the variable x is assigned the expression E_i with probability p_i . In order to define a correct probability distribution, each p_i must be strictly positive and smaller or equal to 1, and they must sum up to 1. Although rational numbers are not natively handled in Event-B, we assume that an adequate context is present. That can be done by defining a "Rational" theory in Rodin using the theory plug-in providing capabilities to define and use mathematical extensions to the Event-B language and the proving infrastructure [6].

Remark that standard deterministic assignments are conserved, but can also be considered as enumerated probabilistic assignments where $m = 1$.

3.2 Syntax

Turning all non-deterministic choices into probabilistic choices has side effects on the syntax of events and models. In probabilistic Event-B, we therefore propose to use the following syntax for a probabilistic event e_i :

$$e_i \hat{=} \text{weight } W_i(\bar{v}) \text{ any } \bar{t} \text{ where } G_i(\bar{t}, \bar{v}) \text{ then } S_i(\bar{t}, \bar{v}) \text{ end}$$

where $W_i(\bar{v})$ is the weight of the event, $G_i(\bar{t}, \bar{v})$ is the guard of the event and $S_i(\bar{t}, \bar{v})$ is a *probabilistic action*, i.e. an action consisting only of deterministic and *probabilistic* assignments which are executed in parallel.

For simplicity reasons we impose, as in standard Event-B, that the initialisation event must be deterministic. The results we present in the rest of the paper can nevertheless easily be extended to probabilistic initialisation events.

Definition 1 *Probabilistic Event-B Model* A probabilistic Event-B model is a tuple $M = (\bar{v}, I(\bar{v}), \text{PEvts}, \text{Init})$ where $\bar{v} = \{v_1 \dots v_n\}$ is a set of variables, $I(\bar{v})$ is the invariant, PEvts is a set of probabilistic events and Init is the initialisation event.




Figure 2: Probabilistic model of the emergency brake system

3.3 Running Example

A probabilistic version of the emergency brake system from Section 2.2 is given in Fig. 2. This model has the same variables `pedal`, `brake` and `wear`, the same invariants and the same events as the Event-B model `StdEmergencyBrake` from Fig. 1. Remark that, unlike in standard Event-B, requirements R2 and R5 can be taken into account in this probabilistic version. R2 is specified in the probabilistic event `PushPedal` by using an enumerated probabilistic assignment instead of a non-deterministic assignment: the variable `pedal` is assigned the value `down` with a probability 9/10 (attempted behaviour) and the value `up` with a probability 1/10 (failure), hence resulting in a risk of pedal failure of 10%. Requirement R5 is taken into account by annotating probabilistic event `ApplyBrake` with a weight $\text{MAX_WEAR} - \text{wear}$ and probabilistic event `ApplyBrakeFailure` with a weight `wear`. As the probabilistic event `ApplyBrake` increases the variable `wear` when it is executed, the weight of the probabilistic event `ApplyBrake` decreases each time it is executed whereas the weight of the probabilistic event `ApplyBrakeFailure` increases. The failure of the brake is modelled by means of a probabilistic choice between `ApplyBrake` and `ApplyBrakeFailure` instead of a non-deterministic choice as in the standard version, which implies that the more `ApplyBrake` is executed, the higher the probability that `ApplyBrakeFailure` occurs instead. In this version, all requirements are therefore taken into account.

4 Consistency of probabilistic Event-B models

As in standard Event-B, the consistency of a probabilistic Event-B model is defined by means of proof obligations (POs). In this section, we therefore introduce new POs specific to probabilistic Event-B and explain how we adapt standard Event-B POs in order to prove the consistency of probabilistic Event-B models.

4.1 Specific POs for Probabilistic Event-B

Numeric weight. For simplicity reasons, we impose that the expression $W_i(\vec{v})$ representing the weight of a given probabilistic event must evaluate to natural numbers.

$$\boxed{I(\vec{v}) \wedge G_i(\vec{f}, \vec{v}) \vdash W_i(\vec{v}) \in \text{NAT}} \quad (\text{event/WGHT/NAT})$$

Parameter values finiteness. In order to be able to use a discrete uniform distribution over the set of parameter valuations ensuring that the guard of a probabilistic event is satisfied, we impose that this set must be finite.

$$\boxed{I(\vec{v}) \vdash \text{finite}(\{\vec{f} \mid G_i(\vec{f}, \vec{v})\})} \quad (\text{event/param/pWD})$$

Enumerated probabilistic assignments well-definedness and feasibility. In all enumerated probabilistic assignments, it is necessary to ensure that the discrete probability values $p_1 \dots p_n$ define a correct probability distribution.

Formally, this leads to two POs:

1. Probability values p_i in enumerated probabilistic assignments are strictly positive and smaller or equal to 1.

$$\boxed{\vdash 0 < p_i \leq 1} \quad (\text{event/assign/pWD1})$$

2. The sum of the probability values $p_1 \dots p_n$ in enumerated probabilistic assignments must be equal to 1.

$$\boxed{\vdash p_1 + \dots + p_n = 1} \quad (\text{event/assign/pWD2})$$

Feasibility of enumerated probabilistic assignments is trivial: as soon as at least one expression $E_i(\bar{t}, \bar{v})$ is present and well-defined, it always returns a value.

Predicate probabilistic assignment well-definedness and feasibility. To define a discrete uniform distribution over the set of values of a variable x making the predicate $Q_x(\bar{t}, \bar{v}, x')$ of the corresponding assignment satisfied, we impose that this set must be finite.

$$\boxed{I(\bar{v}) \wedge G_i(\bar{t}, \bar{v}) \wedge W_i(\bar{v}) > 0 \vdash \text{finite}(\{x' \mid Q_x(\bar{t}, \bar{v}, x')\})} \quad (\text{event/assign/pWD3})$$

Feasibility of predicate probabilistic assignments is ensured by the standard feasibility PO [1] inherited from Event-B. It ensures that the set $\{x' \mid Q_x(\bar{t}, \bar{v}, x')\}$ is not empty.

4.2 Modifications to Standard POs

Where standard Event-B POs are concerned, the main difference in probabilistic Event-B is the condition for a probabilistic event to be enabled. Indeed, while it suffices to show that the guard of an event is satisfied for this event to be enabled in standard Event-B, we also have to show in probabilistic Event-B that its weight is strictly positive. We therefore modify standard Event-B POs as follows.

Invariant preservation. The invariant must be preserved by all enabled probabilistic events.

$$\boxed{I(\bar{v}) \wedge G_i(\bar{t}, \bar{v}) \wedge W_i(\bar{v}) > 0 \wedge S_i(\bar{t}, \bar{v}, \bar{v}') \vdash I(\bar{v}')} \quad (\text{event/pINV})$$

Deadlock freedom. In all acceptable configurations, there must exist at least one enabled probabilistic event.

$$\boxed{I(\bar{v}) \vdash (G_1(\bar{t}, \bar{v}) \wedge W_1(\bar{v}) > 0) \vee \dots \vee (G_n(\bar{t}, \bar{v}) \wedge W_n(\bar{v}) > 0)} \quad (\text{model/pDLF})$$

For the sake of understanding, we hereby insist on the separation between the guard of an event, which reflects the classical notion of enabledness, and the fact that its weight must be strictly positive. Obviously, one could also automatically re-write the guard of all probabilistic events in order to include the condition on its weight. This solution would allow conserving most of the standard Event-B consistency POs without modifications in the probabilistic setting.

5 Semantics

Semantics of standard Event-B models can be expressed in terms of Labelled Transition Systems [5]. Informally, given an Event-B model $M = (\bar{v}, I(\bar{v}), \text{Evts}, \text{Init})$, its semantics is the LTS $\mathcal{M} = (S, s_0, AP, L, \text{Acts}, T)$ where S is a set of states, Acts is the set of actions (event names), $s_0 \in S$ is the initial state obtained by executing the Init event, AP is the set of valuations of the variables in \bar{v} that satisfy the invariant $I(\bar{v})$, $L : S \rightarrow AP$ is a labelling function that provides the valuations of the variables in a given state, and $T \subseteq S \times \text{Acts} \times S$ is the transition relation corresponding to the actions of the events.

In the following, we extend this work by presenting the semantics of probabilistic Event-B models in terms of Discrete Time Markov Chains (DTMC).

Remark that our goal, unlike in [19, 21] is not to translate our models into DTMCs and use standard model-checking techniques to verify them. Instead, we aim at reasoning directly on probabilistic Event-B models and benefiting from the symbolic proof mechanism that is the signature of the Event-B approach. The following DTMC semantics are nevertheless introduced as a demonstration of the correctness of our approach and results.

