

A Probabilistic Extension of UML Statecharts Specification and Verification

David N. Jansen¹, Holger Hermanns², and Joost-Pieter Katoen²

¹ Universiteit Twente, Information Systems Group
Postbus 217, 7500 AE Enschede, The Netherlands
`dnjansen@cs.utwente.nl`

² Universiteit Twente, Formal Methods and Tools Group
Postbus 217, 7500 AE Enschede, The Netherlands
`{hermanns,katoen}@cs.utwente.nl`

Abstract. This paper introduces means to specify system randomness within UML statecharts, and to verify probabilistic temporal properties over such enhanced statecharts which we call probabilistic UML statecharts. To achieve this, we develop a general recipe to extend a statechart semantics with discrete probability distributions, resulting in Markov decision processes as semantic models. We apply this recipe to the requirements-level UML semantics of [8]. Properties of interest for probabilistic statecharts are expressed in PCTL, a probabilistic variant of CTL for processes that exhibit both non-determinism and probabilities. Verification is performed using the model checker PRISM. A model checking example shows the feasibility of the suggested approach.

Keywords: Markov decision processes, model checking, probabilities, semantics, UML statecharts.

1 Introduction

The Unified Modelling Language (UML) is more and more pervading system design and engineering. Accordingly, it is not far fetched to predict that the coming years shall see substantial efforts to extend the UML and the accompanying design methodology towards soft real-time, fault-tolerance, quality of service and the like. First work in this direction has been undertaken, e. g., in [5,9,14,16].

One of the principal modelling paradigms needed to express such aspects is the concept of *probability*, allowing one to quantitatively describe the randomness the system is exposed to, the randomness the system itself exhibits, or both. For instance, probability is a useful means to describe varying workload, to quantify uncertainty in the system timing, as well as to properly model randomised distributed algorithms. Furthermore, probability is also an abstraction means, e. g., it allows one to hide data dependencies by just representing the likelihood of particular branches to be taken.

There are two facets of the probability concept that are worth to be distinguished. On the one hand, each reactive system is exposed to external stimuli that exhibit some kind of randomness. We call this *environmental randomness*.

On the other hand, the system behaviour itself may ask for a probabilistic description, either because the system embodies a randomised algorithm, or because probability is used for abstraction. We call this *system randomness*.

This paper addresses system randomness. It introduces probabilistic UML-statecharts as a means to support the design of probabilistic systems. More concretely, the paper extends statecharts by probabilistic elements: a transition is allowed to lead to one of several states depending on a probability distribution; each probability distribution is guarded by a trigger, inspired by [20]. The interference of probabilities, priorities and nondeterminism raises some subtle semantic issues. We attack these issues in a way that allows one still to employ an arbitrary priority scheme to resolve or reduce nondeterminism. The semantics is formally defined as a mapping on (strictly) alternating probabilistic transition systems [10], a subset of Markov decision processes (MDP) [19]. To allow verification of probabilistic temporal properties over probabilistic UML-statecharts, properties are expressed in the probabilistic branching-time temporal logic PCTL [1], the prime logic for property specification and verification of models that exhibit both probabilities and nondeterminism. These properties can be checked using the model checker PRISM [15].

Care is taken to achieve a conservative extension of standard UML statecharts. Among the various published semantics for statecharts, we take as a representative the requirements-level semantics of Eshuis and Wieringa [8], which is based on the semantics by Damm et al. [6] which in turn formalises the StateMate semantics of statecharts [13]. A requirements-level model focuses on the design; an implementation-level model describes the implementation of a design. Requirements-level semantics mostly use the perfect technology assumption, which abstracts from limitations (in speed and memory) imposed by an implementation [18]. The chosen semantics combines the StateMate semantics with communication and classification. We have chosen it because it is simple and close to the most used semantics for UML. A detailed justification of this semantics and comparisons to StateMate semantics as well as implementation-level semantics can be found in [8]. The setup of our probabilistic extension, however, is independent of the UML basis we take. This means that other formal statechart semantics can equally well be enhanced with a similar probabilistic extension, as long as certain principal conditions are respected. We give an account of these conditions in Sect. 7.

We omit some minor features of the semantics, that could be added easily but would clutter up the exposition. These features include actions on initial transitions and entry and exit actions. We also leave out real-time aspects. To facilitate model-checking of the underlying model, we confine ourselves to bounded integer variables and employ a closed world assumption.

In summary, this paper makes the following contributions. (i) It introduces a generic recipe to conservatively extend a statechart dialect with probabilistic features. (ii) It details this recipe for the requirement-level semantics of [8], and identifies subtle interferences between the imposed priority scheme, and the order of resolving nondeterminism and probabilistic choice. (iii) It proposes to use the

probabilistic logic PCTL to specify properties over statecharts and shows how model checking of probabilistic statecharts can be performed effectively.

Organisation of the paper. Sect. 2 and 3 introduce syntax and semantics of probabilistic statecharts. In Sect. 4, we show that this semantics conservatively extends the non-probabilistic statechart semantics. Sect. 5 presents a logic for P-statecharts and Sect. 6 demonstrates model checking with a larger example. Sect. 7 discusses our approach in the broader context of statechart semantics and of probabilistic models.

2 Probabilistic UML Statecharts

This section introduces probabilistic UML-statecharts (P-statecharts, for short), together with some drawing conventions. We first fix some notations. We assume familiarity with basic measure and probability theory [21].

Notation. The powerset of a set A is denoted by $\mathbb{P}(A)$. A probability space is denoted (Ω, \mathcal{A}, P) where Ω is the set of possible outcomes of the probabilistic experiments, $\mathcal{A} \subseteq \mathbb{P}(\Omega)$ is a σ -algebra of measurable sets and $P : \mathcal{A} \rightarrow [0, 1]$ is a probability measure. For discrete probability spaces, we sometimes write (Ω, P) instead of $(\Omega, \mathbb{P}(\Omega), P)$.

Collection of statecharts. A system consists of a finite collection of communicating statecharts. In the following, we assume a given finite collection of P-statecharts, denoted by $\{PSC_1, \dots, PSC_n\}$.

Syntax. A single P-statechart PSC_i consists of the following elements:

- A finite set $Nodes_i$ of nodes with a tree structure, described by a function $children_i : Nodes_i \rightarrow \mathbb{P}(Nodes_i)$. (If $x' \in children_i(x)$, then x is the parent of x' . Of course, $children_i$ has to fulfil several constraints to make it describe a tree structure. For simplicity, these are omitted here.) Descendants are children or children of children, etc. The root is denoted $root_i$. The function $type_i : Nodes_i \rightarrow \{\text{BASIC}, \text{AND}, \text{OR}\}$ assigns to every node its type. Nodes that are leaves of the tree have type BASIC; children of AND nodes have type OR; $type_i(root_i) = \text{OR}$; other nodes have type OR or AND. The partial function $default_i : Nodes_i \dashrightarrow Nodes_i$ identifies for each OR node one of its children as the default (or initial) node.
- A finite set $Events$ of events. (Note that the set of events is identical for all P-statecharts in the collection.) We will use the symbol \perp to denote “no event required”; $\perp \notin Events$.
- A finite set $Vars_i$ of variables together with an initial valuation $V_{0,i} : Vars_i \rightarrow \mathbb{Z}$, which assigns initial values to the variables. (We only allow bounded integer variables.)
- A set $Guards_i$ of guard expressions. Guard expressions are boolean combinations of the atoms $j.isin(x)$, for $j \in \{1, \dots, n\}$ and $x \in Nodes_j$ (with the intuitive meaning “the P-statechart PSC_j is in node x ”), and comparisons like $expr \leq expr$ and $expr > expr$, for arithmetic expressions made up from the variables and integer constants.

