Comparative Analysis of Bisimulation Relations on Alternating and Non-Alternating Probabilistic Models

Roberto Segala, Andrea Turrini
Dipartimento di Informatica - Università di Verona

Abstract

We consider bisimulation and weak bisimulation relations in the context of the Labeled Markov Chains of Hansson and Jonsson, the Concurrent Labeled Markov Chains of Philippou, Lee, and Sokolsky, and the Probabilistic Automata of Segala. We identify a taxonomy of bisimulation relations that captures the existing definitions for each one of the three models, and we compare the relations within each model and across models. The comparison across models is given according to a notion of embedding, where we order the three models by generality and we view objects in less general models as objects of more general models.

1 Introduction

The literature on concurrency theory shows an increasing number of results aimed at extending known formalisms with stochastic behavior, e.g., labeled transition systems, automata, process algebras, Petri nets, event structures, domains. In this paper we are interested in extensions of labeled transition systems with discrete stochastic behavior, the reason being that there are several proposals for extensions and yet not a fully clear understanding of the relationships between the models.

The models that we consider are the Labeled Markov Chains of Hansson and Jonsson [6], the Concurrent Labeled Markov Chains of Philippou, Lee, and Sokolsky [8], and the simple Probabilistic Automata of Segala [10]. The Labeled Markov Chains of Hansson and Jonsson and the Concurrent Labeled Markov Chains of Philippou, Lee, and Sokolsky derive from a proposal of Vardi [13], where nondeterminism is added to discrete Markov chains by introducing a new class of *nondeterministic* states, where nondeterministic choices occur. In other words, the states of an automaton are partitioned into *probabilistic* and *nondeterministic* states: a probabilistic state describes a probability measure over nondeterministic states, while a nondeterministic state describes nondeterminism by enabling several ordinary transitions. The difference between Labeled Markov Chains and Concurrent Labeled Markov Chains is that in the first model there is a strict alternation between nondeterministic and probabilistic states (i.e., transitions leaving from nondeterministic states lead to probabilistic states only), while in the second model, as well as in the original model of Vardi, the alternation is not enforced between nondeterministic states. For this reason the first model is often referred to as the model of *strictly alternating automata* and the second model is often referred to as the model of *alternating automata*. The simple Probabilistic Automata of Segala do not distinguish states, but introduce a notion of *probabilistic transition* whose target is a discrete measure over states rather than a single state. A deterministic version of probabilistic automata was introduced by Rabin [9] in the context of language theory. Given the absence of a distinction between nondeterministic and probabilistic states, and in contraposition to the alternating models, the model of simple Probabilistic Automata is often referred to as the model of *non-alternating automata*.

Several equivalence and preorder relations are studied within the models above. In this paper we are interested in bisimulation relations in their strong and weak version. Indeed, there are several definitions of such relations, given in different styles, and it is not clear how such definitions are related within and across models. The relations across models are important to understand the expressive power of alternating and non-alternating automata.

The only relation that was proposed for strictly alternating automata is *strong bisimulation* [6], which is an equivalence relation on nondeterministic states, that induces an equivalence relation on probabilistic states, such that, whenever two states are related, each transition from one state can be matched by a transition from the other state. Two matching transitions are labeled by the same action and reach the same equivalence classes with the same probability. This idea was first proposed by Larsen and Skou [7].

The relations proposed for alternating automata are strong and weak bisimulation [8]. Again strong bisimulation is an equivalence relation on nondeterministic states; however, a special treatment of probabilistic states ensures that an ordinary transition between nondeterministic states can be matched by a transition to a probabilistic state that describes a Dirac measure. Somehow, probabilistic states are considered as a technicality to represent the probabilistic transitions of Segala. Weak bisimulations, instead, are defined on all states; however, the measure described by a probabilistic state is studied conditional on leaving the equivalence class of the probabilistic state. The question then is whether such a discrepancy between strong and weak bisimulation is fundamental.

Historically, weak bisimulation was introduced on Probabilistic Automata by Segala and Lynch [11] and then studied by Baier and Hermanns on fully probabilistic models [1]. The definition of [11] does not use conditional measures, while the definition of [1] uses them. Evidently the definition of [8] was influenced by [11] and [1], but what is the reason for using conditional measures in [1] and [8]?

The relations proposed for non-alternating automata are again strong and weak bisimulation, where the matching between transitions and weak transitions is the classical one according to [7]; however in [11] it is proposed that a transition from a state may be matched by a convex combination of transitions from another state, thus leading to weaker notions of *probabilistic bisimulation* and *weak probabilistic bisimulation*. In this paper we show that the weak bisimulation of [8] coincides with the probabilistic weak bisimulation of [11], thus understanding why conditional measures are not needed in the non-alternating model.

Our comparative study is carried out according to two scenarios that we call *embedding* and *transformation*. According to the embedding scenario, we view strictly alternating automata as special cases of alternating automata and alternating automata as special cases of non-alternating automata. Indeed, transitions from nondeterministic states are just ordinary transitions, which are a special case of probabilistic transitions, while the probability measures associated with probabilistic states can be represented by probabilistic transitions labeled by an ad-hoc internal action, which in this paper we call ν. Thus, we define appropriate functions to embed restricted models to more general models, obtaining a way to propagate relations from general models to restricted models.

The transformation scenarios formalizes the idea that is behind some folklore constructions to represent elements of a formalism within other formalisms. In particular, probabilistic states can be seen as a technicality to represent probabilistic transitions. Thus, a strictly alternating automaton can be represented by (transformed into) a non-alternating automaton by removing all probabilistic states and collapsing into a unique probabilistic transition all pairs of consecutive transitions that pass through a probabilistic state. Vice versa, the converse transformation introduces new probabilistic states and splits probabilistic transitions. The transformations that involve alternating automata ensure that no splitting occurs on transitions that lead to Dirac measures. Again, via transformations it is possible to propagate definitions from one model to the other.

For our study we take the non-alternating model as our reference model, and we start by understanding the main ideas behind the definitions proposed in the literature. We classify relations on the alternating models according to how nondeterministic and probabilistic states are treated. In the definition of strong bisimulation for strictly alternating automata, nondeterministic states and probabilistic states are treated separately, and the step condition, i.e., the condition on matching transitions, is standard. This suggests a typology of bisimulation, which we call *divided*, where we impose a total separation between probabilistic and nondeterministic states and we check the step condition via embedding on non-alternating automata. In the definition of strong bisimulation for alternating automata, the equivalence relation is studied on nondeterministic states only, while probabilistic states play only a technical role in representing a probabilistic transition. This suggests a typology of bisimulation, which we call *nondeterministic*, where we consider an equivalence relation on nondeterministic states and we check the step condition via transformation on non-alternating automata. Finally, in the definition of weak bisimulation for alternating automata there is no distinction between probabilistic and nondeterministic states, and the step condition, except for self loops, is standard. This suggests a third typology of bisimulation, which we call *mixed*, where we consider an equivalence relation on states and we check the step condition via embedding on non-alternating automata.