5.1 Notations

Let $M = (\bar{v}, I(\bar{v}), \text{PEvts}, \text{Init})$ be a probabilistic Event-B model and σ be a valuation its variables. Given a variable $x \in \bar{v}$, we write $[\sigma]x$ for the value of x in σ . Given an expression $E(\bar{v})$ over variables in \bar{v} , we write $[\sigma]E(\bar{v})$ (or $[\sigma]E$ when clear from the context) for the evaluation of $E(\bar{v})$ in the context of σ .

Given a probabilistic event e_i with a set of parameters \bar{t} and a valuation σ of the variables, we write $T_\sigma^{e_i}$ for the set of parameter valuations θ such that the guard of e_i evaluates to true in the context of σ and θ . Formally, $T_\sigma^{e_i} = \{\theta \mid [\sigma, \theta]G_i(\bar{t}, \bar{v}) = \text{true}\}$. Recall that parameter valuations are chosen uniformly on this set. We write $P_{T_\sigma^{e_i}}$ for the uniform distribution on the set $T_\sigma^{e_i}$.

Given a valuation σ of the variables and a probabilistic event e_i , we say that e_i is *enabled* in σ iff (a) the weight of e_i evaluates to a strictly positive value in σ and (b) either e_i has no parameter and its guard evaluates to true in σ or there exists at least one parameter valuation θ such that the guard of e_i evaluates to true in the context of σ and θ , i.e. $T_\sigma^{e_i} \neq \emptyset$.

Given a probabilistic event e_i , we write $\text{Var}(e_i)$ for the set of variables in \bar{v} that are modified by the action of e_i , i.e. the variables that appear on the left side of an assignment in $S_i(\bar{t}, \bar{v})$. Recall that a variable $x \in \text{Var}(e_i)$ must be on the left side of either a predicate probabilistic assignment or an enumerated probabilistic assignment. In both cases, given an original valuation σ of the variables, a valuation θ of the parameters of e_i and a target valuation σ' of the variables, we write $P_{\sigma, \theta}^{e_i}(x, \sigma')$ for the probability that x is assigned the new value $[\sigma']x$ when executing e_i from the valuation σ and with parameter valuation θ . If e_i is not equipped with parameters, this is written $P_\sigma^{e_i}(x, \sigma')$. In the following, we always use the more general notation and assume that it is replaced with the specific one when there are no parameters. The formal definition of $P_{\sigma, \theta}^{e_i}(x, \sigma')$ is given in Appendix A.1.

5.2 DTMC semantics

Informally, the semantics of a probabilistic Event-B model $M = (\bar{v}, I(\bar{v}), \text{PEvts}, \text{Init})$ is a Probabilistic LTS $\llbracket M \rrbracket = (S, s_0, AP, L, \text{Acts}, P)$ where the states, labels, actions, atomic propositions and initial state are similarly obtained as for the standard LTS semantics of Event-B. The only difference with the standard LTS semantics is that the transitions are equipped with probabilities, which we explain below. In the following, we identify the states with the valuations of the variables defined in their labels.

Intuitively, the transition probabilities are obtained as follows: Let $e_i \in \text{PEvts}$ be a probabilistic event, $x \in \bar{v}$ be a variable and s, s' be two states of $\llbracket M \rrbracket$ such that (s, e_i, s') is a transition in the standard LTS semantics, i.e. where e_i is enabled in s and there exists a parameter valuation $\theta \in T_s^{e_i}$, if any, such that the action of e_i may take the system from s to s' . The probability assigned to transition (s, e_i, s') is then equal to the product of (1) the probability that the event e_i is chosen from the set of enabled events in state s , (2) the probability of choosing each parameter valuation θ , and (3) the overall probability that each modified variable is assigned the value given in s' under parameter valuation θ .

Definition 2 *Probabilistic Event-B Semantics* The semantics of a probabilistic Event-B model $M = (\bar{v}, I(\bar{v}), \text{PEvts}, \text{Init})$ is a PLTS $\llbracket M \rrbracket = (S, s_0, AP, L, \text{Acts}, P)$ where S is a set of states where each state is uniquely identified by its label, $s_0 \in S$ is the initial state obtained after the execution of the Init event, AP represents the valuations of all variables that satisfy the invariant of the model: $AP = \{\sigma \mid [\sigma]I(\bar{v}) = \text{true}\}$, $L : S \rightarrow AP$ is the labelling function that assigns to each state the corresponding valuation of the variables, Acts is the alphabet of actions (event names), and $P : S \times \text{Acts} \times S \rightarrow [0, 1]$ is the transition probability function such that for a given state s , for all $e_i, s' \in \text{Acts} \times S$, we have $P(s, e_i, s') = 0$ if $e_i \notin \text{Acts}(s)$ or $\exists x \in X \setminus \{\text{Var}(e_i)\}$ st $[s]x \neq [s']x$ and otherwise

$$P(s, e_i, s') = \underbrace{\frac{[s]W_i(\bar{v})}{\sum_{e_j \in \text{Acts}(s)} [s]W_j(\bar{v})}}_{(1)} \times \sum_{\theta \in T_s^{e_i}} \underbrace{(P_{T_s^{e_i}}(\theta))}_{(2)} \times \underbrace{\prod_{x \in \text{Var}(e_i)} P_{s, \theta}^{e_i}(x, s')}_{(3)}$$

In the following proposition, we show that the semantics of a probabilistic Event-B model as defined above is indeed a DTMC. For space reasons, the proof of this proposition is given in Appendix A.3.

Proposition 1 *The semantics of a probabilistic Event-B model M satisfying the POs given in Section 4.1 is a DTMC.*

For space reasons, the DTMC of the probabilistic emergency brake system is given in Appendix A.2.

6 Convergence

The development process in Event-B is inherently based on refinement. As said earlier, systems are typically developed progressively using an ordered sequence of models, where each model contains more details than its predecessor. One key aspect of refinement is the addition, in one refinement step, of new variables and new events that characterize the evolution of those variables. In order to preserve certain properties, it is then necessary to show that the introduction of these new events in a refined model cannot prevent the system from behaving as specified in the abstract model. In particular, it is necessary to show that such new events are “convergent”, in the sense that they cannot keep control indefinitely: at some point the system has to stop executing new events in order to follow the behaviour specified in its abstract model.

Although this paper does not address refinement in probabilistic Event-B, we propose a solution in order to prove that a given set of events almost-certainly converges in a probabilistic Event-B model, which is a necessary step for addressing refinement in the future. We therefore start this section with a brief recall of how events can be proven convergent in standard Event-B and then propose a set of sufficient conditions, expressed as POs, that allow proving that a set of events is almost-certainly convergent in probabilistic Event-B.

Convergence in Standard Event-B. In order to prove that a set of events is convergent in Event-B, one has to show that it is not possible to keep executing convergent events infinitely, and therefore that a non-convergent event is eventually performed from any state. The classical solution is therefore to introduce a natural number expression $V(\bar{v})$, called a *variant*, and show that all convergent events strictly decrease the value of this variant. As a consequence, when the variant hits zero, it is guaranteed that no convergent event can be performed. In practice, this is expressed using two POs:

1. **Numeric variant.** Under the guard $G_i(\bar{t}, \bar{v})$ of each convergent event e_i , the variant $V(\bar{v})$ is greater than 0.

$I(\bar{v}) \wedge G_i(\bar{t}, \bar{v}) \vdash V(\bar{v}) \in \text{NAT}$	(event/var/NAT)
--	-----------------

2. **Convergence.** The action $S_i(\bar{t}, \bar{v})$ of each convergent event e_i *must always* decrease the variant $V(\bar{v})$.

$I(\bar{v}) \wedge G_i(\bar{t}, \bar{v}) \vdash \forall \bar{v}'. S_i(\bar{t}, \bar{v}, \bar{v}') \Rightarrow V(\bar{v}') < V(\bar{v})$	(event/VAR)
---	-------------

Almost-certain Convergence in the literature. In the context of probabilistic Event-B, instead of proving that a given set of events necessarily converges as in standard Event-B, we are interested in showing that a given set of events *almost-certainly* converges. In other words, we are interested in showing that, in all states of the system where convergent events can be executed, the probability of eventually taking a non-convergent event or reaching a deadlock is 1 (the probability of infinitely executing convergent events is 0).