- A set $Actions_i$ of actions. Actions are $v := expr$, which denotes an assignment to $v \in Vars_i$, and **send** $j.e$ (for $j \in \{1, \dots, n\}$ and $e \in Events$) with the intuitive meaning “send event e to the P-statechart PSC_j ”.
- A finite set $PEdges_i$ of P-edges. A P-edge is a tuple (X, e, g, P) where $X \subseteq Nodes_i$ is a non-empty set of source state nodes, $e \in Events \cup \{\perp\}$, $g \in Guards_i$ is a guard, and P is a probability measure in the discrete probability space $(\mathbb{P}(Actions_i) \times (\mathbb{P}(Nodes_i) \setminus \{\emptyset\}), P)$. We assume that there is a bijective index function $\iota : \{1, \dots, |PEdges_i|\} \rightarrow PEdges_i$ to simplify the identification of P-edges.

Note that $Guards_i$ and $Actions_i$ are defined in terms of the other components. We will therefore denote a P-statechart simply by $(Nodes_i, Events, Vars_i, PEdges_i)$.

The above definition differs from the usual statecharts syntaxes (e. g., [8]) in the way edges are refined into P-edges. A P-edge (X, e, g, P) can be considered as a hyperedge with source node set X , and possibly multiple targets (A, Y) , each target having a certain probability $P(A, Y)$. Note that a target is an action set A together with a non-empty set Y of successor nodes. Once the P-edge is triggered by event e and guard g holds in node X , a target (A, Y) is selected with probability $P(A, Y)$.

Drawing a P-Statechart. We adopt the following drawing conventions. The root node is not drawn. Nodes that are not children of an AND-node are drawn as rectangles with rounded corners. A parent node encloses its children. Children of an AND-node partition the node by dashed lines. Each OR-node encloses a black dot and indicates its default node by an arrow directed from the dot to the default node. A trivial P-edge (where probability 1 is assigned to a unique action set/node set) with event e , guard g and action set A is denoted as an arrow $\xrightarrow{e[g]/A}$. A P-edge possessing a non-trivial probability space consists of two parts: first an arrow with event and guard $\xrightarrow{e[g]}$ that points to a symbol \textcircled{P} (a so-called P-pseudonode), then several arrows emanating from the P-pseudonode, each with a probability and an action set $\xrightarrow{p/A}$. This notation is inspired by C-pseudonodes \textcircled{C} , used for case selection purposes e. g., in [12]. If the event on a P-edge is \perp , we may omit it. Further, if the guard is *true*, we may omit it. Similarly an empty set of actions may be omitted.

Example 1. Figure 1 depicts a P-statechart which shows the behaviour when playing with an *unreliable, but fair coin*: the event “flip” may or may not be ignored. If the system reacts, it outputs “heads” or “tails”, each with 50% chance. If the output is heads, the system stops playing. It is unspecified how (un)reliable the system is.

3 P-Statechart Semantics

This section discusses the semantics of P-Statechart, which is an adaptation of the nonprobabilistic semantics in [8]. The semantics is defined in two phases.




Fig. 1. Example P-statechart. The labels printed in italics (*ignore* etc.) are not part of the diagram, but are included to facilitate referencing the edges near the labels.

First, it is defined what will be a step. This encompasses the resolution of non-determinism, probabilistic choice and priorities within a single P-statechart. Subsequently, these steps are used as the building blocks in a mapping of a collection of P-statecharts onto a Markov decision process.

Closed-world assumption. Opposed to [8] we assume a closed system model in the sense that we do *not* consider the system to be subject to inputs from the environment. This “closed-world assumption” simplifies model checking. The “open-world” semantics of [8] provides for input from the environment of the statecharts, and therefore, their state transitions consist of several phases, cluttering up the semantics. However, if one wants to consider specific environmental inputs, one can easily add another component that generates the desired events.

Intuitive semantics for a single P-statechart. The intuitive behaviour of a P-statechart can be described as follows. The statechart is always in some state (which consists of one or several nodes). A P-edge is taken if the P-statechart is in the source node(s), the event of the edge happens and its guard holds. Then, the system chooses one of the possible results (probabilistically and non-deterministically); it leaves the source nodes, executes the chosen action and enters the chosen target nodes of the P-edge. More than one edge may be taken simultaneously, even within a single statechart.

3.1 Step Construction

This section describes how a step is constructed for a single P-statechart PSC_i .

Configurations and states. A configuration C_i of P-statechart PSC_i is a set of nodes that fulfils the conditions:

- $root_i \in C_i$.
- If an OR-node is in C_i , then exactly one of its children is in C_i .
- If an AND-node is in C_i , then all its children are in C_i .

The set of all configurations of PSC_i is denoted $Conf_i$. A state of PSC_i is a triple (C_i, I_i, V_i) where C_i is a configuration, $I_i \subseteq Events$ is a set of events (to which the statechart still has to react), and $V_i : Vars_i \rightarrow \mathbb{Z}$ is a valuation of

the variables. The set of all valuations of PSC_i is denoted Val_i . The validity of guard g in a state depends on the configurations C_1, \dots, C_n and the valuations V_1, \dots, V_n . We write $(C_{1\dots n}, V_{1\dots n}) \models g$ iff g holds in the state of the collection of P-statecharts.

Edges. An edge is a triple (j, A, Y) , where j identifies a P-edge, $A \subseteq Actions_i$ is a set of actions and $Y \subseteq Nodes_i$ is a set of target nodes. The set $Edges_i$ is defined as: $\{(j, A, Y) \mid \exists X, e, g, P : \iota(j) = (X, e, g, P) \in PEdges_i \wedge P(\{(A, Y)\}) > 0\}$.

Scope. The *scope* of an edge (j, A, Y) is the smallest (in the parent-child-hierarchy) OR-node that contains both the source nodes $\iota(j).X$ and the target nodes Y . Since this node only depends on source and target nodes, we refer to it as $scope(\iota(j).X, Y)$. (The scope is the smallest node that is not affected when the edge is executed.)

Steps. A step is a set of edges that are taken together as a reaction to events. The edges in a step for P-statechart PSC_i depend on its current state (C_i, I_i, V_i) . A step has to obey several constraints, which are in close correspondance to [8]:

Enabledness. All edges in the step must be enabled. A P-edge (X, e, g, P) is enabled if the current configuration C_i contains its source state nodes X , the event e is in the current input set I_i and the guard g holds: $(C_{1\dots n}, V_{1\dots n}) \models g$. An edge is enabled if its corresponding P-edge is enabled¹.