Our study works as follows. On each of the alternating models we define the nondeterministic, divided, and mixed typologies of strong bisimulation, weak bisimulation, strong probabilistic bisimulation, and weak probabilistic bisimulation, and we show how the original definitions fit into this classification. Then we compare the different typologies within each model and across models via embedding. In the strong case many relations collapse; however, it is clear that probabilistic relations are useless in the alternating models due to the presence of probabilistic states. In the weak case the divided typology turns out to be too strong and practically useless. The other two typologies are incomparable, mainly for technical reasons that we clarify later.

There are other relevant comparative studies in the literature. In [5] probabilistic models are classified into reactive, generative, stratified, and are ordered according to how a bisimulation in the style of [7] in one model can be abstracted into a bisimulation in another model. There is no real nondeterminism in the models considered by [5]; indeed, the models considered in this paper can be seen as a nondeterministic extension of reactive systems. In [3] the mixed relations are compared across models via transformation by studying axiomatizations for a process algebra of finite trees. In [2] there is a comparative study based on bisimulations on discrete and continuous Markov chains with no nondeterminism. Finally, in [12] there is an extensive comparison of several models, including the three models considered in this paper. The comparison is based on strong bisimulation only, and formalisms are ordered via a criterion that considers a model more general than another model if the embedding from the least general model to the most general model preserves and reflects bisimulation. The embeddings considered in [12] are more conservative than ours since they do not change the transition relation; thus, alternating and non-alternating models are not comparable. In our case we use a similar ordering criterion; however the study is carried out with several bisimulation relations, including weak and probabilistic bisimulations, and we study embeddings as well as transformations. On the counter part, we consider only three models.

The rest of the paper is organized as follows. Section 2 outlines some background on measure theory; Section 3 gives formal definitions of the models under study; Section 4 formalizes the embedding and transformation functions; Section 5 gives the formal definitions of the bisimulation relations under study; Section 6 introduces the typologies of bisimulations; Section 7 compares relations within and across models; Section 8 gives some concluding remarks.

2 Preliminaries on Measure Theory

A *σ-field* over a set X is a set $\mathcal{F} \subseteq 2^X$ that includes X and is closed under complement and countable union. A *measurable space* is a pair $(X, \mathcal{F})$ where X is a set, also called *sample space*, and $\mathcal{F}$ is a σ-field over X. A measurable space $(X, \mathcal{F})$ is called *discrete* if $\mathcal{F} = 2^X$. A *measure* over a measurable space $(X, \mathcal{F})$ is a function $\mu\colon \mathcal{F} \to \mathbb{R}^{\geq 0}$ such that, for each countable collection $\{X_i\}_{i \in I}$ of pairwise disjoint elements of $\mathcal{F}$, $\mu[\cup_I X_i] = \sum_I \mu[X_i]$. A *probability measure* over a measurable space $(X, \mathcal{F})$ is a measure μ over $(X, \mathcal{F})$ such that $\mu[X] = 1$. A *sub-probability measure* over $(X, \mathcal{F})$ is a measure over $(X, \mathcal{F})$ such that $\mu[X] \leq 1$. A measure over a discrete measurable space $(X, 2^X)$ is called a *discrete measure* over X.

Given a set X, denote by $\mathit{Disc}(X)$ the set of discrete probability measures over X, and by $\mathit{SubDisc}(X)$ the set of discrete sub-probability measures over X. We call a discrete (sub-)probability measure a *Dirac* measure if it assigns measure 1 to exactly one object x (denote this measure by δ_x). We also call Dirac a sub-probability measure that assigns measure 0 to all objects. In the sequel discrete sub-probability measures are used to describe progress. If the measure of a sample space is not 1, then it means that with some non-zero probability the system does not progress.

3 Probabilistic Automata

In this section we give the formal definitions of the three kinds of models under study and we introduce some notational conventions. We assume a set Σ of actions partitioned into *external* and *internal*. For the purpose of this paper we have only two internal actions: an action τ and a special action ν that we use in our embeddings.

3.1 Strictly Alternating Automata

A *strictly alternating automaton* [6] is a tuple $(S, \bar{s}, \Sigma, \mathcal{D})$ where S is a set of *states*, $\bar{s} \in S$ is the *start state*, Σ is a set of *actions* and $\mathcal{D}$ is the *transition relation*. The set of states S is partitioned into two sets N and P of *nondeterministic* and *probabilistic* states, respectively. We require that $\bar{s} \in N$. The set of transitions $\mathcal{D}$ is partitioned into two sets $\mathcal{N} \subseteq N \times \Sigma \times P$ and $\mathcal{P} \in P \to \mathit{Disc}(N)$ of *nondeterministic* and *probabilistic* transitions, respectively.

In strictly alternating automata there is a strict alternation between nondeterministic and probabilistic states. The start state is always nondeterministic. The fact that the transition relation from probabilistic states is a function ensures that for each probabilistic state there is exactly one associated probability measure. In the sequel we denote by ***SA*** the class of strictly alternating automata.

3.2 Alternating Automata

An *alternating automaton* [8] is a tuple $(S, \bar{s}, \Sigma, \mathcal{D})$ where S is a set of *states*, $\bar{s} \in S$ is the *start state*, Σ is a set of *actions* and $\mathcal{D}$ is the *transition relation*. The set of states S is partitioned into two sets N and P of *nondeterministic* and *probabilistic* states, respectively. We require that $\bar{s} \in N$. The set of transitions $\mathcal{D}$ is partitioned into two sets $\mathcal{N} \subseteq N \times \Sigma \times S$ and $\mathcal{P} \in P \to \mathit{Disc}(N)$ of *nondeterministic* and *probabilistic* transitions, respectively.

The formal definition of alternating automata differs from the definition of strictly alternating automata only in the third component of $\mathcal{N}$: in addition to strictly alternating automata an alternating automaton may contain transitions from nondeterministic states that reach nondeterministic states directly, thus skipping the intermediate probabilistic state. Indeed, there is no need to include a probabilistic state if the target measure is Dirac. In the sequel we denote by ***A*** the class of alternating automata.