This property has already been investigated in [8] and [11], in the context of events having probabilistic actions but where non-determinism is still present between events. In this context, Hallerstede and Hoang propose in [8] sufficient conditions for a set of events to almost-certainly converge. These conditions can be summarized as follows: As in standard Event-B, one needs to exhibit a natural number expression $V(\bar{v})$ called a variant, but unlike in the standard setting, only one resulting valuation of the execution of *each* convergent event needs to decrease this variant. Indeed, in this case, the probability of decreasing the variant is strictly positive. Unfortunately, using such a permissive condition is not sufficient: there might also be a strictly positive probability of increasing the variant. Therefore, Hallerstede and Hoang require the introduction of another natural number expression $U(\bar{v})$ which must maximise the variant $V(\bar{v})$ and never increase. The proposition from [8] is refined in [11], where Hoang requires in addition that the probabilities considered in probabilistic assignments are bounded away from 0. This is ensured by requiring that the set of values that can be returned by a probabilistic assignment is finite.

Adaptation to fully probabilistic Event-B. We now show how to adapt the results proposed in [8] and [11] to our fully probabilistic Event-B setting. Since there are no non-deterministic choices between enabled events, it is not anymore necessary to require that *all* enabled events in a given configuration may decrease the variant. We therefore start by relaxing the condition proposed in [8]: we only require that, in all configurations where a convergent event is enabled, there is *at least one* convergent event for which at least one resulting valuation decreases the variant.

1. **Almost-certain convergence.** In all configurations where at least one convergent event is enabled, there must exist at least one valuation \bar{v}' obtained after the execution of one of these enabled events which decreases the variant.

$$\boxed{I(\bar{v}) \wedge ((G_1(\bar{t}, \bar{v}) \wedge W_1(\bar{v}) > 0) \vee \dots \vee (G_i(\bar{t}, \bar{v}) \wedge W_i(\bar{v}) > 0)) \vdash (\exists \bar{v}'. G_1(\bar{t}, \bar{v}) \wedge W_1(\bar{v}) > 0 \wedge S_1(\bar{t}, \bar{v}, \bar{v}') \wedge V(\bar{v}') < V(\bar{v})) \vee \dots \vee (\exists \bar{v}'. G_i(\bar{t}, \bar{v}) \wedge W_i(\bar{v}) > 0 \wedge S_i(\bar{t}, \bar{v}, \bar{v}') \wedge V(\bar{v}') < V(\bar{v}))} \quad (\text{model/pVar})$$

As in [8], we also require that convergent events can only be enabled when the variant is positive and that the variant is bounded above. In order to simplify the reasoning, we propose to use a constant bound U , as in [11].

2. **Numeric variant.** Convergent events can only be enabled when the variant is greater or equal to 0.

$$\boxed{I(\bar{v}) \wedge G_i(\bar{t}, \bar{v}) \wedge W_i(\bar{v}) > 0 \vdash V(\bar{v}) \in \text{NAT}} \quad (\text{event/var/pNAT})$$

3. **Bounded variant.** Convergent events can only be enabled when the variant is less or equal to U .

$$\boxed{I(\bar{v}) \wedge G_i(\bar{t}, \bar{v}) \wedge W_i(\bar{v}) > 0 \vdash V(\bar{v}) \leq U} \quad (\text{event/pBOUND})$$

Finally, the finiteness of the set of values that can be returned by a probabilistic assignment is already ensured by the syntax for enumerated probabilistic assignments and by PO (event/assign/pWD3) for predicate probabilistic assignments and their non-emptiness is ensured by the standard feasibility POs.

Inadequacy of adapted POs. Unfortunately, as we deal with potentially infinite-state systems, POs 1–3 presented above are not anymore sufficient for proving that the probability of eventually executing a non-convergent event or reaching a deadlock state is 1. Indeed, although the probability of decreasing the variant is always strictly positive because of PO (model/pVar) and although the number of values that can be returned by a given probabilistic assignment is always finite, the combination of event weights and parameter choice can make this value infinitely small in some cases. In this case, it is well known that almost-certain reachability/convergence is not ensured. This problem is a direct consequence of the unboundedness of the weights of convergent events, which, by getting arbitrarily big, cause the probability of decreasing the variant to get arbitrarily small. Examples illustrating this fact are given in Appendix B.

Additional Proof Obligations. We therefore adapt classical results from infinite-state DTMC to our setting and propose sufficient conditions in terms of proof obligations to prove the almost-certain convergence of a given set of events. Informally, the following POs ensure that the probability of decreasing the variant cannot get infinitely small by requiring that both the weights of convergent events and the number of potential values given to parameters in convergent events are bounded.

4. **Bounded weight.** The weight of all convergent events must be bounded above by a constant upper bound BW.

$$\boxed{I(\bar{v}) \wedge G_i(\bar{t}, \bar{v}) \vdash W_i(\bar{v}) \leq \text{BW}} \quad (\text{event/wght/BOUND})$$

5. **Bounded parameters.** The number of potential values for parameters in convergent events must be bounded above by a constant upper bound BP.

$$\boxed{I(\bar{v}) \vdash \text{card}(\{\bar{t} \mid G_i(\bar{t}, \bar{v})\}) \leq \text{BP}} \quad (\text{event/param/BOUND})$$

We now formally prove that the conditions presented above are sufficient for guaranteeing the almost-certain convergence of a given set of events in a probabilistic Event-B model.

Theorem 1 *Let $M = (\bar{v}, I(\bar{v}), V(\bar{v}), \text{PEvts}, \text{Init})$ be a probabilistic Event-B model and $\text{PEvts}_c \subseteq \text{PEvts}$ a set of convergent events. If M satisfies the above POs (1-5), then the set PEvts_c almost-certainly converges.*

Proof. We consider the DTMC semantics $\llbracket M \rrbracket$ of the probabilistic Event-B model M and use the global coarseness property of infinite-state DTMC [14] to show that, from all states of $\llbracket M \rrbracket$, the probability of eventually taking a non-convergent event or reaching a deadlock is 1. The full proof is presented in Appendix B.3. \square

7 Conclusion

As suggested by Morgan *et al.* in [16], the ideal probabilistic extension of Event-B should allow using probabilities as a refinement of non-deterministic choices in all places where such choices exist. In Event-B, non-determinism occurs in several places and, to the best of our knowledge, existing works on extending Event-B with probabilities




Figure 3: Probabilistic plugin to the Rodin platform

have only focused on refining non-deterministic assignments into probabilistic assignments [8, 19, 21] while leaving other sources of non-determinism such as the choice between enabled events or the choice between admissible parameter values untouched.

In this paper, we have proposed a fully probabilistic extension of Event-B where probabilistic choices are introduced as replacement of *all* non-deterministic choices, be it between enabled events, parameter values or assignments as suggested by Morgan *et al.* in their seminal work. Our long term goal is to produce a probabilistic extension of Event-B where the developer can choose at his convenience where to refine non-deterministic choices with probabilities and where to keep non-deterministic choices intact. However, this long-term goal is clearly ambitious and will require several years of study to be achieved. In this paper, we have therefore focused on a more reasonable objective, restricting ourselves to purely probabilistic systems where probabilities appear in the last step of refinement. Although the long-term goal presented above is not yet achieved, this is clearly a first step in the right direction.

In particular, we have introduced new notations and semantics, along with novel and adapted POs dedicated to the *consistency* of probabilistic Event-B models. We have shown that, when these POs are satisfied, the semantics of a probabilistic Event-B model is a discrete time Markov chain. Finally, we have provided sufficient conditions, expressed in terms of POs, to show that a probabilistic Event-B model satisfies the *almost-certain convergence* of a given set of events, which is a necessary step for addressing refinement in the future.

In parallel, we have started the development of a prototype plugin for the Rodin Platform. This plugin currently allows the specification of fully probabilistic Event-B models and the semi-automatic generation of a probabilistic Event-B model from a standard Event-B model as shown in Fig 3. It also supports the generation of several consistency proof obligations on probabilistic Event-B models.

Future work. As the development in Event-B is intrinsically based on a refinement process, we plan on studying the refinement of probabilistic Event-B models, including (but not restricting to) the "probabilisation" of non-deterministic models, the introduction of new probabilistic events, and, the merge and the split of probabilistic events. We also plan to study how to handle Event-B models combining non-deterministic and probabilistic events as well as the (probabilistic) refinement of such models.

Most of the properties of interest that are verified in standard Event-B are safety-related. They are most of the time expressed by means of invariants and discharged as POs. We therefore plan to consider *probabilistic invariants*, i.e. invariants related to probabilistic distributions [12]. In addition, critical systems must also satisfy some liveness properties. In this paper, we have studied the *almost-certain convergence* of a given set of events, but other probabilistic liveness properties could be considered. Indeed, the verification of other liveness properties on standard Event-B models using refinement and proof obligations have been considered in [9, 4]. We will pursue these works and extend them to the verification of probabilistic liveness properties on probabilistic Event-B models.