Consistency. All edges in the step must be pairwise consistent. This means that they are either identical or that their scopes are different children of some AND-node or their descendants (in the latter case, the scopes are called orthogonal in [8]). If two edges are not consistent, they are called inconsistent.

Priority. We assume a given priority scheme (a partial order on the edges) that resolves some of the inconsistencies: If an enabled edge e is not in the step, then there must be an edge in the step that is inconsistent with e and does not have lower priority than e .

Maximality. A step must be maximal. This means that adding any edge leads to a violation of the above conditions².

We now give an algorithm to construct a step of a single statechart which – by construction – satisfies the conditions above. The algorithm employs a specific order with respect to the resolution of nondeterminism and probabilities. Assume that the current state is (C_i, I_i, V_i) .

1. Calculate the set of enabled P-edges: for $j \in \{1, \dots, |PEdges_i|\}$,

$$j \in EnP(C_i, I_i, V_i) \text{ iff } \iota(j).X \subseteq C_i \wedge (\iota(j).e \in I_i \cup \{\perp\}) \wedge (C_{1\dots n}, V_{1\dots n}) \models \iota(j).g$$

¹ It may happen that no P-edge, and consequently no edge, is enabled. Because of the closed world assumption, this is a deadlock.

² If there is no unique maximum, the system is nondeterministic.

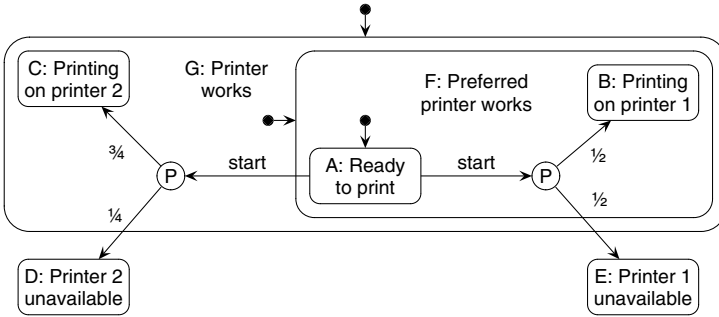


Fig. 2. Example of priority depending on target state

2. Draw samples from the probability spaces of the enabled P-edges, reducing the set $EnP(C_i, I_i, V_i)$ to a set $En(C_i, I_i, V_i)$ of enabled edges.
3. Calculate $Steps(En(C_i, I_i, V_i))$, where $Steps(E)$ (for $E \subseteq Edges_i$) contains all maximal, prioritized, consistent sets of edges $\subseteq E$.
4. Choose nondeterministically an element of $Steps(En(C_i, I_i, V_i))$.

Task 2 can be formalised as follows: For every state (C_i, I_i, V_i) , we define a discrete probability space $\mathcal{PR}_{(C_i, I_i, V_i)}$ over $\mathbb{P}(Edges_i)$. Its probability measure is the lift of the following probability weight to sets of sets: for any selection of A_j and Y_j (for $j \in EnP(C_i, I_i, V_i)$),

$$P(\{(j, A_j, Y_j) \mid j \in EnP(C_i, I_i, V_i)\}) = \prod_{j \in EnP(C_i, I_i, V_i)} (\iota(j) \cdot P)(A_j, Y_j)$$

and $P(E) = 0$ otherwise. Note that if $EnP(C_i, I_i, V_i) = \emptyset$, then P is the trivial probability weight such that $P(\emptyset) = 1$.

Tasks 3 and 4 can be achieved by applying the original algorithm for *nextstep* (see [8]) to the calculated set $En(C_i, I_i, V_i)$. It is a nondeterministic algorithm that calculates a maximal, prioritized, consistent step, given a set of enabled edges. As a consequence, the above algorithm (consisting of Tasks 1–4) leads to a step that is enabled, consistent, prioritized and maximal.

It is worth to highlight that (after calculating the enabled possibilities in Task 1), we first choose probabilistically (in Task 2) according to the probabilities given by the P-edges. Only after that, in Tasks 3 and 4, we resolve the nondeterminism between the remaining possibilities. This order – first resolve probabilism, then nondeterminism – is essential, as shown by the following two examples, and can only be reversed by sacrificing the expressivity of P-statecharts, e.g., by disallowing arrows to cross node boundaries from a P-pseudonode to a target node.

Priority depends on probabilistic choices. In the mostly used priority schemes, priority depends on the scope. For example, in the UML priority scheme, smaller (in the parent–child–hierarchy) scopes have higher priority, according to [12]. The P-statechart in Fig. 2 describes a fragment of a system with two printers, of which




Fig. 3. Example of consistency depending on target state

the preferred printer (printer 1) is available only with probability $\frac{1}{2}$, and the other (printer 2) is available in $\frac{3}{4}$ of the cases. The probabilities are independent. If a print request is started, the system directs it to the best available printer. The edge leading from A: Ready to print to C: Printing on printer 2 (denoted $A \rightarrow C$) with scope G: Printer works has higher priority than $A \rightarrow D$ and $A \rightarrow E$ (with scopes *root*), but $A \rightarrow C$ has lower priority than the edge $A \rightarrow B$ (with scope F: Preferred printer works), because F: Preferred printer works is a descendant of G: Printer works. So, if in configuration $\{A, F, G, \text{root}\}$ event start happens, the step produced by the above algorithm is:

- $\{A \rightarrow B\}$ with probability $P(\{A \rightarrow B, A \rightarrow C\}) + P(\{A \rightarrow B, A \rightarrow D\}) = \frac{1}{2}$, as edge $A \rightarrow B$ has priority over all other edges.
- $\{A \rightarrow C\}$ with probability $P(\{A \rightarrow E, A \rightarrow C\}) = \frac{1}{2} \cdot \frac{3}{4} = \frac{3}{8}$, as $A \rightarrow C$ has priority over $A \rightarrow E$.
- Either $\{A \rightarrow D\}$ or $\{A \rightarrow E\}$ with probability $P(\{A \rightarrow E, A \rightarrow D\}) = \frac{1}{2} \cdot \frac{1}{4} = \frac{1}{8}$. The choice between these two steps is purely nondeterministic.

As the edges that belong to one P-edge have different scopes, it is impossible to resolve the priorities prior to resolving the probabilistic choice. Although this is exemplified using the priority scheme of [12], a similar P-statechart can be drawn for the Statemate priority scheme (which also depends on the scope).

The priority scheme of [22] does not depend on the scope, but only on the source nodes. For such a priority scheme, the above phenomenon is not a problem; but the next example is independent from priorities.

Consistency depends on probabilistic choices. The P-statechart in Fig. 3 shows a system which reacts to an event e in two independent components, of which one causes an error with probability $\frac{1}{4}$. The edge $A \rightarrow \text{Error}$ is inconsistent with $B \rightarrow D$, as the scopes *root* and G are not (descendants of) different children of an AND-node (orthogonal). So, if in configuration $\{A, B, F, G, H, \text{root}\}$ event e happens, the step produced by the above algorithm is:

- $\{C, D, F, G, H, \text{root}\}$ with probability $P(A \rightarrow C) = \frac{3}{4}$.
- Either $\{A, D, F, G, H, \text{root}\}$ or $\{\text{Error}, \text{root}\}$ with probability $P(A \rightarrow \text{Error}) = \frac{1}{4}$; most priority schemes only allow one of the two cases.