3.3 Non-Alternating Automata

A *non-alternating automaton* [10] is a tuple $(S, \bar{s}, \Sigma, \mathcal{D})$ where S is a set of *states*, $\bar{s} \in S$ is the *start state*, Σ is a set of *actions* and $\mathcal{D} \subseteq S \times \Sigma \times \mathit{Disc}(S)$ is *transition relation*.

In non-alternating automata there is no distinction between nondeterministic and probabilistic states and there is a unique transition relation whose elements, called probabilistic transitions, describe directly the target measure without any need to move through special states. In the sequel we denote by ***NA*** the class of non-alternating automata.

3.4 Notational Conventions

Throughout the paper we let $\mathcal{A}$ range over automata, r, q, s range over states, a, b, c range over actions, and μ range over discrete measures. We also denote the generic elements of an automaton $\mathcal{A}$ by S, $\bar{s}$, Σ, $\mathcal{D}$, etc., and we propagate primes and indices when necessary. Thus, the elements of a non-alternating automaton $\mathcal{A}'_i$ are S'_i, $\bar{s}'_i$, Σ'_i and $\mathcal{D}'_i$. Also, the set of nondeterministic states of an alternating automaton $\mathcal{A}'_i$ is N'_i.

An element of a transition relation $\mathcal{D}$ is called a *transition*. A transition $tr = (s, a, \mu)$ is said to leave from state s, to be labeled by a, and to lead to μ. We denote it alternatively by $s \xrightarrow{a} \mu$. We also say that s enables action a, that action a is enabled from s, and that (s, a, μ) is enabled from s. We denote by *source*(tr) the state s, by *action*(tr) the label a, and by *target*(tr) the measure μ. We apply a similar convention to transitions of the kind (s, a, s') and (s, μ) of (strictly) alternating automata. Finally, we denote by $\mathcal{D}(s)$ the set of transitions with source state s, i.e., $\mathcal{D}(s) = \{tr \in \mathcal{D} \mid \textit{source}(tr) = s\}$.

4 Embeddings and Transformations

In this section we define the embedding and transformation functions on automata. The embedding functions map objects of restricted models to objects in more general models; the transformation functions represent objects of one model in any other model.

4.1 Embeddings

We start from the point of view that the non-alternating model is the most general, in the sense that the structure implicit in the other models is a sub-structure of that of probabilistic automata. The embedding functions represent how the sub-structure of a model can be mapped into the structure of a more general model.

The embedding of the strictly alternating model onto the alternating model, denoted by $\mathcal{E}_{\boldsymbol{SA}\to\boldsymbol{A}}$, is simply the identity function since the strictly alternating model has just a few more restrictions on its transition relation.

The embedding of the alternating model onto the non-alternating model, denoted by $\mathcal{E}_{\boldsymbol{A}\to\boldsymbol{NA}}$, adds a special action ν to the transitions leaving probabilistic states. Formally, if $\mathcal{A}'$ is the embedding of $\mathcal{A}$, all components of $\mathcal{A}'$ are the same as $\mathcal{A}$ except for Σ' and $\mathcal{D}'$, which are defined as follows: $\Sigma' = \Sigma \cup \{\nu\}$, and $\mathcal{D}' = \{(s, a, \delta_{s'}) \mid (s, a, s') \in \mathcal{N}\} \cup \{(s, \nu, \mu) \mid (s, \mu) \in \mathcal{P}\}$.

Finally, the embedding of the strictly alternating model onto the non alternating model, denoted by $\mathcal{E}_{\boldsymbol{SA}\to\boldsymbol{NA}}$, is the composition of the previous two embeddings, which is essentially the same as $\mathcal{E}_{\boldsymbol{A}\to\boldsymbol{NA}}$.

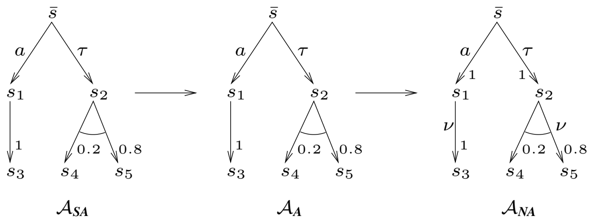

Figure 1. Embeddings of automata.

Figure 1 gives an example of application of the embedding functions. The alternating automaton $\mathcal{A}_{\boldsymbol{A}}$ is the embedding of the strictly alternating automaton $\mathcal{A}_{\boldsymbol{SA}}$, and the non-alternating automaton $\mathcal{A}_{\boldsymbol{NA}}$ is the embedding of both $\mathcal{A}_{\boldsymbol{A}}$ and $\mathcal{A}_{\boldsymbol{SA}}$. Observe that in $\mathcal{A}_{\boldsymbol{NA}}$ action ν is used to label transitions from probabilistic states.

4.2 Transformations

The transformation functions, already used in [3], describe the folklore idea that a transition in the non-alternating model corresponds to two transitions in the strictly alternating model and vice versa. When moving to the alternating model a transition is split only if its target is not a Dirac measure. We omit the formal definitions of the six transformation functions, which we denote by $\mathcal{T}$. Figure 2 gives examples of transformation functions. The

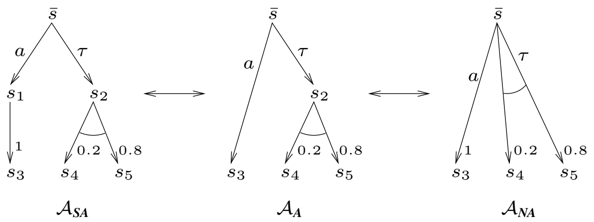

Figure 2. Transformations of automata.

three automata $\mathcal{A}_{\boldsymbol{SA}}$, $\mathcal{A}_{\boldsymbol{A}}$ and $\mathcal{A}_{\boldsymbol{NA}}$ are transformations of each other. Observe, indeed, how from $\mathcal{A}_{\boldsymbol{SA}}$ to $\mathcal{A}_{\boldsymbol{NA}}$ all probabilistic states have disappeared. Similarly, observe how from $\mathcal{A}_{\boldsymbol{SA}}$ to $\mathcal{A}_{\boldsymbol{A}}$ all probabilistic states with associated Dirac measures have disappeared. If we follow the reverse transformation, then the probabilistic states appear again.