References

- [1] J.-R. Abrial. *Modeling in Event-B: system and software engineering*. Cambridge University Press, 2010.
- [2] J.-R. Abrial, M. Butler, S. Hallerstede, T. S. Hoang, F. Mehta, and L. Voisin. Rodin: an open toolset for modelling and reasoning in event-b. *International journal on software tools for technology transfer*, 12(6):447–466, 2010.
- [3] J.-R. Abrial, D. Cansell, and D. Méry. A mechanically proved and incremental development of ieee 1394 tree identify protocol. *Formal aspects of computing*, 14(3):215–227, 2003.
- [4] J.-R. Abrial, D. Cansell, and D. Méry. Refinement and reachability in event-b. In *ZB 2005: Formal Specification and Development in Z and B*, volume 3455 of *LNCS*, pages 222–241. Springer, 2005.
- [5] D. Bert and F. Cave. Construction of finite labelled transition systems from b abstract systems. In *Integrated Formal Methods*, volume 1945 of *LNCS*, pages 235–254. Springer, 2000.
- [6] M. Butler and I. Maamria. Practical theory extension in event-b. In *Theories of Programming and Formal Methods*, volume 8051 of *LNCS*, pages 67–81. 2013.
- [7] W. W. Chu and C.-M. Sit. Estimating task response time with contentions for real-time distributed systems. In *Real-Time Systems Symposium, 1988., Proceedings.*, pages 272–281. IEEE, 1988.
- [8] S. Hallerstede and T. S. Hoang. Qualitative probabilistic modelling in event-b. In *Integrated Formal Methods*, pages 293–312. Springer, 2007.
- [9] T. Hoang and J.-R. Abrial. Reasoning about liveness properties in event-b. In *Formal Methods and Software Engineering*, volume 6991 of *LNCS*, pages 456–471. Springer, 2011.
- [10] T. S. Hoang. *The development of a probabilistic B-method and a supporting toolkit*. PhD thesis, The University of New South Wales, 2005.
- [11] T. S. Hoang. Reasoning about almost-certain convergence properties using event-b. *Sci. Comput. Program.*, 81:108–121, 2014.
- [12] T. S. Hoang, Z. Jin, K. Robinson, A. McIver, and C. Morgan. Probabilistic invariants for probabilistic machines. In *ZB 2003: Formal Specification and Development in Z and B*, pages 240–259. Springer, 2003.
- [13] T. S. Hoang, Z. Jin, K. Robinson, A. McIver, and C. Morgan. Development via refinement in probabilistic b - foundation and case study. In *International Conference of B and Z Users*, pages 355–373. Springer, 2005.
- [14] R. Mayr, N. B. Henda, and P. A. Abdulla. Decisive markov chains. *Logical Methods in Computer Science*, 3, 2007.
- [15] A. McIver, C. Morgan, and T. S. Hoang. Probabilistic termination in b. In *International Conference of B and Z Users*, pages 216–239. Springer, 2003.
- [16] C. Morgan, T. S. Hoang, and J.-R. Abrial. The challenge of probabilistic event b—extended abstract—. In *ZB 2005: Formal Specification and Development in Z and B*, pages 162–171. Springer, 2005.
- [17] C. Morgan and A. McIver. Abstraction, refinement and proof for probabilistic systems. *Monographs in Computer Science*. Springer, 2005.
- [18] R. Motwani and P. Raghavan. *Randomized algorithms*. Chapman & Hall/CRC, 2010.
- [19] A. Tarasyuk, E. Troubitsyna, and L. Laibinis. Reliability assessment in event-b development. *NODES 09*, page 11, 2009.
- [20] A. Tarasyuk, E. Troubitsyna, and L. Laibinis. Towards probabilistic modelling in event-b. In *Integrated Formal Methods*, pages 275–289. Springer, 2010.
- [21] A. Tarasyuk, E. Troubitsyna, and L. Laibinis. Integrating stochastic reasoning into event-b development. *Formal Aspects of Computing*, 27(1):53–77, 2015.

- [22] K. S. Trivedi, S. Ramani, and R. Fricks. Recent advances in modeling response-time distributions in real-time systems. *Proceedings of the IEEE*, 91(7):1023–1037, 2003.
- [23] A. Villemeur. *Reliability, Availability, Maintainability and Safety Assessment, Assessment, Hardware, Software and Human Factors*, volume 2. Wiley, 1992.
- [24] E. Yilmaz. *Tool support for qualitative reasoning in Event-B*. PhD thesis, Master Thesis ETH Zürich, 2010, 2010.

The appendix is given as complementary material for the understanding and correctness of our propositions and results. They can be consulted at the discretion of the reviewers for review purpose and will be removed in the final version of the paper.

A Complements to Probabilistic Semantics

A.1 Notations

In this section, we provide some basic notations specific to the DTMC semantics of probabilistic Event-B models. Let $M = (\bar{v}, I(\bar{v}), \text{PEvts}, \text{Init})$ be a probabilistic Event-B model. Let e_i be a probabilistic event in PEvts and let $x \in \text{Var}(e_i)$. Recall that x can be modified only by one assignment within the action of e_i . If x is modified by an enumerated probabilistic assignment ($x := E_1(\bar{t}, \bar{v}) @_{p_1} \oplus \dots \oplus E_m(\bar{t}, \bar{v}) @_{p_m}$ ($m \geq 1$)), then we write $\mathcal{E}_{e_i}(x)$ for the set of all expressions that can be assigned to the variable x by this assignment.

$$\mathcal{E}_{e_i}(x) = \{E_1(\bar{t}, \bar{v}), \dots, E_m(\bar{t}, \bar{v})\}$$

The probability of choosing an expression E_i among all others expressions is written $P_x^{e_i}(E_i) = p_i$.

Let $e_i \in \text{PEvts}$ be a probabilistic event, $x \in \text{Var}(e_i)$ be a variable, σ, σ' two valuations of the variables \bar{v} and θ a valuation of the parameter values associated to the event e_i such that e_i is enabled in σ w.r.t parameter valuation θ and leads the system to σ' .

If x is modified by an enumerated probabilistic assignment of e_i , then we write $\mathcal{E}_{e_i}(x)|_{\sigma, \theta}^{\sigma'}$ for the set of expressions in $\mathcal{E}_{e_i}(x)$ such that their evaluation in the valuation σ with parameter valuation θ returns the value of x in the valuation σ' .

Formally,

$$\mathcal{E}_{e_i}(x)|_{\sigma, \theta}^{\sigma'} = \{E \in \mathcal{E}_{e_i}(x) \mid [\sigma, \theta](E(\bar{t}, \bar{v})) = [\sigma']x\}$$

If e_i is not equipped with parameters, then this subset is written $\mathcal{E}_{e_i}(x)|_{\sigma}^{\sigma'}$.

If x is modified by a predicate probabilistic assignment ($x : \oplus Q_x(\bar{t}, \bar{v}, x')$), then we write $\mathcal{V}_{\theta, \sigma}^{e_i}(x)$ for the set of values x' that make the predicate $Q_x(\bar{t}, \bar{v}, x')$ true when evaluated in σ and θ .

$$\mathcal{V}_{\theta, \sigma}^{e_i}(x) = \{x' \mid [\sigma, \theta]Q_x(\bar{t}, \bar{v}, x') = \text{true}\}$$

If e_i is not equipped with parameters, then this subset is written $\mathcal{V}_{\sigma}^{e_i}(x)$.

In Section 5, we have defined the probability $P_{\sigma, \theta}^{e_i}(x, \sigma')$ that the variable x is assigned the new value $[\sigma']x$ when executing e_i from the valuation σ with parameter valuation θ . Formally, this probability is given by:

1. if x is modified by an enumerated probabilistic assignment, then:

$$P_{\sigma, \theta}^{e_i}(x, \sigma') = \sum_{E \in \mathcal{E}_{e_i}(x)|_{\sigma, \theta}^{\sigma'}} P_x^{e_i}(E)$$

2. if x is modified by a predicate probabilistic assignment, then:

$$P_{\sigma, \theta}^{e_i}(x, \sigma') = \frac{1}{\text{card}(\mathcal{V}_{\theta, \sigma}^{e_i}(x))} \text{ if } [\sigma']x \in \mathcal{V}_{\theta, \sigma}^{e_i}(x) \text{ and } 0 \text{ otherwise.}$$

A.2 DTMC semantics of the probabilistic Emergency Brake system

Fig. 4 presents the DTMC semantics of the probabilistic Event-B model of the emergency brake system given in Fig. 2. The states of this DTMC correspond to the valuations of the variables *pedal*, *brake* and *wear*.