Thus, the probability of taking edge $B \rightarrow D$ as a reaction to event e depends on the resolution of the probabilistic choice in the parallel node F . It is impossible to resolve the nondeterminism first, as there may or may not be inconsistent edges.

In summary, the influence of the target state in the construction of a step, as present in both the consistency definition and the priority scheme, forces us to resolve probabilism *prior* to establishing consistency and priority.

3.2 Step Execution

After having settled how steps are selected within a single P-statechart, we now consider their joint execution in the collection $\{PSC_1, \dots, PSC_n\}$. The execution of a step is the same as in [8], as probabilistic aspects are not involved anymore. On the level of a single statechart, executing a step consists of two parts: updating the variables and events occurring in the actions and determining the new state. As the actions of one P-statechart may influence the sets of events of other P-statecharts, we describe the step execution of the complete collection of P-statecharts.

Default completion. The *default completion* C' of some set of nodes C is the smallest superset of C such that C' is a configuration. If C' contains an OR-node x but C contains none of its descendants, C' contains its default node $default_i(x)$.

Executing a step. Given configurations (C_1, \dots, C_n) , steps (T_1, \dots, T_n) , and valuations (V_1, \dots, V_n) , we define for P-statechart PSC_i the new state (C'_i, I'_i, V'_i) by:

- C'_i is the default completion of the union of $\bigcup_{(j,A,Y) \in T_i} Y$ and $\{x \in C_i \mid \forall (j, A, Y) \in T_i : x \text{ is not a descendant of } scope(\iota(j).X, Y)\}$.
- $I'_i = \bigcup_{k=1}^n \{e \mid \exists (j, A, Y) \in T_k : \mathbf{send} \ i.e \in A\}$
- $V'_i = V_i[\{v := expr \mid v := expr \in A, (j, A, Y) \in T_i\}]$, the same valuation as before except for the assignments in any action of the step. If these assignments are inconsistent, V'_i is undefined.

We denote this as: $Execute(C_{1..n}, T_{1..n}, V_{1..n}) = ((C'_1, I'_1, V'_1), \dots, (C'_n, I'_n, V'_n))$.

3.3 Markov Decision Process Semantics

Recall that in the step construction algorithm we resolve probabilistic choices prior to resolving non-determinism. A semantic model – that contains both non-determinism and discrete probabilities – preferably obeys the same order. Bundle probabilistic transition systems (BPTS) [7] is one of the rare models for which

this is indeed the case. Although this model would be appropriate as semantical model, we prefer to use the (more standard) model of Markov decision processes (MDP) [19]. This slightly complicates the semantics, but has the nice property that it facilitates model checking of probabilistic properties.

We now embed the step semantics described above in a global semantics, mapping a collection of P-statecharts onto a finite Markov decision process. To allow the interpretation of temporal logic formulas later on, we equip an MDP with a state-labelling that assigns a set of atomic propositions to states. We assume a given fixed set AP of atomic propositions.

Markov decision processes. An MDP is a quadruple $(S, Distr, L, s_0)$ where:

- S is a finite, non-empty set of states.
- $Distr$ assigns to each state a finite, non-empty set of distributions on S .
- $L : S \rightarrow \mathbb{P}(AP)$ assigns to each state a set of atomic propositions.
- $s_0 \in S$ is the initial state.

In state s , the atomic propositions in $L(s)$ hold. Informally speaking, an MDP exhibits the following behaviour. Whenever the system is in state s , a probability distribution $\mu \in Distr(s)$ is chosen nondeterministically. Then, the system chooses probabilistically the next state according to the selected distribution μ .

Paths in an MDP. A path is an infinite sequence of states (s_0, s_1, \dots) such that s_0 is the initial state and the probability that s_{i+1} is reached from s_i is > 0 , for each i . A path represents a possible behaviour of an MDP.

MDP semantics of a collection of P-statecharts. In an MDP, first a non-deterministic choice is made (among the available distributions) after which a next state is selected probabilistically. This order is reversed for the construction of a step of a P-statechart. To overcome this difference, we add auxiliary states to the MDP. Recall that (original) states consist of, per P-statechart, a configuration, a set of events, and a valuation, written (C, I, V) . Auxiliary states will correspond to the outcome of Task 2 of the step construction algorithm and consist of, per P-statechart, a configuration, a set of enabled edges and a valuation, written (C, E, V) . Each auxiliary state will be labelled with the distinguished atomic proposition Δ . It offers a non-deterministic choice of trivial distributions (assigning probability 1 to a single state) only, such that each successor state is an original state (not labelled Δ). Original states, in turn, do only possess singleton sets of probability distributions (hence there is no non-determinism), and all states with positive probability will be states labelled Δ . This type of MDPs is also known as (strictly) alternating probabilistic transition systems [10].

Let the set $AP = \{\Delta\} \cup \bigcup_{i=1}^n \{i.isin(x) \mid x \in Nodes_i\}$. For a finite collection of P-statecharts (indexed from 1 to n), the MDP $(S, Distr, L, s_0)$ is defined by:

- $S = O \cup A$ where

$$O = \prod_{i=1}^n (Conf_i \times \mathbb{P}(Events) \times Val_i) \quad \text{and} \quad A = \prod_{i=1}^n (Conf_i \times \mathbb{P}(Edges_i) \times Val_i).$$