5 Bisimulations

In this section we give the formal definitions of the bisimulation relations in the alternating and non-alternating models as they are formulated in the original papers, where we refer the reader for intuitions and justifications. Later we show that the definitions of this section can be expressed in our taxonomy, thus providing the reader with alternative, and sometimes easier to understand, characterizations.

We give a preliminary definition of lifting of an equivalence relation to discrete probability measures, since this concept is common to all the definitions given in this section. Given an equivalence relation $\mathcal{R}$ on a set of states Q, we say that two probability measures μ_1 and μ_2 of $Disc(Q)$ are equivalent according to $\mathcal{R}$, denoted by $\mu_1 \equiv_{\mathcal{R}} \mu_2$, if for each equivalence class $\mathcal{C}$ of $\mathcal{R}$, $\mu_1[\mathcal{C}] = \mu_2[\mathcal{C}]$.

5.1 Strong bisimulation in SA

The definition of strong bisimulation given in [6] relates nondeterministic states and derives an induced equivalence relation on probabilistic states.

Given two strictly alternating automata $\mathcal{A}_1$ and $\mathcal{A}_2$ and an equivalence relation $\mathcal{R}$ on $N_1 \cup N_2$, we say that two probabilistic states p_1 and p_2 are $\mathcal{R}$-equivalent, denoted by $p_1 \equiv_{\mathcal{R}} p_2$, if $\mathcal{P}_1(p_1) \equiv_{\mathcal{R}} \mathcal{P}_2(p_2)$.

Relation $\mathcal{R}$ is a *strong bisimulation* if, for each pair of states $q, r \in N_1 \cup N_2$ such that $q \ \mathcal{R} \ r$, if $q \xrightarrow{a} q'$ for some $q' \in P_1 \cup P_2$ then there exists $r' \in P_1 \cup P_2$ such that $r \xrightarrow{a} r'$ and $q' \equiv_{\mathcal{R}} r'$.

The strictly alternating automata $\mathcal{A}_1$ and $\mathcal{A}_2$ are said to be strongly bisimilar if there exists a strong bisimulation $\mathcal{R}$ on $N_1 \cup N_2$ such that $\bar{s}_1 \ \mathcal{R} \ \bar{s}_2$. We denote this strong bisimulation relation by $\sim_{\boldsymbol{SA}}$.

5.2 Strong bisimulation in A

The definition of strong bisimulation given in [8] requires two auxiliary functions. Given an alternating automaton $\mathcal{A}$, $s, s' \in S$, and $M \subseteq S$, define two functions pr and π as follows.

1. $\mathrm{pr}(s, s') = \begin{cases} \mu[s'] & \text{if } \mathcal{P}(s) = \mu \\ 1 & \text{if } s = s', s \in N \\ 0 & \text{otherwise} \end{cases}$

2. $\pi(s, M) = \sum_{s' \in M} \mathrm{pr}(s, s')$.

Function pr computes the probability of moving from a state s to a state s'. The function, however, states that the probability to move from a nondeterministic state to itself is 1, which, combined with the definition below, implies that a nondeterministic state n can be related to a probabilistic state p provided that the measure associated with p is concentrated on the equivalence class of n. In practice, splitting a transition between two nondeterministic states produces a bisimilar automaton.

Let $\mathcal{A}_1, \mathcal{A}_2$ be two alternating automata. An equivalence relation $\mathcal{R}$ on $S_1 \cup S_2$ is a *strong bisimulation* if, for each pair of states $q, r \in S_1 \cup S_2$ such that $q \ \mathcal{R} \ r$,

1. for all $a \in \Sigma$, if $q, r \in N_1 \cup N_2$ and $q \xrightarrow{a} q'$ then $r \xrightarrow{a} r'$ and $q' \ \mathcal{R} \ r'$;

2. for all $\mathcal{C} \in (S_1 \cup S_2)/\ \mathcal{R}$, $\pi(q, \mathcal{C}) = \pi(r, \mathcal{C})$.

The alternating automata $\mathcal{A}_1$ and $\mathcal{A}_2$ are said to be strongly bisimilar if there exists a strong bisimulation $\mathcal{R}$ on $S_1 \cup S_2$ such that $\bar{s}_1 \ \mathcal{R} \ \bar{s}_2$. We denote this strong bisimulation relation by $\sim_{\boldsymbol{A}}$.

5.3 Weak bisimulation in A

The definition of weak bisimulation of [8] requires some extra machinery to identify weak transitions. We use the same notation as in [8]; however, it is possible to reformulate definitions in terms of the weak transitions that we define for the non-alternating model.

A *computation* of an alternating automaton $\mathcal{A}$ is either a finite or infinite sequence of states, possibly divided by actions, $s_0 a_0 s_1 a_1 s_2 s_3 s_4 \ldots$ such that, for each $i \geq 0$, if $s_i \in N$, then it is followed by an action a_i and a state s_{i+1} such that $(s_i, a_i, s_{i+1}) \in \mathcal{N}$, and if $s_i \in P$, then it is followed by a state s_{i+1} such that $\mathcal{P}(s_i)[s_{i+1}] > 0$.

The *length* of c, denoted by $|c|$, is the number of occurrences of actions in c. If c is infinite, then $|c| = \infty$. We denote by $\mathrm{comps}(\mathcal{A})$ the set of all computations of $\mathcal{A}$ and by $\mathrm{comps}^*(\mathcal{A})$ the set of all finite computations of $\mathcal{A}$. Given $c \in \mathrm{comps}^*(\mathcal{A})$, we define $\mathrm{trace}(c)$ to be the sub-sequence of external actions of c, and $\mathrm{last}(c)$ to be the last state that occurs in c.

A *scheduler* for an alternating automaton $\mathcal{A}$ is a function $\sigma\colon \mathrm{comps}^*(\mathcal{A}) \to \mathcal{N} \cup \{\perp\}$ such that, for each $c \in \mathrm{comps}^*(\mathcal{A})$, if $\sigma(c) \neq \perp$, then $\sigma(c) = (\mathrm{last}(c), a, s)$ for some action a and state s. The meaning of $\sigma(c) = \perp$ is that σ does not schedule anything after c. We denote by $\mathrm{Sched}(\mathcal{A})$ the set of all schedulers for an alternating automaton $\mathcal{A}$.

Given an alternating automaton $\mathcal{A}$ and a scheduler $\sigma \in \mathrm{Sched}(\mathcal{A})$, define the set of *scheduled computations* $\mathrm{Scomps}(\mathcal{A}, \sigma) \subseteq \mathrm{comps}(\mathcal{A})$, to be the set of computations $c = s_0 a_0 s_1 \ldots$ such that $\sigma(c) = \perp$ if c is finite, and for each $i < |c|$, if $s_i \in N$ then $\sigma(s_0 a_0 \ldots s_i) = (s_i, a_i, s_{i+1})$.