The transitions correspond to the possible occurrence of the events, labelled with their probability value. In this example, we set the constant *MAX_WEAR* constant to 4. The probabilities are computed as follows:

In the state $(d, r, 0)$, three events are enabled: *ApplyBrake*, *ApplyBrakeFailure* and *ReleasePedal*. The event *ReleasePedal* leads to the state $(u, r, 0)$ with probability $\frac{1}{3} = \frac{4}{4+4+4}$, where $\frac{4}{4+4+4}$ corresponds to the probability of choosing the event *ReleasePedal* rather than the events *ApplyBrake* and *ApplyBrakeFailure*. The events *ApplyBrake* and *ApplyBrakeFailure* have the same probability value, this value is similarly calculated as for the event *ReleasePedal*. In the state $(u, r, 2)$, the event *PushPedal* is enabled, it leads to the state $(u, r, 2)$ with probability $\frac{1}{10}$ and the state $(d, r, 2)$ with probability $\frac{9}{10}$. $\frac{9}{10}$ corresponds to the probability of assigning the value *down* to the variable *pedal* and $\frac{1}{10}$ corresponds to the probability of assigning the value *up* to the variable *pedal*. The probabilities of all the transitions in this DTMC are computed in a similar manner.

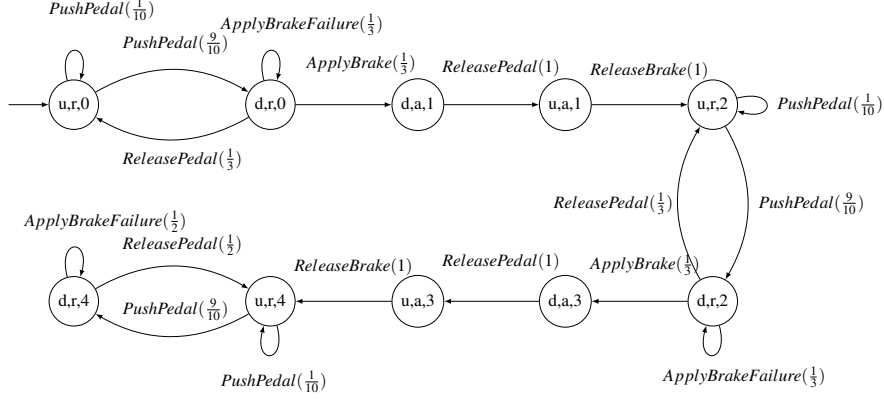


Figure 4: DTMC of the probabilistic EBS with $MAX_WEAR = 4$

A.3 Proof of Proposition 1

Given a probabilistic Event-B model M , the semantics $\llbracket M \rrbracket$ of M is a DTMC.

We must prove that for each state s in $\llbracket M \rrbracket$, the sum of probabilities of the outgoing transitions from s is equal to one. Let M be a probabilistic Event-B model, $\bar{v} = (x_1, x_2, \dots, x_n)$ the set of variables of M and $s \in S$ a state of $\llbracket M \rrbracket$. We assume that each variable x_i in \bar{v} takes its value from a set X_i .

Recall that the probability of a transition (s, e_i, s') is 0 if $e_i \notin \text{Acts}(s)$ or $\exists x \in \bar{v} \setminus \{\text{Var}(e_i)\} \mid [s]x \neq [s']x$ and otherwise:

$$P(s, e_i, s') = \frac{[s]W_i(\bar{v})}{\sum_{e_j \in \text{Acts}(s)} [s]W_j(\bar{v})} \times \sum_{\theta \in T_s^{e_i}} (P_{T_s^{e_i}}(\theta) \times \prod_{x \in \text{Var}(e_i)} P_{s, \theta}^{e_i}(x, s'))$$

We must therefore show that $\sum_{e_i \in \text{Acts}(s)} \sum_{s' \in S} P(s, e_i, s') = 1$.

$$\begin{aligned} \sum_{s' \in S, e_i \in \text{Acts}(s)} P(s, e_i, s') &= \sum_{e_i \in \text{Acts}(s)} \sum_{s' \in S} \frac{[s]W_i(\bar{v})}{\sum_{e_j \in \text{Acts}(s)} [s]W_j(\bar{v})} \times \sum_{\theta \in T_s^{e_i}} (P_{T_s^{e_i}}(\theta) \times \prod_{x \in \text{Var}(e_i)} P_{s, \theta}^{e_i}(x, s')) \\ \sum_{s' \in S, e_i \in \text{Acts}(s)} P(s, e_i, s') &= \sum_{e_i \in \text{Acts}(s)} \frac{[s]W_i(\bar{v})}{\sum_{e_j \in \text{Acts}(s)} [s]W_j(\bar{v})} \times \sum_{\theta \in T_s^{e_i}} (P_{T_s^{e_i}}(\theta) \times \sum_{s' \in S, x \in \text{Var}(e_i)} \prod P_{s, \theta}^{e_i}(x, s')) \end{aligned}$$

Let $S_1 = \{s' \in S \mid \forall x \in \bar{v} \setminus \text{Var}(e_i). [s]x = [s']x\}$.

$$\sum_{s' \in S_1, e_i \in \text{Acts}(s)} P(s, e_i, s') = \sum_{e_i \in \text{Acts}(s)} \frac{[s]W_i(\bar{v})}{\sum_{e_j \in \text{Acts}(s)} [s]W_j(\bar{v})} \times \sum_{\theta \in T_s^{e_i}} (P_{T_s^{e_i}}(\theta) \times \sum_{s' \in S_1, x \in \text{Var}(e_i)} \prod P_{s, \theta}^{e_i}(x, s'))$$

$\forall x \in \text{Var}(e_i)$, we recall that $P_{s, \theta}^{e_i}(x, s') = \sum_{E \in \mathcal{E}_{e_i}(x) \mid s'_\theta} P_x^{e_i}(E)$ if x is modified by an enumerated probabilistic assignment and $P_{s, \theta}^{e_i}(x, s') = \frac{1}{\text{card}(\mathcal{V}_{\theta, s}^{e_i}(x))}$ if x is modified by a predicate probabilistic assignment.

We then remark that $P_{s, \theta}^{e_i}(x, s')$ does not really depend on s' but only depends on $v'_x = [s']x$ (as s' corresponds to the valuations of the variables x_i in the state s').

Given $x \in \bar{v}$ and $v'_x \in X$, we therefore write $F_x^{s, \theta, e_i}(v'_x) = P_{s, \theta}^{e_i}(x, s')$ if $x \in \text{Var}(e_i)$.

For $\bar{v} = \{x_1, \dots, x_n\}$, we have $S_1 = \{(v'_1, \dots, v'_n) \mid v'_i = [s']x_i \text{ if } x_i \in \text{Var}(e_i) \text{ and } v'_i \in X_i \text{ otherwise}\}$.

We assume that $\text{Var}(e_i) = \{x_1, \dots, x_k\}$ with $k \leq n$,

Then for all expression α with $\alpha = [v_{x_{k+1}} = [s]x_{k+1}, \dots, v_{x_n} = [s]x_n]$ we have:

$$\sum_{s' \in S_1} \alpha = \sum_{v_{x_1} \in X_1} \left(\sum_{v_{x_2} \in X_2} \left(\dots \sum_{v_{x_k} \in X_k} \alpha \right) \right)$$

As a consequence, we have:

$$\begin{aligned} \sum_{s' \in S_1} \prod_{x_i \in \text{Var}(e_i)} F_x^{s, \theta, e_i}([s']x_i) &= \sum_{v'_{x_1} \in X_1} \left(\sum_{v'_{x_2} \in X_2} \left(\dots \sum_{v'_{x_k} \in X_k} \left(\prod_{i=1}^k F_x^{s, \theta, e_i}(v'_{x_i}) \right) \right) \right) \\ &= \sum_{v'_{x_1} \in X_1} \left(\sum_{v'_{x_2} \in X_2} \left(\dots \sum_{v'_{x_k} \in X_k} (F_{x_1}^{s, \theta, e_i}(v'_{x_1}) \cdot F_{x_2}^{s, \theta, e_i}(v'_{x_2}) \dots F_{x_k}^{s, \theta, e_i}(v'_{x_k})) \right) \right) \\ &= \left[\sum_{v'_{x_1} \in X_1} F_{x_1}^{s, \theta, e_i}(v'_{x_1}) \right] \cdot \left[\sum_{v'_{x_2} \in X_2} F_{x_2}^{s, \theta, e_i}(v'_{x_2}) \right] \dots \left[\sum_{v'_{x_k} \in X_k} F_{x_k}^{s, \theta, e_i}(v'_{x_k}) \right] \end{aligned}$$

$$= \prod_{x_i \in \text{Var}(e_i)} \left[\sum_{v'_i \in X_i} F_{x_i}^{s, \theta, e_i}(v'_i) \right]$$