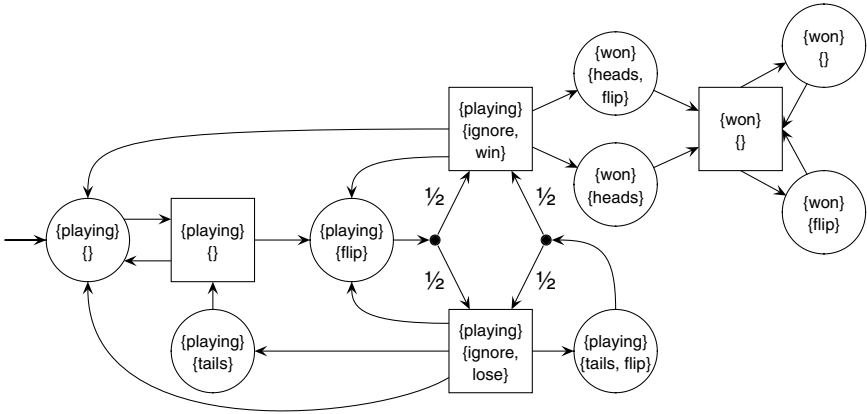


Fig. 4. MDP semantics of the P-statechart of Fig. 1

- $Distr((s_1, \dots, s_n)) =$
 - $\{(s'_1, \dots, s'_n) \mapsto \prod_{i=1}^n \mu_i(s'_i)\}$, if $(s_1, \dots, s_n) \in O$, where μ_i is the distribution corresponding to $\mathcal{PR}_{s_i} = (\mathbb{P}(Edges_i), P)$, defined by $\mu_i(C_i, E', V_i) = P(\{E'\})$ for $E' \subseteq Edges_i$ and $\mu_i(s) = 0$ for other states³.
 - $\{\mu_s^1 \mid \exists T_i \in Steps(E_i) : s = Execute(C_{1\dots n}, T_{1\dots n}, V_{1\dots n})\}$ otherwise. μ_s^1 denotes the trivial distribution that assigns probability 1 to state s .
- $L((s_1, \dots, s_n)) = \begin{cases} \bigcup_{i=1}^n \{i.isin(x) \mid x \in C_i\} & \text{if } (s_1, \dots, s_n) \in O \\ \{\Delta\} & \text{otherwise} \end{cases}$
- $s_0 = ((C'_{0,1}, \emptyset, V_{0,1}), \dots, (C'_{0,n}, \emptyset, V_{0,n})) \in O$, where $C'_{0,i}$ is the default completion of $\{root_i\}$ and $V_{0,i}$ is the initial valuation in the i th P-statechart.

where $s_i = (C_i, I_i, V_i)$ if $(s_1, \dots, s_n) \in O$ and $s_i = (C_i, E_i, V_i)$ otherwise.

Example 2. To illustrate how P-statecharts are mapped onto MDPs, we consider the “unreliable, but fair coin” P-statechart from Fig. 1. We compose this P-statechart with an event generator, which generates “flip” events at random. (If we don’t add a component which generates “flip” events, the only reachable state of the system would be the initial state, due to the closed world assumption.) The MDP semantics of this collection of two P-statecharts is illustrated in Fig. 4. Here, original states of the form (C, I, \emptyset) (where I is a set of events) are shown by circles with the sets C and I inscribed. Auxiliary states of the form (C, E, \emptyset) (where E is a set of edges) are shown by boxes with the sets C and E inscribed. The names used for edges are shown in italics in Fig. 1; the nodes and edges of the event generator are omitted.

³ If no edges are enabled in P-statechart i , this will lead to $\mu_i(C_i, \emptyset, V_i) = 1$.

4 Comparison to Non-probabilistic Semantics

In this section, we compare our P-statechart semantics to the semantics of the corresponding non-probabilistic statechart. For the sake of simplicity, we adapt the semantics of [8] to our notation and abstract from some minor aspects of their semantics (the very same aspects mentioned in Sect. 1).

A (traditional) statechart is a tuple $(Nodes, Events, Vars, Edges)$ where $Edges \subseteq \mathbb{P}(Nodes) \times (Events \cup \{\perp\}) \times Guards \times \mathbb{P}(Actions) \times \mathbb{P}(Nodes)$ is a set of edges, and the other components are as for a P-statechart. A step is – like before – an enabled, consistent, prioritized and maximal set of edges. The execution of steps in a finite collection of statecharts leads to a new state of the statecharts similar to the procedure described in Sect. 3.2. The semantics of a collection of statecharts is a Kripke structure (instead of an MDP). A Kripke structure KS is a quadruple (S, T, L, s_0) , where S, L and s_0 are defined as for MDPs and $T \subseteq S \times S$ is the (non-probabilistic) transition relation.

Projections. To facilitate the comparison, we define three projections: one that abstracts from the probabilistic choices in a P-statechart (called α_1), and one that abstracts from the probabilistic choices in an MDP (called α_2). The projections replace probabilities by nondeterminism. Let $\alpha_1(PSC_i) = SC_i$ where SC_i is obtained from PSC_i by replacing the set of P-edges $PEdges_i$ by the set of edges $Edges'_i$ and by adding some variables to handle the interplay between probabilities and priorities correctly: For every nontrivial P-edge where edges have different priorities, we add one variable to make a nondeterministic choice between the possible continuations of the P-edge. So, $Edges'_i$ has the form $\{(X, e, g \wedge \dots, A \cup \{\dots\}, Y) \mid \exists d : \exists P : \iota(d) = (X, e, g, P) \wedge (d, A, Y) \in Edges_i\}$, where the \dots stand for checks of and assignments to new variables, where necessary. We also add an extra state *Init* to initialize the variables. Later on, we will use a third projection π to remove the additional variables and *Init* again. Further, $\alpha_1(\{PSC_1, \dots, PSC_n\}) = \{\alpha_1(PSC_1), \dots, \alpha_1(PSC_n)\}$.

Example 3. Figure 5 illustrates α_1 : it contains a possible translation of Fig. 2 to a statechart. As both P-edges in the P-statechart have edges of different priority, α_1 adds two variables.

The projection α_2 is defined by: $\alpha_2(S, Distr, L, s_0) = (S', T, L \upharpoonright_{S'}, s_0)$, where:

- $S' = \{((C_1, I_1, V_1), \dots, (C_n, I_n, V_n)) \in S \mid I_i \subseteq Events\}$
- $T = \{(s, s') \mid \exists \mu \in Distr(s) : (\exists \tilde{s} \in S : \mu(\tilde{s}) > 0 \wedge \exists \mu' \in Distr(\tilde{s}) : \mu'(s') > 0)\}$

State set S' contains all non-auxiliary states, and $(s, s') \in T$ whenever s can move to s' via some auxiliary state \tilde{s} in the original MDP with positive probability.

Theorem 1. *The following diagram commutes:*

$$\begin{array}{ccc}
 \{PSC_1, \dots, PSC_n\} & \xrightarrow{\alpha_1} & \{SC_1, \dots, SC_n\} \\
 \text{\scriptsize } sem_1 \downarrow & & \text{\scriptsize } \pi \circ sem_2 \downarrow \\
 MDP & \xrightarrow{\alpha_2} & KS
 \end{array}$$

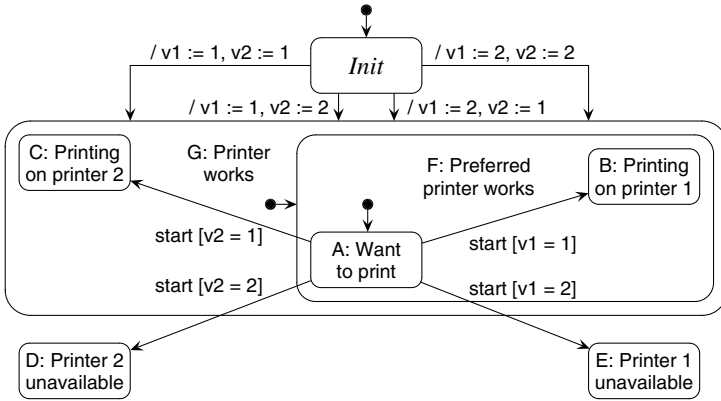


Fig. 5. A translation of Fig. 2 to a statechart

where sem_1 denotes our MDP-semantics and sem_2 denotes the KS-semantics of [8].

Proof. The proof is straightforward (but tedious) by checking the components of the resulting Kripke structures. The central idea is sketched below: The sets of steps in both semantics correspond to each other.