Given an alternating automaton $\mathcal{A}$, a set of traces $\Phi \subset \Sigma^*$, a set of states $M \subseteq S$, and a scheduler $\sigma \in \mathrm{Sched}(\mathcal{A})$, define $\mathrm{Paths}(\mathcal{A}, \Phi, M, \sigma)$ to be the set of finite scheduled computations that have trace in Φ and end in a state of M, that is, $\mathrm{Paths}(\mathcal{A}, \Phi, M, \sigma) = \{c \in \mathrm{Scomps}(\mathcal{A}, \sigma) \mid \mathrm{last}(c) \in M, \mathrm{trace}(c) \in \Phi\}$. The probability of $\mathrm{Paths}(\mathcal{A}, \Phi, M, \sigma)$ assuming start state $\bar{s}$, denoted by $\Pr(\bar{s}, \Phi, M, \sigma)$, is given by the smallest solution to $X(\bar{s}, \Phi, M, \sigma, \bar{s})$ defined by the following set of equations:

$$X(s, \Phi, M, \sigma, c) = \begin{cases} 1 & \text{if } \varepsilon \in \Phi,\, s \in M,\, \sigma(c) = \bot \\ 0 & \text{if } \varepsilon \notin \Phi,\, s \notin M,\, \sigma(c) = \bot \\ 0 & \text{if } \Phi = \emptyset \\ \sum_{s'} \mu[s'] X(s', \Phi, M, \sigma, cs') & \\ \quad \text{if } s \in P,\, \mathcal{P}(s) = \mu & \\ X(s', \Phi - a, M, \sigma, cas') & \\ \quad \text{if } s \in N,\, \sigma(c) = (s, a, s') & \end{cases}$$

where ε is the empty string and $\Phi - a = \{\phi \mid a\phi \in \Phi\}$.

Before defining weak bisimulation, we need an extra function that extends the π function used for strong bisimulation. Formally, given an alternating automaton $\mathcal{A}$ and an equivalence relation $\mathcal{R}$ on S, we denote by $\pi_{\mathcal{R}}(s, \mathcal{C})$ the probability of reaching $\mathcal{C} \subseteq S$ from state s conditional on leaving the equivalence class $[s]_{\mathcal{R}}$:

$$\pi_{\mathcal{R}}(s, \mathcal{C}) = \begin{cases} \pi(s, \mathcal{C}) & \text{if } \pi(s, [s]_{\mathcal{R}}) = 1; \\ \pi(s, \mathcal{C})/(1 - \pi(s, [s]_{\mathcal{R}})) & \text{otherwise.} \end{cases}$$

Technically speaking $\pi_{\mathcal{R}}(s, [s]_{\mathcal{R}})$ should be defined to be 0. However, this value is never used, so the original definition in [8] has the form stated above.

Let $\mathcal{A}_1, \mathcal{A}_2$ be two alternating automata. An equivalence relation $\mathcal{R}$ on $S_1 \cup S_2$ is a *weak bisimulation* if, for each pair of states $s, t \in S_1 \cup S_2$ such that $s \;\mathcal{R}\; t$,

1. for all $a \in \Sigma$, if $s \in N_1 \cup N_2$ and $s \xrightarrow{a} s'$, then there exists a scheduler σ such that $\Pr(t, \tau^* a \tau^*, [s']_{\mathcal{R}}, \sigma) = 1$;

2. there exists a scheduler σ such that for all $\mathcal{C} \in (S_1 \cup S_2)/\mathcal{R} - [s]_{\mathcal{R}}$, $\pi_{\mathcal{R}}(s, \mathcal{C}) = \Pr(t, \tau^*, \mathcal{C}, \sigma)$.

The alternating automata $\mathcal{A}_1$ and $\mathcal{A}_2$ are said to be weakly bisimilar if there exists a weak bisimulation $\mathcal{R}$ on $S_1 \cup S_2$ such that $\bar{s}_1 \;\mathcal{R}\; \bar{s}_2$. We denote this weak bisimulation relation by $\approx_{\boldsymbol{A}}$.

5.4 Strong Bisimulations in NA

There are two notions of strong bisimulation for the non-alternating model, depending on whether a transition is matched by a transition or by a convex combination of transitions, called a combined transition. For this reason we define first the notion of combined transition.

Given a non-alternating automaton $\mathcal{A}$, a state s, and a measure $\gamma \in \mathit{SubDisc}(\mathcal{D}(s))$, define the *combined transition* according to γ to be the pair (s, μ_γ), where $\mu_\gamma \in \mathit{SubDisc}(\Sigma \times S)$ is defined for each pair (a, q) as $\mu_\gamma[(a, q)] = \sum_{(s,a,\mu)} \gamma[(s, a, \mu)] \mu[q]$. We denote a combined transition (s, μ) alternatively by $s \longrightarrow_{\mathrm{C}} \mu$. Whenever there exists an action a such that $\mu[(a, S)] = 1$ we denote the corresponding combined transition alternatively by $s \xrightarrow{a}_{\mathrm{C}} \mu'$ where $\mu' \in \mathit{Disc}(S)$ is defined for each state q as $\mu'[q] = \mu[(a, q)]$.

Let $\mathcal{A}_1, \mathcal{A}_2$ be two non-alternating automata. An equivalence relation $\mathcal{R}$ on $S_1 \cup S_2$ is a *strong (probabilistic) bisimulation* if, for each pair of states $q, r \in S_1 \cup S_2$ such that $q \;\mathcal{R}\; r$, if $q \xrightarrow{a} \mu$ for some measure μ, then there exists a measure μ' such that $\mu \equiv_{\mathcal{R}} \mu'$ and $r \xrightarrow{a} \mu'$ ($r \xrightarrow{a}_{\mathrm{C}} \mu'$).

The non-alternating automata $\mathcal{A}_1$ and $\mathcal{A}_2$ are said to be strongly (probabilistic) bisimilar if there exists a strong (probabilistic) bisimulation $\mathcal{R}$ on $S_1 \cup S_2$ such that $\bar{s}_1 \;\mathcal{R}\; \bar{s}_2$. We denote this strong (probabilistic) bisimulation relation by $\sim_{\boldsymbol{NA}}$ ($\sim^p_{\boldsymbol{NA}}$).