By construction, for $x_i \in \text{Var}(e_i)$, we have $\sum_{v'_i \in X_i} F_{x_i}^{s, \theta, e_i}(v'_i) = 1$. Therefore, $\sum_{s' \in S_1} \prod_{x \in \text{Var}(e_i)} F_x^{s, \theta, e_i}([s']x) = 1$.
As a consequence,

$$\begin{aligned} \sum_{s' \in S, e_i \in \text{Acts}(s)} P(s, e_i, s') &= \sum_{e_i \in \text{Acts}(s)} \left[\frac{[s]W_i(\bar{v}) \cdot \sum_{\theta \in T_s^{e_i}} P_{T_s^{e_i}}(\theta)}{\sum_{e_j \in \text{Acts}(s)} [s]W_j(\bar{v})} \right] \\ &= \frac{\sum_{e_i \in \text{Acts}(s)} [s]W_i(\bar{v}) \cdot \left(\sum_{\theta \in T_s^{e_i}} P_{T_s^{e_i}}(\theta) \right)}{\sum_{e_j \in \text{Acts}(s)} [s]W_j(\bar{v})} \end{aligned}$$

By construction, we have $\sum_{\theta \in T_s^{e_i}} P_{T_s^{e_i}}(\theta) = 1$ and thus:

$$\sum_{s' \in S, e_i \in \text{Acts}(s)} P(s, e_i, s') = \frac{\sum_{e_i \in \text{Acts}(s)} [s]W_i(\bar{v})}{\sum_{e_j \in \text{Acts}(s)} [s]W_j(\bar{v})} = 1$$

As a conclusion, we have that $\forall s, \sum_{s' \in S, e_i \in \text{Acts}(s)} P(s, e_i, s') = 1$ and then $\llbracket M \rrbracket$ is a DTMC.

B Complements to Almost-certain Convergence

B.1 Necessity of bounding event weights

In this Section, we show by means of an example of a probabilistic Event-B model the necessity of the new PO (event/wght/BOUND) introduced in order to prove the almost-certain convergence of a set of probabilistic convergent events.

Consider the probabilistic Event-B model M1 given in Figure 5. This model has two variables: x and y and three events $evt1$, $evt2$ and $evt3$, two of which ($evt1$ and $evt2$) are convergent. The variant of this model is x and the bound on the variant is clearly $U = 2$.




Figure 5: Probabilistic Event-B model M1

The DTMC semantics of M1 is given in Figure 6.




Figure 6: DTMC part of the model M1

In states where $x = 1$, only convergent events $evt1$ and $evt2$ are enabled and the local probability of choosing $evt1$ is $\frac{1}{y}$ while the local probability of choosing $evt2$ is $\frac{y-1}{y}$.

In states where $x = 2$, only $evt1$ can be chosen with probability 1.

In states where $x = 0$, the only enabled event is the (non-convergent) event $evt3$.

Clearly, the model M1 satisfies proof obligations (model/pVar), (event/var/pNAT) and (event/pBOUND). However, as we show below, the probability of eventually taking a non-convergent event is strictly smaller than 1 from all states where $x > 0$ because the probability of decreasing the variant, although strictly positive in all states, gets infinitely small from states where $x = 1$ as y increases.

W.l.o.g., we compute the probability of eventually taking $evt3$ from the initial state where $x = 1$ and $y = 2$. The reasoning starting from other states is similar. This probability is equal to the sum of

- (1) the probability of directly taking $evt1$ from $(1,2)$,
- (2) the probability of reaching $(1,4)$ and taking $evt1$ from $(1,4)$,
- (3) the probability of reaching $(1,8)$ and taking $evt1$ from $(1,8)$
- (4) ...

Clearly, (1) is equal to $\frac{1}{2}$, (2) is equal to $\frac{1}{2} \cdot \frac{1}{4} = \frac{1}{8}$, (3) is equal to $\frac{1}{2} \cdot \frac{3}{4} \cdot \frac{1}{8} < \frac{1}{16}$ and in general, the probability of reaching state $(1, 2^i)$ with $i > 2$ and taking evt1 from this state is strictly smaller than $\frac{1}{2^{i+1}}$. As a consequence, the probability of eventually taking evt3 from the initial state is strictly smaller than

$$\frac{1}{2} + \sum_{i=2}^{\infty} \frac{1}{2^{i+1}} = \frac{3}{4}$$

Therefore, M1 does not almost-certainly converge.

The behaviour we expose here is a direct consequence of the unboundedness of the weights of convergent events, which, by getting arbitrarily big, cause the probability of decreasing the variant to get arbitrarily small.

B.2 Necessity of bounding event parameter values

As in the appendix B.1, we show by means of an example that bounding parameter values through the new PO (event/param/BOUND) is necessary for the almost-certain convergence of probabilistic Event-B models. The model M2, given in Fig. 7 and its semantics, given in Fig. 8 are similar to the ones presented in appendix B.1.

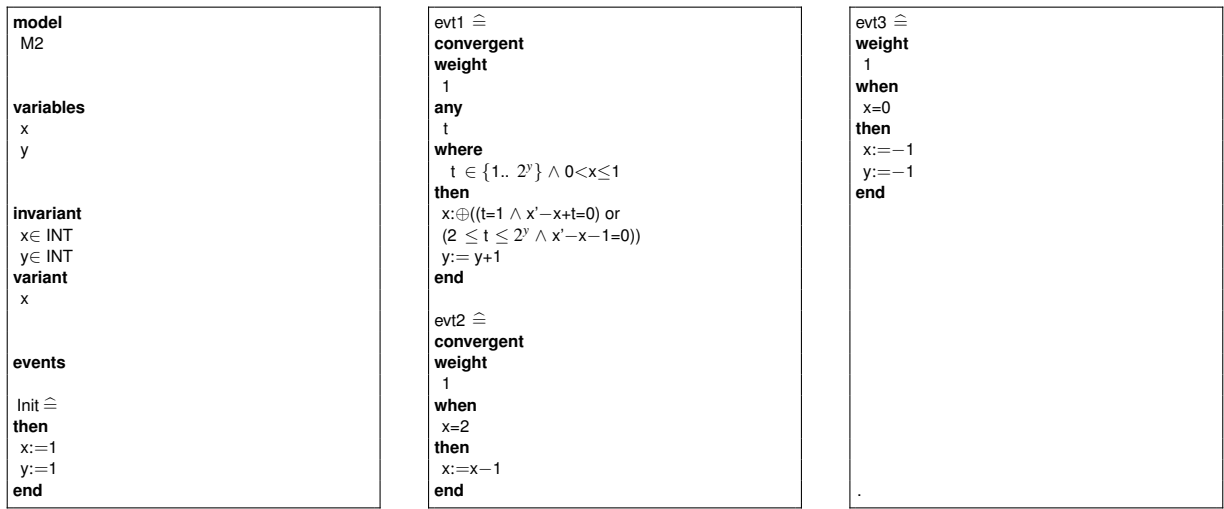


Figure 7: Probabilistic Event-B model M2

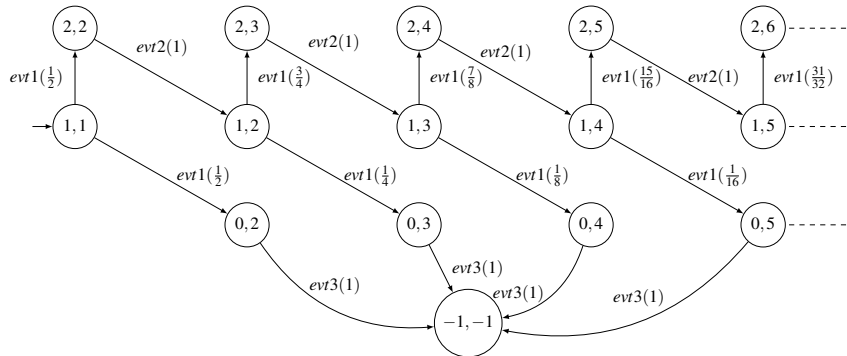


Figure 8: DTMC part of the model M2

In this case also, we observe that the probability of eventually executing non-convergent event evt3 from the initial state is strictly smaller than $\frac{3}{4}$. The main difference is that, in M2, only the choice of parameter values is responsible for infinitely decreasing the probabilities of decreasing the variant. Bounding parameter values through (event/param/BOUND) prevents this problem from happening.

B.3 Proof of Theorem 1

Let $M = (\bar{v}, I(\bar{v}), V(\bar{v}), \text{Evts}, \text{Init})$ be a probabilistic Event-B model. $\text{Evts} = \text{Evts}_c \cup \text{Evts}_{nc}$ is the partition of the set of events Evts into convergent events $\text{Evts}_c = \{e_1, \dots, e_i\}$ and non convergent events $\text{Evts}_{nc} = \{e_{i+1}, \dots, e_n\}$ ($1 \leq i < n$).