Assume given a state $s_i = (C_i, I_i, V_i)$ of one P-statechart (according to sem_1) and a step T_i with positive probability. We have to prove that there is a corresponding step T'_i in a corresponding state s'_i of $sem_2(\alpha_1(S))$.

The state s'_i is (C_i, I_i, V'_i) ⁴, where V'_i extends V_i with suitable values for the new variables, such that the guards added to the edges in T_i by α_1 become true. The edges in T'_i are extensions of the edges in T_i by assignments to the new variables. We have some freedom of choice in these assignments, as T_i does not prescribe anything. (This is the point where nondeterminism replaces probabilism.) If one considers a path, one can choose the target that enables the next step in the path. The set T'_i is a step: it contains only edges enabled in s'_i and it is consistent, prioritized and maximal because T_i is a step.

The other parts of the proof are easy.

Corollary 1. *The P-statechart semantics is a conservative extension of the statechart semantics of [8].*

5 Property Specification for P-Statecharts

As a property specification language for P-statecharts we propose to use the probabilistic branching time logic PCTL, which extends CTL with probabilistic features. PCTL was originally interpreted over fully probabilistic systems, i. e.,

⁴ States are not subjected to probabilities, that's why s_i and s'_i are so similar.

systems that do not exhibit any non-determinism [11]. We use the interpretation of PCTL over MDPs defined by Baier and Kwiatkowska [1,15]⁵, similar to pCTL and pCTL* [3]. PCTL allows one to express properties such as (Ψ) “the probability that a system crashes within 13 steps without ever visiting certain states is at most 10^{-5} ”. In order to decide these properties, the non-determinism is resolved by means of schedulers (also known as adversaries or policies). Temporal formulas are then interpreted with respect to all schedulers or some schedulers. Here, we restrict ourselves to the fragment of PCTL for which actual model-checking tool-support is available; i. e., we only consider path properties interpreted for all fair schedulers. Formulas of the form “There is a fair scheduler such that ...” can be checked via duality. The model-checking algorithm thus returns “true” for property (Ψ) , iff (Ψ) holds for all fair schedulers that resolve the non-determinism. For simplicity, we do not consider next-formulas.

Syntax and informal semantics. The syntax of PCTL is given by the following grammar, where a denotes an atomic proposition, $p \in [0, 1]$ denotes a probability and \sqsubseteq is a placeholder for a comparison operator $<, \leq, =, \geq, >$:

$$\varphi, \psi ::= \mathbf{true} \mid \mathbf{false} \mid a \mid v \leq k \mid v \geq k \mid \varphi \wedge \psi \mid \neg\varphi \mid \mathcal{P}_{\sqsubseteq p}[\varphi \mathcal{U}^{\leq k} \psi] \mid \mathcal{P}_{\sqsubseteq p}[\varphi \mathcal{U} \psi]$$

The meaning of **true**, comparisons, conjunction and negation is standard. Recall from Sect. 3.3 that atomic propositions are Δ and $i.isin(x)$, which holds in states where P-statechart i is in node x . Formula $\mathcal{P}_{\sqsubseteq p}[\varphi \mathcal{U}^{\leq k} \psi]$ holds in a state if the probability of the set of paths that reach a ψ -state in at most k steps while passing only through φ -states is $\sqsubseteq p$. Property Ψ , e. g., is expressed as $\mathcal{P}_{\leq 10^{-5}}[\neg\varphi \mathcal{U}^{\leq 13} \mathit{crash}]$ where φ describes the states that should be avoided. $\mathcal{P}_{\sqsubseteq p}[\varphi \mathcal{U} \psi]$ has the same meaning, but does not put a bound on the number of steps needed to reach the ψ -state. The temporal operator \diamond can be defined e. g., as $\mathcal{P}_{\sqsubseteq p}[\diamond^{\leq k} \varphi] = \mathcal{P}_{\sqsubseteq p}[\mathbf{true} \mathcal{U}^{\leq k} \varphi]$. (A formal interpretation on MDPs is omitted here, and can be found in [1]).

Schedulers and fair schedulers. The above explanation is ambiguous if non-determinism is present, because the probability will (in general) depend on the resolution of non-determinism. Non-determinism is resolved by schedulers. A *scheduler* selects, for each initial fragment of a path through the MDP, one of the possible (non-deterministic) continuations. It does not resolve probabilistic choices. Several types of schedulers do exist, see [1]. Here, we consider *fair* schedulers. A fair scheduler only selects fair paths. A path π is *fair* if, for each state s that appears infinitely often in π , each of the possible non-deterministic continuations in s also appears infinitely often. Thus, for instance, $\mathcal{P}_{\leq 10^{-5}}[\neg\varphi \mathcal{U}^{\leq 13} \mathit{crash}]$ is valid if for all fair schedulers the probability to reach a crash-state within 13 steps (without visiting a φ -state) is at most 10^{-5} . From now on, we assume all PCTL-formulas to be interpreted over all fair schedulers.

Example 4. For the P-statechart in Fig. 1, we express the following properties:

⁵ Baier and Kwiatkowska sometimes call the logic PBTL.

- “The probability that eventually the game will be over is 1”:

$$\mathcal{P}_{=1}[\diamond \neg 1.isin(playing)]$$

- “In less than 50 % of the cases, the game will be won within at most 20 steps”:

$$\mathcal{P}_{<0.5}[\diamond^{\leq 20} 1.isin(won)]$$

PCTL interpreted over P-statecharts. In our setting, a formula is interpreted over a finite collection of P-statecharts $\{PSC_1, \dots, PSC_n\}$ and one of its states (s_1, \dots, s_n) where $s_i = (C_i, I_i, V_i)$. Formally, the semantics is defined via the MDP semantics, i. e., $\{PSC_1, \dots, PSC_n\} \models \varphi$ iff the corresponding MDP satisfies $\overline{\varphi}$. Here $\overline{}$ denotes a syntactic translation needed to “hop along” the auxiliary (Δ -labelled) MDP states. It is defined by induction over the structure of formulas. For elementary PCTL-formulas, such as atomic propositions and variable constraints, this translation is simply the identity, e. g., $\overline{\mathbf{true}} = \mathbf{true}$. For the remaining operators we have:

$$\begin{aligned} \overline{\varphi \wedge \psi} &= \overline{\varphi} \wedge \overline{\psi} \\ \overline{\neg \varphi} &= \overline{\varphi} \\ \overline{\mathcal{P}_{\exists p}[\varphi \mathcal{U}^{\leq k} \psi]} &= \mathcal{P}_{\exists p}[(\overline{\varphi} \vee \Delta) \mathcal{U}^{\leq 2k} \overline{\psi}] \\ \overline{\mathcal{P}_{\exists p}[\varphi \mathcal{U} \psi]} &= \mathcal{P}_{\exists p}[(\overline{\varphi} \vee \Delta) \mathcal{U} \overline{\psi}] \end{aligned}$$

6 Example: Hawks and Doves

This section applies P-statecharts and PCTL to the specification and verification of the behaviour of a small example taken from theoretical biology. Conflicts between animals are often analysed using simulation techniques. We consider the following variant of the hawk–dove-game [4,17]. In a population of animals, individuals combat for some advantage (such as food, dominance, or mates), their success being measured in points. Individuals may fight using several strategies. In particular, we consider

Hawk strategy: Hawk-like individuals will fight with great effort, until they win the contest (+5 points) or are severely injured (−3 points).