5.5 Weak Bisimulations in NA

Again, the definition of weak bisimulations is based on some auxiliary machinery to describe weak transitions.

An *execution fragment* of a non-alternating automaton $\mathcal{A}$ is a finite or infinite sequence of alternating states and actions $\xi = s_0 a_1 s_1 a_2 s_2 \ldots$ starting from a state and, if the sequence is finite, ending with a state, such that, for each i there exists a transition (s_{i-1}, a_i, μ_i) of $\mathcal{D}$ where $\mu_i[s_i] > 0$. If the sequence ξ is finite, then denote by $\mathrm{last}(\xi)$ the last state of ξ. The *length* of ξ, denoted by $|\xi|$, is the number of occurrences of actions in ξ. If ξ is infinite, then $|\xi| = \infty$. Denote by $\mathit{frags}(\mathcal{A})$ the set of execution fragments of $\mathcal{A}$ and by $\mathit{frags}^*(\mathcal{A})$ the set of finite execution fragments of $\mathcal{A}$. An execution fragment ξ is a *prefix* of an execution fragment ξ', denoted by $\xi \preceq \xi'$, if the sequence ξ is a prefix of the sequence ξ'. The *trace* of an execution fragment ξ, denoted by $\mathrm{trace}(\xi)$, is the sub-sequence of external actions of ξ.

A *scheduler* for a non-alternating automaton $\mathcal{A}$ is a function $\sigma\colon \mathit{frags}^*(\mathcal{A}) \to \mathit{SubDisc}(\mathcal{D})$ such that for each finite execution fragment ξ, $\sigma(\xi) \in \mathit{SubDisc}(\mathcal{D}(\mathrm{last}(\xi)))$. A scheduler is *Dirac* if it assigns a Dirac measure to each execution fragment.

Given a scheduler σ and a finite execution fragment ξ, the measure $\sigma(\xi)$ describes how transitions are chosen to move from $\mathrm{last}(\xi)$. The resulting combined transition is the combined transition according to $\sigma(\xi)$. We denote by $\mu_{\sigma(\xi)}$ the corresponding target measure.

A scheduler σ and a start state s induce a measure on execution fragments as follows. The sample space is the set of execution fragments that start with s; the σ-field is the σ-field generated by the set of *cones*, sets of the form $C_\xi = \{\xi' \in \mathit{frags}(\mathcal{A}) \mid \xi \preceq \xi'\}$; the measure is the unique extension $\mu_{\sigma,s}$ of the measure defined on cones by the following equation: $\mu_{\sigma,s}[C_{sa_1s_1\ldots a_ns_n}] = \prod_{i\in\{1,\ldots,n\}} \mu_{\sigma(sa_1s_1\ldots a_{i-1}s_{i-1})}[(a_i, s_i)]$.

We say that there is a *weak transition* from a state s to a measure over states μ labeled by an action a, denoted by $s \stackrel{a}{\Longrightarrow} \mu$, if there is a Dirac scheduler σ such that the following holds for the induced measure $\mu_{\sigma,s}$:

1. $\mu_{\sigma,s}[\mathit{frags}^*(\mathcal{A})] = 1$;
2. $\forall_{\xi \in \mathit{frags}^*(\mathcal{A})}$ if $\mu_{\sigma,s}[\xi] > 0$ then $\mathrm{trace}(\xi) = \mathrm{trace}(a)$;
3. $\forall_{q \in S}\ \mu_{\sigma,s}[\{\xi \in \mathit{frags}^*(\mathcal{A}) \mid \mathrm{last}(\xi) = q\}] = \mu[q]$.

If we remove the Dirac condition on the scheduler σ, then we say that there is a *combined weak transition* from s to μ labeled by a, denoted by $s \stackrel{a}{\Longrightarrow}_{\mathrm{C}} \mu$.

Let $\mathcal{A}_1, \mathcal{A}_2$ be two non-alternating automata. An equivalence relation $\mathcal{R}$ on $S_1 \cup S_2$ is a *weak (probabilistic) bisimulation* if, for each pair of states $q, r \in S_1 \cup S_2$ such that $q\ \mathcal{R}\ r$, if $q \xrightarrow{a} \mu$ for some measure μ, then there exists a measure μ' such that $\mu \equiv_{\mathcal{R}} \mu'$ and $r \stackrel{a}{\Longrightarrow} \mu'$ ($r \stackrel{a}{\Longrightarrow}_{\mathrm{C}} \mu'$).

The non-alternating automata $\mathcal{A}_1$ and $\mathcal{A}_2$ are weakly (probabilistic) bisimilar if there exists a weak (probabilistic) bisimulation $\mathcal{R}$ on $S_1 \cup S_2$ such that $\bar{s}_1\ \mathcal{R}\ \bar{s}_2$. Denote this weak (probabilistic) bisimulation relation by $\approx_{\boldsymbol{NA}}$ ($\approx^{p}_{\boldsymbol{NA}}$).

6 Taxonomy

The existing definitions of bisimulation relations suggest three typologies of relations on the alternating models, all with respect to the non-alternating model. We will see that the three typologies capture the existing definitions.

6.1 Nondeterministic Typology

The definition of strong bisimulation for the alternating model is based essentially on nondeterministic states; however, in order not to distinguish automata with ordinary transitions from automata that reach Dirac measures via probabilistic states, the definition of [8] considers probabilistic states as well, and treats them in a special way to capture the idea that they are describing a probabilistic transition. This suggests a typology of bisimulation relation where only nondeterministic states are related and where equivalence is verified on the non-alternating model via transformation.

Let $\mathcal{A}_1$, $\mathcal{A}_2$ be two (strictly) alternating automata. An equivalence relation $\mathcal{R}$ on $N_1 \cup N_2$ is a *(weak)(probabilistic) nondeterministic bisimulation* if $\mathcal{R}$ is a (weak) (probabilistic) bisimulation between $\mathcal{T}(\mathcal{A}_1)$ and $\mathcal{T}(\mathcal{A}_2)$.

Denote by $\sim^{N}_{X}$, $\sim^{pN}_{X}$, $\approx^{N}_{X}$, and $\approx^{pN}_{X}$ the four bisimulations on model X, where X is either ***SA*** or ***A***.

6.2 Divided Typology

The definition of strong bisimulation for the strictly alternating model relates nondeterministic states to nondeterministic states and probabilistic states to probabilistic states; then the step condition is the one of the non-alternating model. This suggests a typology of bisimulation where nondeterministic and probabilistic states are separated and equivalence is verified on the non-alternating model via embedding.