We show that if M satisfies the following convergence POs:

1. event/var/pNAT

$$\forall e \in \text{Evts}_c. I(\bar{v}) \wedge W_e(\bar{v}) > 0 \wedge G_e(\bar{t}, \bar{v}) \vdash V(\bar{v}) \in \text{NAT}$$

2. event/pBOUND

$$\forall e \in \text{Evts}_c. I(\bar{v}) \wedge W_e(\bar{v}) > 0 \wedge G_e(\bar{t}, \bar{v}) \vdash V(\bar{v}) \leq U$$

3. event/wght/BOUND

$$\forall e \in \text{Evts}_c. I(\bar{v}) \wedge G_e(\bar{t}, \bar{v}) \vdash W(\bar{v}) \leq BW$$

4. event/param/BOUND

$$\forall e \in \text{Evts}_c. I(\bar{v}) \vdash \text{card}(\{\bar{t} \mid G_e(\bar{t}, \bar{v})\}) \leq BP$$

5. model/pVar

$$I(\bar{v}) \wedge (G_1(\bar{t}, \bar{v}) \vee \dots \vee G_i(\bar{t}, \bar{v})) \vdash (\exists \bar{v}'. W_1(\bar{v}) \wedge G_1(\bar{t}, \bar{v}) \wedge S_1(\bar{t}, \bar{v}) \wedge V(\bar{v}') < V(\bar{v})) \vee \dots \\ \vee (\exists \bar{v}'. W_i(\bar{v}) \wedge G_i(\bar{t}, \bar{v}) \wedge S_i(\bar{t}, \bar{v}) \wedge V(\bar{v}') < V(\bar{v}))$$

then M almost-certainly converges (with probability 1).

Recall that almost-certain convergence of M consists in proving that, from all valuation of the variables of M where a convergent event is enabled, the probability of eventually taking a non-convergent event or reaching a deadlock is 1. In order to prove this result, we consider a slightly modified version of the DTMC semantics of M and use classical results on DTMCs in order to show that the probability of eventually reaching a given set of states is 1 from all states where non-convergent events are enabled.

Proof.

In order to take into account the difference between convergent and non-convergent events, we propose the following slightly modified version of the DTMC semantics of M . In this version, all the states are replicated in order to “remember” the last event executed.

Formally, consider the probabilistic Event-B model M introduced above and let $\llbracket M \rrbracket = (S, s_0, AP, L, \text{Acts}, P)$ be the DTMC semantics of M as introduced in Definition 2. We build the DTMC $\llbracket M \rrbracket' = (T, t_0, AP, L', \text{Acts}, P')$ where

- $T \subseteq S \times (\text{Acts} \cup \{\varepsilon\})$ is the set of extended states, consisting in pairs (s, a) where s is a state of $\llbracket M \rrbracket$ and a is an action (event name),
- $t_0 = (s_0, \varepsilon)$ is the initial state,
- L' is such that $L'((s, a)) = L(s)$ for all $s \in S$ and $a \in \text{Acts}$, and
- P' is such that $P'((s, a), e, (s', b)) = P(s, e, s')$ if $e = b$ and 0 otherwise for all action a .

It is easy to see that M almost-certainly converges iff the probability of eventually reaching either a deadlock state or an extended state of the form $t = (s, e)$ where e is a non-convergent event is 1 in $\llbracket M \rrbracket'$ from all (extended) states where convergent events are enabled.

Since $\llbracket M \rrbracket$ has a potentially infinite set of states, showing such a result is not trivial. In order to prove it, we therefore exploit existing results from the theory of DTMCs. In particular, we focus on the global coarseness property introduced in [14], which is a sufficient condition for the “decisiveness” of infinite-state Markov Chains. Formally, given a Markov Chain $\mathcal{M} = (\mathcal{S}, \mathcal{P})$ and a target set of states $\mathcal{F} \subseteq \mathcal{S}$, we say that \mathcal{M} is globally coarse w.r.t. \mathcal{F} iff there exists some minimal bound $\alpha > 0$ such that for all state $s \in \mathcal{S}$, the probability of eventually reaching \mathcal{F} from s is either 0 or greater or equal to α . It is then shown in [14] that whenever a Markov Chain \mathcal{M} is globally coarse w.r.t. the set \mathcal{F} , the probability of eventually reaching either \mathcal{F} or a set of states $(\tilde{\mathcal{F}})$ from which \mathcal{F} cannot be reached is 1 from any state of \mathcal{M} .

In the following, we will apply this result to the DTMC $\llbracket M \rrbracket'$ in order to prove that M almost-certainly converges. We therefore proceed as follows:

- (a) We start with introducing notations that will be used throughout the proof.
- (b) We then propose a partition of the extended states T of $\llbracket M \rrbracket'$ and introduce our goal set $F \subseteq T$.
- (c) We show that all states from each partition of T satisfy the global coarseness property w.r.t. F .
- (d) We finally show that the set \tilde{F} is empty and conclude.

We now detail each step of this proof.

- (a) Consider the following notations:

- In the DTMC $\llbracket M \rrbracket'$, we partition the set of actions (event names) as follows: $\text{Acts} = \text{Acts}_{nc} \cup \text{Acts}_c$, where Acts_{nc} is the set of non convergent actions and Acts_c is the set of convergent actions.
- Given an extended state t and a set of states $G \subseteq T$, we write $P(t \models \diamond G)$ for the probability of eventually reaching G from t .
- Given a predicate P and an extended state $t = (s, a)$ of $\llbracket M \rrbracket'$, we write $P(t)$ for the evaluation of P in the state s .
- Given an extended state $t = (s, a) \in T$, we write $\text{Acts}(t)$ for the set of events enabled in s . Similarly, we write $\text{Acts}_c(t)$ for the set of convergent events enabled in s and $\text{Acts}_{nc}(t)$ for the set of non convergent events enabled in s .
- Given a set of events E and a state $t = (s, a) \in T$, we write $W^t(E)$ (or $W^s(E)$ when clear from the context) for the sum of the weights of the events from E that are enabled in s .
- Given a state $t = (s, a) \in T$, we write $\text{Succ}(t)$ for the set of extended states that are reached from t :

$$\text{Succ}(t) = \{t' \in T \mid \exists e \in \text{Acts}(t). P^t(t, e, t') > 0\}$$

- Given a finite execution $\sigma = t_0, e_0, t_1, \dots, t_{n-1}, e_{n-1}, t_n$ of $\llbracket M \rrbracket'$, the length of σ is written $L(\sigma)$ and is equal to the number of transitions executed in σ . In the above example case, $L(\sigma) = n$.

- (b) We now introduce the following sets of extended states T :

- $T_1 = \{t = (s, a) \in T \mid \exists e \in \text{Evts}_c, \exists \theta \in T_s^e, G_e(s, \theta) \wedge \forall e' \in \text{Evts}_{nc}, \forall \theta \in T_s^{e'}, \neg G_{e'}(s, \theta)\}$ is the set of extended states where only convergent events are enabled.
- $T_2 = \{t = (s, a) \in T \mid \exists e \in \text{Evts}_c, \exists \theta \in T_s^e, G_e(s, \theta) \wedge \exists e' \in \text{Evts}_{nc}, \exists \theta \in T_s^{e'}, G_{e'}(s, \theta)\}$ is the set of states where both convergent and non convergent events are enabled.
- $T_3 = \{t = (s, a) \in T \mid \forall e \in \text{Evts}_c, \forall \theta \in T_s^e, \neg G_e(s, \theta)\}$ is the set of states where no convergent events are enabled.
- $T_4 = \{t = (s, a) \in T \mid a \in \text{Evts}_{nc}\}$ is the set of states reached after performing a non convergent event.

It is easy to see that $T = T_1 \cup T_2 \cup T_3$ defines a partition of T . The convergence property for our probabilistic Event-B model M clearly concerns states from T_3 and T_4 . We therefore define our target set as $F = T_3 \cup T_4$. As in [14], we write \tilde{F} for the subset of states of T from which it is impossible to reach F . We show later that \tilde{F} is empty.

- (c) We now show that all extended states in T_1 and T_2 and T_3 satisfy the global coarseness property w.r.t F , i.e. that there exists a minimal bound $\alpha > 0$ such that for each extended state $t \in T$, the probability of eventually reaching F is either 0 or greater or equal to α .