Dove strategy: Dove-like individuals will fight with limited effort, until they win the contest (+5 points) or give up after some fight (−1 point). When facing a hawk, they immediately give up (± 0 points).

We consider a small scenario with three individuals and an arbiter. In every round, the arbiter chooses nondeterministically a pair of individuals; they will be opponents in the next contest. The two individuals select probabilistically the hawk or dove strategy. The arbiter decides who wins. Figures 6 and 7 show the P-statechart for one individual and the arbiter, respectively. The players all start off with 17 points and the individual scores may float in the interval $[0, 55]$ (otherwise they stop). Applying the MDP semantics of Sect. 3 together with




Fig. 6. Statechart of a contestant in the hawk–dove-game

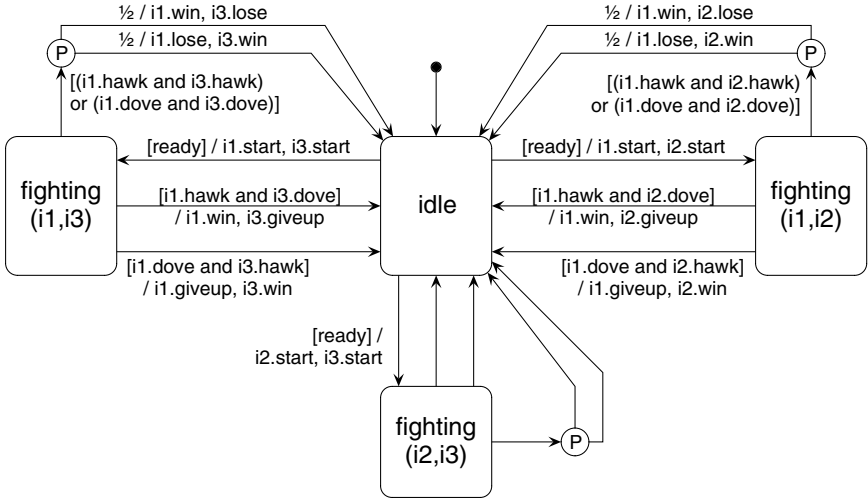


Fig. 7. Statechart of the arbiter in the hawk–dove-game. We have omitted some P-edge labels from and to node `fighting(i2,i3)`, which are analogous to the other `fighting` nodes.

some further optimisations (leaving out trivial intermediary states, encoding the configuration efficiently) leads to a system of 3,147,947 reachable states. The size of the state space is mainly dominated by the integer variables storing the scores. Different scenarios were checked with the model checker PRISM [15] where each scenario consisted of different types of animals. These types were generated by taking different values for p , the probability to behave like a dove. Formulas are checked for the initial state. The three considered scenarios are the following.

One daring and two careful players. This is a scenario with two individuals (c_1 and c_2) for which $p = 0.75$ and one individual (d) with $p = 0.5$. The probability that any individual dies (its points drop below 0) turns out to be very small, with the daring individual running a higher risk of being killed, since

$$\mathcal{P}_{\leq 10^{-7}}[(\neg dead(c_1) \wedge \neg dead(d)) \mathcal{U} dead(c_2)] \text{ is refuted}$$

(and likewise with c_1 and c_2 reversed), but

$$\mathcal{P}_{\leq 10^{-7}}[(\neg dead(c_1) \wedge \neg dead(c_2)) \mathcal{U} dead(d)] \text{ holds.}$$

The actual probability of d dying first is at most (depending on the scheduler) $7.206 \cdot 10^{-7}$, while the probability of the careful one dying first is at most $5.923 \cdot 10^{-8}$ (each). On the other hand, the daring individual is likely to outperform the others on accumulating a certain number of points, say 37. This follows from verifying:

$$\mathcal{P}_{<0.5}[(points_{c_1} < 37 \wedge points_d < 37) \mathcal{U} points_{c_2} \geq 37] \text{ which is valid, and}$$

$$\mathcal{P}_{\leq 0.75}[(points_{c_1} < 37 \wedge points_{c_2} < 37) \mathcal{U} points_d \geq 37] \text{ which is refuted.}$$

Three aggressive players. In this scenario each animal (d_1, d_2, d_3) plays hawk with probability 0.9 (i. e., $p = 0.1$). The probability that some of the individuals dies is relatively high, e. g.,

$$\mathcal{P}_{<0.01}[(\neg dead(d_1) \wedge \neg dead(d_2)) \mathcal{U} dead(d_3)] \text{ is refuted}$$

(and likewise for the permutations of the d_i). So, there are schedulers which will lead to d_3 dying first with more than 1% chance. The probability that one of the individuals gets more than 37 points within 100 steps is always less than 0.75, as

$$\mathcal{P}_{<0.75}[\diamond^{\leq 100} (points_{d_1} \geq 37)] \text{ holds.}$$

Three careful players. In the opposite situation (the three individuals play dove with probability 0.9), the individuals (c_1, c_2, c_3) are less likely to die and more likely to get a reward fast. The probability that any of the individuals dies is rather low as, e. g.,

$$\mathcal{P}_{\leq 10^{-10}}[(\neg dead(c_1) \wedge \neg dead(c_2)) \mathcal{U} dead(c_3)] \text{ holds.}$$

So, for any scheduler, the probability of c_3 dying first never exceeds 10^{-10} . The probability that one of the individuals gets more than 37 points within 100 steps turns out to be greater than 0.8, since

$$\mathcal{P}_{\leq 0.8}[\diamond^{\leq 100} (points_{c_1} \geq 37)] \text{ is refuted.}$$

Conclusion. As a general conclusion of the experiments we may state that it is good for a population as a whole if the animals are careful; but an individual may be at an advantage if it is more daring than the others.

7 Discussion and Conclusion

This section discusses our approach in the broader context of statechart semantics and of probabilistic models.

Contribution. This paper has developed a recipe to conservatively extend a statechart dialect with probabilistic features. We have applied this recipe to the requirement-level UML semantics of [8], mapping the probabilistic extension onto Markov decision processes as semantic models. Further, we have shown how to use the probabilistic logic PCTL to specify properties over probabilistic statecharts and how model checking of probabilistic statecharts can be performed effectively.

Adaptation to other statechart semantics. Various semantics have been published for statecharts, [2] lays out the spectrum of the many possibilities in defining a semantics. The extension to P-statecharts described in this paper can be applied to a wide range of other semantic definitions. The main idea of our extension is:

1. Syntactically, probabilities are trigger-guarded, i. e., reactions to triggers may depend on the result of a probabilistic experiment, whereas the triggers themselves are not subjected to probabilities. This restricts our approach to describing system randomness, opposed to environmental randomness.
2. Semantically, we reduce the P-statechart probabilistically to a (traditional) statechart, and this is done just before a step. The step is constructed and executed in the traditional statechart setting, and the step's result is interpreted in the P-statechart again. Such a reduction is possible as long as the effects of probabilistic experiments are encapsulated in the steps.