Let $\mathcal{A}_1, \mathcal{A}_2$ be two (strictly) alternating automata and let N be $N_1 \cup N_2$, P be $P_1 \cup P_2$. An equivalence relation $\mathcal{R} \subseteq (N \times N) \cup (P \times P)$ is a *divided (weak) (probabilistic) bisimulation* if $\mathcal{R}$ is a (weak) (probabilistic) bisimulation between $\mathcal{E}(\mathcal{A}_1)$ and $\mathcal{E}(\mathcal{A}_2)$.

Denote by $\sim^{D}_{X}$, $\sim^{pD}_{X}$, $\approx^{D}_{X}$, and $\approx^{pD}_{X}$ the four bisimulations on model X, where X is either ***SA*** or ***A***.

6.3 Mixed Typology

The definition of weak bisimulation for the alternating model does not distinguish between nondeterministic and probabilistic states except for the use of conditional measures from probabilistic states. It turns out that these conditional measures are the key aspect of the formal definition of [8] that leads to the probabilistic variant of weak bisimulation. Besides this, equivalence is verified via embedding.

Let $\mathcal{A}_1$, $\mathcal{A}_2$ be two (strictly) alternating automata. An equivalence relation $\mathcal{R}$ on $S_1 \cup S_2$ is a *(weak) (probabilistic) mixed bisimulation* if $\mathcal{R}$ is a (weak) (probabilistic) bisimulation between $\mathcal{E}(\mathcal{A}_1)$ and $\mathcal{E}(\mathcal{A}_2)$.

Denote by $\sim^{M}_{X}$, $\sim^{pM}_{X}$, $\approx^{M}_{X}$, and $\approx^{pM}_{X}$ the four bisimulations on model X, where X is either ***SA*** or ***A***.

7 Comparative Analysis

In this section we compare the relations defined so far. We analyze separately, within and across models, strong bisimulations, strong probabilistic bisimulations, and weak probabilistic bisimulations. We do not analyze weak bisimulations because they are not all transitive (cf. [4]).

7.1 Strong Bisimulations

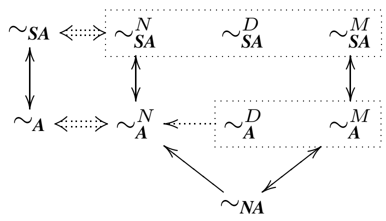

Figure 3. Taxonomy for $\sim$.

Figure 3 compares strong bisimulations within models and across models via embedding. The dotted boxes identify relations that are equivalent within the same model, while dotted arrows are directed from stronger to weaker relations within the same model. The double dotted double arrows show the correspondence between the relations of our taxonomy and the relations already defined in the literature. The other arrows identify preservation and reflection of relations across models via embedding. Thus, the arrow from $\sim_{\mathit{SA}}$ to $\sim_{\mathit{A}}$ means that if two strictly alternating automata $\mathcal{A}_1$ and $\mathcal{A}_2$ are bisimilar according to $\sim_{\mathit{SA}}$, then their embeddings into the alternating model are bisimilar according to $\sim_{\mathit{A}}$, and the arrow from $\sim_{\mathit{A}}$ to $\sim_{\mathit{SA}}$ means that if the embeddings into the alternating model of two strictly alternating automata $\mathcal{A}_1$ and $\mathcal{A}_2$ are bisimilar according to $\sim_{\mathit{A}}$, then the two automata are bisimilar according to $\sim_{\mathit{SA}}$. Other arrows follow by transitivity provided that no upward arrow is combined with a downward arrow.

Observe first that the strong bisimulation relation of [6] coincides with the three typologies in the strictly alternating model and that the strong bisimulation of [8] coincides with the nondeterministic typology in the alternating model, thus confirming our previous intuitions about the typologies.

In the strictly alternating model all typologies identify the same relation. Indeed, strong bisimulations distinguish all kinds of actions, which forces us to relate probabilistic states to probabilistic states only, and an equivalence relation on nondeterministic states induces uniquely an equivalence relation on probabilistic states, which equates equivalences via embedding and via transformations. In the alternating model mixed and divided typologies are equivalent and are stronger than the nondeterministic typology. This is due to the fact that via transformation it is possible to relate ordinary automata with their strictly alternating version, while via embedding such relationships are not possible. This is confirmed also by the relationships across models: the mixed (and divided) relations are equivalent to the non-alternating relation via embedding, while the nondeterministic relations are weaker.

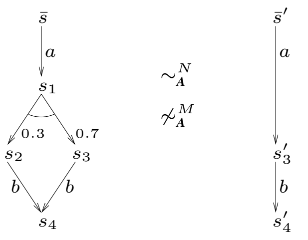

Figure 4. Nondeterministic vs. mixed in A.

Figure 4 gives an example of two alternating automata that are equivalent according to the nondeterministic typology but not according to the mixed typology. Since state s_1 does not appear in the transformations of the two automata, it is enough to put s_2, s_3, s_3' in the same equivalence class. On the other hand, in the embeddings of the two automata there is no state to which s_1 can be related. Observe that the automata of Figure 4 are not the embedding of any strictly alternating automaton. The same figure is a counterexample for the missing arrow from $\sim_{\mathit{A}}^{N}$ to $\sim_{\mathit{NA}}$ in Figure 3.

7.2 Strong Probabilistic Bisimulations

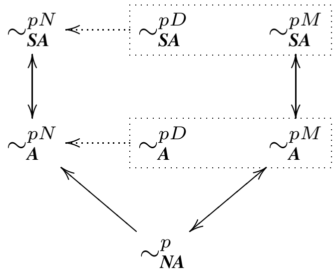

Figure 5. Taxonomy for $\sim^p$.

Figure 5 compares strong probabilistic bisimulations within models and across models via embeddings. The main difference compared to strong bisimulations is that in the strictly alternating model the nondeterministic typology is weaker than the divided and mixed typologies. This is due to the fact that, as observed already in [3], only the removal of probabilistic states permits to use the power of combining transitions. The two strictly alternating automata

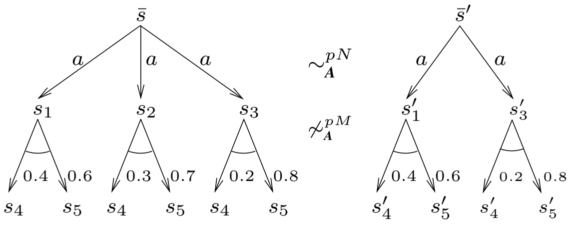

Figure 6. Probabilistic nondeterministic vs. probabilistic mixed.

of Figure 6 give a counterexample for the missing implication from $\sim_{\mathit{A}}^{pN}$ to $\sim_{\mathit{A}}^{pM}$ provided that states s_4, s_4' and states s_5, s_5' are in different equivalence classes, e.g., by enabling different actions. Indeed, the non-alternating transformations of the two automata are equivalent since the middle a-labeled transition of the left automaton can be matched by a uniform combination of the a-labeled transitions of the right automaton; however, the two automata are not equivalent since there is no state to which s_2 can be related.