- We begin with states in T_2 . Let $t_2 = (s_2, a) \in T_2$. Let F_2 be the subset of states that are reached from t_2 by non convergent events. Obviously, $F_2 \subseteq T_4 \subseteq F$. Formally,

$$F_2 = \{t' = (s', a') \in T \mid t' \in \text{Succ}(t_2) \wedge a' \in \text{Acts}_{nc}\}$$

By definition of T_2 , at least one convergent event is enabled in t_2 , therefore we have $W^{t_2}(\text{Acts}_c) > 0$. Likewise, at least one non convergent event can be enabled in t_2 , thus $W^{t_2}(\text{Acts}_{nc}) > 0$. Therefore $W^{t_2}(\text{Acts}) > 0$.

Recall from section 5 that the probability of a transition (t_2, e, t') where $e \in \text{Acts}_{nc}(t_2)$ and $t' = (s', e) \in F_2$ is given by:

$$P'(t_2, e, t') = P(s_2, e, s') = \frac{W_e(s_2)}{W^{s_2}(\text{Acts})} \times \sum_{\theta \in T_{s_2}^e} [P_{T_{s_2}^e}(\theta) \times \prod_{x \in \text{Var}(e)} P_{s_2, \theta}^e(x, s')]$$

By definition, all non convergent events e take the system in states in F_2 regardless of the probabilistic choice made inside the action of e . Therefore:

$$\sum_{e \in \text{Acts}_{nc}(s_2), t' \in F_2} P(t_2, e, t') = \sum_{e \in \text{Acts}_{nc}(s_2)} \frac{W_e(s_2)}{W^{s_2}(\text{Acts})} \times 1$$

Therefore, the probability of eventually reaching F_2 from t_2 is above $\frac{W^{s_2}(\text{Acts}_{nc})}{W^{s_2}(\text{Acts})}$.

We now show by contradiction that there exists $\alpha_2 > 0$ s.t $\forall t_2 \in T_2, P(t_2 \models \diamond F_2) \geq \alpha_2$.

Assume the contrary, i.e. $\forall \alpha_2 > 0, \exists t_2 \in T_2$ s.t $P(t_2 \models \diamond F_2) < \alpha_2$.

Let α_2 be such that $(\frac{1}{\alpha_2} - 1) > BW \times \text{card}(\text{Acts}_c)$. There must exist $t_2 = (s_2, a) \in T_2$ such that $P(t_2 \models \diamond F_2) < \alpha_2$. By the result above, we know that $P(t_2 \models \diamond F_2) \geq \frac{W^{s_2}(\text{Acts}_{nc})}{W^{s_2}(\text{Acts})}$. As a consequence, we must have:

$$\frac{W^{s_2}(\text{Acts}_{nc})}{W^{s_2}(\text{Acts})} < \alpha_2$$

Recall that $W^{s_2}(\text{Acts}) = W^{s_2}(\text{Acts}_{nc}) + W^{s_2}(\text{Acts}_c)$. Therefore,

$$\frac{W^{s_2}(\text{Acts})}{W^{s_2}(\text{Acts}_{nc})} = 1 + \frac{W^{s_2}(\text{Acts}_c)}{W^{s_2}(\text{Acts}_{nc})} > \frac{1}{\alpha_2}$$

As a consequence,

$$W^{s_2}(\text{Acts}_c) > W^{s_2}(\text{Acts}_{nc}) \cdot \left(\frac{1}{\alpha_2} - 1\right)$$

By definition of T_2 , we have $W^{s_2}(\text{Acts}_{nc}) \geq 1$, therefore

$$W^{s_2}(\text{Acts}_c) > \left(\frac{1}{\alpha_2} - 1\right)$$

Finally, by definition of α_2 , we have $W^{s_2}(\text{Acts}_c) > BW \times \text{card}(\text{Acts}_c)$, which is clearly in contradiction with PO event/wght/BOUND.

We therefore conclude that there exists $\alpha_2 > 0$ such that $\forall t_2 \in T_2, P(t_2 \models \diamond F_2) \geq \alpha_2$.

- We now move to extended states in T_1 : we show that there exists α_1 such that for all extended states $t_1 \in T_1, P(t_1 \models \diamond F) \geq \alpha_1$.

Recall that the probability function of $\llbracket M \rrbracket'$ is expressed as follows: For all $t_1 = (s_1, a) \in T_1, e \in \text{Acts}$, and $t' = (s', a) \in T$, we have

$$P(t_1, e, t') = P(s_1, e, s') = \frac{W_e(s_1)}{W^{s_1}(\text{Acts})} \times \sum_{\theta \in T_{s_1}^e} [P_{T_{s_1}^e}(\theta) \times \prod_{x \in \text{Var}(e)} P_{s_1, \theta}^e(x, s')]$$

Since $t_1 \in T_1$, this expression can only be non-zero if e is a convergent event. In this case, PO event/wght/BOUND ensures that $W^{s_1}(\text{Acts}) \leq BW \cdot \text{card}(\text{Acts}_c)$. Therefore, for all convergent events enabled in t_1 , we have $\frac{W_e(s_1)}{W^{s_1}(\text{Acts})} \geq \frac{1}{BW \cdot \text{card}(\text{Acts}_c)}$.

Moreover, PO event/param/BOUND ensures that the number of parameter valuations satisfying the guard of e in s_1 is bounded by BP . As a consequence,

$$\sum_{\theta \in T_{s_1}^e} [P_{T_{s_1}^e}(\theta) \times \prod_{x \in \text{Var}(e)} P_{s_1, \theta}^e(x, s')] \geq \frac{1}{BP} \times \sum_{\theta \in T_{s_1}^e} [\prod_{x \in \text{Var}(e)} P_{s_1, \theta}^e(x, s')]$$

Finally, since the probabilities inside each probabilistic assignment ($P_x^e(E)$) are constant and in finite number, there is a minimal value $\beta > 0$ (which we do not detail here) such that for all $t_1 = (s_1, a) \in T_1, e \in \text{Acts}_c$, and $t' = (s', e) \in T$, whenever $P(t_1, e, t') > 0$, we have

$$\sum_{\theta \in T_{s_1}^e} [\prod_{x \in \text{Var}(e)} P_{s_1, \theta}^e(x, s')] \geq \beta$$

As a consequence, there exists a minimal value $\gamma > 0$ such that $P'(t_1, e, t') \geq \gamma$ for all $t_1 \in T_1, e \in \text{Acts}_c$, and $t' \in T$ such that $P'(t_1, e, t') > 0$.

Now, let $t_0 = (s_0, a_0) \in T_1$ be an extended state. By definition of T_1 and because of POs *event/p-BOUND*, *event/var/pNAT* and *model/pVar*, the value of the variant in t_0 is between 0 and U and there must exist a transition that leads the system to an extended state $t_1 = (s_1, a_1)$ s.t. $V(t_1) < V(t_0)$. Necessarilly, we have $t_1 \in T_1$ or $t_1 \in T_2$ or $t_1 \in T_3$, therefore there must exist a finite execution $\sigma = t_0, e_0, t_1, \dots, t_{n-1}, t_{n-1}, t_n$ with $t_n \in T_2 \cup T_3$ and $\forall i < n, t_i \in T_1$ and $L(\sigma) \leq U + 1$.

If $t_n \in T_3 \subseteq F$, then $P(t_0 \models \diamond F) \geq \gamma^{U+1}$. Otherwise, we have $t_n \in T_2$ and $P(t_n \models \diamond F) \geq \alpha_2$, therefore $P(t_0 \models \diamond F) \geq \alpha_2 \cdot \gamma^{U+1}$.

As a consequence, since $\alpha_2 \leq 1$, we have $\gamma^{U+1} \geq \alpha_2 \cdot \gamma^{U+1}$ and there exists $\alpha_1 = \alpha_2 \cdot \gamma^{U+1} > 0$ such that for all extended states $t_1 \in T_1, P(t_1 \models \diamond F) \geq \alpha_1$.

- Finally, since $T_3 \subseteq F$, we have $P(t_3 \models \diamond F) = 1$ for all extended states $t_3 \in T_3$.

We therefore conclude that $\llbracket M \rrbracket'$ is globally coarse w.r.t F . As a consequence, $\forall t \in T, P(t \models \diamond F \vee \diamond \tilde{F}) = 1$.

- (d) We have shown above that for all extended states either in T_1, T_2 or T_3 , we have $P(t \models \diamond F) > 0$. Since $T = T_1 \cup T_2 \cup T_3$, \tilde{F} is therefore necessarilly empty.

Since $\llbracket M \rrbracket'$ is globally coarse w.r.t F and \tilde{F} is empty, we have that for all extended state $t \in T$, the probability of eventually reaching the target set F is 1. As a consequence, the probability of eventually reaching either a deadlock state or an extended state of the form $t = (s, e)$ where e is a non-convergent event is 1 in $\llbracket M \rrbracket'$ from all (extended) states where convergent events are enabled, which concludes our proof. \square