In principle it is possible to define a (traditional) statechart semantics which – if interpreted in the probabilistic extension – would break the encapsulation of probabilities within a step. For instance, one could imagine a semantics where a state variable depends on the enabledness of specific outgoing edges (which could only be decided after resolving the probabilism). However, such a feature appears to be rarely used, the overview given in [2] does not mention any feature like this.

Possible simplifications. Sect. 3.3 has mentioned BPTS as the most natural model for a P-statechart semantics. Thus, we could have simplified the semantics if there were a BPTS model checker available.

On the other hand, observe that the examples in Sect. 3.1 depend on the fact that some P-edge allows a probabilistic choice between target nodes with different parents. For most priority schemes, and for consistency, not the actually entered nodes are relevant, but their parents. If we disallow probabilistic choices between target nodes with different parents, we could resolve nondeterminism and probabilism in a different order and simplify several points: Theorem 1 can be formulated and proved simpler. It becomes feasible to give a direct semantics in terms of MDPs (i. e., without intermediary states), which is closer to the intuition behind P-statecharts. However, it is no more possible to express behaviours like the examples in Sect. 3.1.

Lessons learnt. It is easy to formulate an intuitive extension of statecharts with probabilities. However, when we started formalising and detailing it, a delicate balance had to be found with the other features of statecharts. Our first version of Theorem 1, for example, didn't work properly in the case that some edge has a higher priority than another edge which belongs to the same P-edge.

We have tried to formulate the extension as powerful as possible. This also revealed the problems of extending statecharts more clearly. For some applications, a simpler extension, or a simpler variant of basic statecharts is enough. In that case, one should ask whether the result is worth the effort.

Future work. Apart from exploring the specification and verification approach to system randomness on larger case studies, we are intending to investigate the very same approach in the context of environmental randomness. This asks for modelling and verification support for probabilistic and timing aspects of external stimuli a reactive system is exposed to.

Acknowledgements

The authors thank Rik Eshuis for pointing out a flaw in an earlier version of Theorem 1. Roel Wieringa is thanked for his comments on a draft version of the paper.

References

1. Christel Baier and Marta Kwiatkowska. Model checking for a probabilistic branching time logic with fairness. *Distributed Computing*, 11(3):125–155, 1998.
2. Michael von der Beeck. A comparison of statecharts variants. In H. Langmaack, W.-P. de Roever, and J. Vytupil, editors, *Formal Techniques in Real-Time and Fault-Tolerant Systems : . . . proceedings*, pages 128–148, Berlin, 1994. Springer.
3. Andrea Bianco and Luca de Alfaro. Model checking of probabilistic and nondeterministic systems. In P. S. Thiagarajan, editor, *Foundations of Software Technology and Theoretical Computer Science : . . . proceedings*, volume 1026 of *LNCS*, pages 499–513, Berlin, 1995. Springer.
4. Philip H. Crowley. Hawks, doves, and mixed-symmetry games. *Journal of Theoretical Biology*, 204(4):543–563, June 2000.
5. M. Dal Cin, G. Huszerl, and K. Kosmidis. Evaluation of dependability critical systems based on guarded statechart models. In *4th Int. Symp. on High Assurance Systems Engineering (HASE)*, pages 37–45. IEEE, 1999.
6. Werner Damm, Bernhard Josko, Hardi Hungar, and Amir Pnueli. A compositional real-time semantics of statemate designs. In Willem-Paul de Roever, Hans Langmaack, and Amir Pnueli, editors, *Compositionality : the significant difference. COMPOS '97*, volume 1536 of *LNCS*, pages 186–238, Berlin, 1998. Springer.
7. Pedro R. D'Argenio, Holger Hermanns, and Joost-Pieter Katoen. On generative parallel composition. In Christel Baier, Michael Huth, Marta Kwiatkowska, and Mark Ryan, editors, *PROBMIV'98 First International Workshop on Probabilistic Methods in Verification : Indianapolis, Ind. . . . 1998*, volume 22 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 2000. URL: <http://www.elsevier.nl/locate/entcs/volume22.html>.
8. Rik Eshuis and Roel Wieringa. Requirements-level semantics for UML statecharts. In Scott F. Smith and Carolyn L. Talcott, editors, *Formal Methods for Open Object-Based Distributed Systems IV : . . . FMOODS*, pages 121–140, Boston, 2000. Kluwer Academic Publishers.
9. Stefania Gnesi, Diego Latella, and Mieke Massink. A stochastic extension of a behavioural subset of UML statechart diagrams. In L. Palagi and R. Bilof, editors, *Fifth International Symposium on High-Assurance Systems Engineering (HASE)*, pages 55–64. IEEE Computer Society Press, 2000.
10. H. A. Hansson. *Time and Probability in Formal Design of Distributed Systems*. PhD thesis, University of Uppsala, 1991.

11. Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6:512–535, 1994.
12. David Harel and Eran Gery. Executable object modeling with statecharts. *Computer*, 30(7):31–42, July 1997. IEEE.
13. David Harel and Amnon Naamad. The STATEMATE semantics of statecharts. *ACM Transactions on Software Engineering and Methodology*, 5(4):293–333, 1996.
14. Peter King and Rob Pooley. Derivation of Petri net performance models from UML specifications of communications software. In Boudewijn R. Haverkort, Henrik C. Bohnenkamp, and Connie U. Smith, editors, *Computer Performance Evaluation : Modelling Techniques and Tools ; ... TOOLS 2000*, volume 1786 of *LNCS*, pages 262–276, Berlin, 2000. Springer.
15. Marta Kwiatkowska, Gethin Norman, and David Parker. Probabilistic symbolic model checking with prism: A hybrid approach. In Joost-Pieter Katoen and Perdita Stevens, editors, *Tools and Algorithms for the Construction and Analysis of Algorithms : ... TACAS*, volume 2280 of *LNCS*, pages 52–66, Berlin, 2002. Springer.
16. C. Lindemann, A. Thümmler, A. Klemm, M. Lohmann, and O. P. Waldhorst. Quantitative system evaluation with DSPNexpress 2000. In *Workshop on Software and Performance (WOSP)*, pages 12–17. ACM, 2000.
17. J. Maynard Smith and G. R. Price. The logic of animal conflict. *Nature*, 246, November 1973.
18. S. McMenamin and J. Palmer. *Essential Systems Analysis*. Yourdon Press, New York, 1984.
19. Martin L. Puterman. *Markov Decision Processes : Discrete Stochastic Dynamic Programming*. Wiley, New York, 1994.
20. R. Segala and N. Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2:250–273, 1995.
21. A. N. Shiryaev. *Probability*, volume 95 of *Graduate texts in mathematics*. Springer, New York, 1996.
22. UML Revision Task Force. *OMG UML Specification 1.3*. Object Management Group, 1999. <http://www.omg.org/cgi-bin/doc?ad/99-06-08>.