The relations across models via embedding do not change compared to strong bisimulations.

7.3 Weak Probabilistic Bisimulations

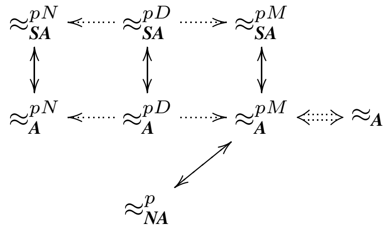

Figure 7. Taxonomy for $\approx^p$.

Figure 7 compares weak probabilistic bisimulations within models and across models via embeddings. In this case all typologies are different within models and in particular the nondeterministic and mixed typologies are incomparable. The divided typology implies both; however, the forced separation between nondeterministic and probabilistic states of the divided typology leads to a very strict relation that is very close to strong probabilistic bisimulation. Considering that in the strong case the divided typology coincides with the mixed typology and that in the weak case the divided typology is too fine, we conclude that this typology is not particularly interesting and we concentrate on the other two, which are not comparable.

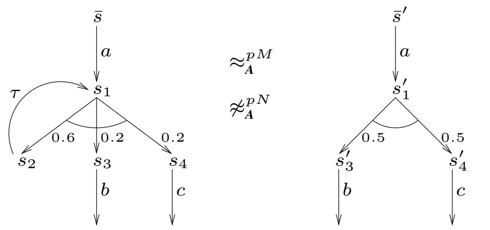

Figure 8. Weak mixed vs. nondeterministic.

Figure 8 shows an example of two automata that are equivalent according to the mixed typology but not according to the nondeterministic typology. The example is valid both for the strictly alternating and alternating models. The example can also be used to develop intuition why the weak bisimulation of [8], which uses conditional measures, is equivalent to the mixed probabilistic typology. The mixed bisimulation for the two automata relates all states with their primed version and adds s_2 to the equivalence class of s_1. The crucial part is to show that s_1 and s_1' match each other's transitions. According to the mixed typology, the transition from s_1 can be matched from s_1' by scheduling the only available transition with probability 0.4 and the transition from s_1' can be matched from s_1 by scheduling transitions until s_3 or s_4 are reached; according to the definition of [8], the measure reachable from s_1 and s_1' conditional on leaving the equivalence class of s_1 should be equivalent. In this case both measures are uniform over the classes of s_3 and s_4. The two automata of Figure 8, though, are not bisimilar according to the nondeterministic typology because there is no nondeterministic state to which state s_2 can be related. Therefore, the transition labeled by a from $\bar{s}$ cannot be matched from $\bar{s}'$.

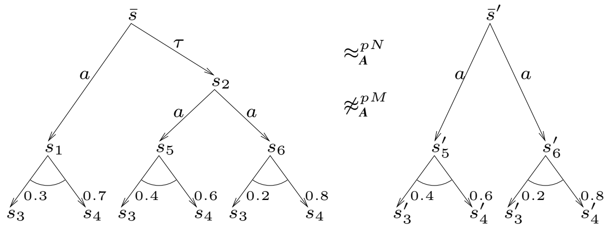

Figure 9. Weak nondeterministic vs. mixed.

Figure 9 shows an example of two automata that are equivalent according to the nondeterministic typology but not according to the mixed typology. The example is valid for the alternating model, but can be adapted to the strictly alternating model by splitting the transition from $\bar{s}$ to s_2. The nondeterministic bisimulation between the two automata relates all nondeterministic states to their primed version and adds s_2 to the equivalence class of $\bar{s}$. The only critical part in showing bisimilarity is matching the left transition from $\bar{s}$ in the transformed left automaton. Since the transformation of the automata removes all probabilistic states, it is enough to combine the two transitions from $\bar{s}'$ with probability $1/2$ each. The two automata, on the other hand, are not bisimilar according to the mixed typology because there is no state to which s_1 can be related.

If we compare relations across models, then the only observation is that the bisimulation relation in the non-alternating model coincides with the mixed typology via embedding, as shown in Figure 7.

7.4 All Bisimulations Together

Figure 10 summarizes all results presented so far, excluding the divided typology, and in addition compares all relations of the same typology within each model. As expected, within each model and each typology, strong bisimulation implies strong probabilistic bisimulation which implies weak probabilistic bisimulation. However, in the mixed typology, strong and strong probabilistic bisimulations coincide since probabilistic states do not allow us to

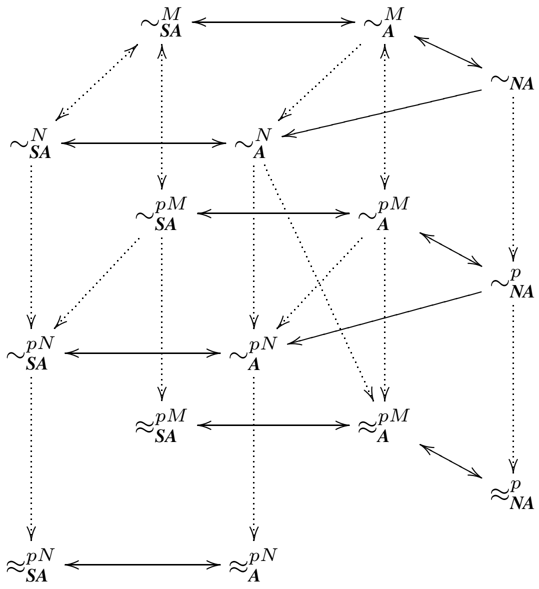

Figure 10. Complete taxonomy.

combine transitions. We have also included an arrow in the alternating model from strong nondeterministic bisimulation to weak probabilistic mixed bisimulation to show that, although the strong and weak bisimulation of [8] are defined in an inconsistent way according to our classification, strong bisimulation implies weak bisimulation as well.

8 Concluding Remarks

We have considered three models for nondeterministic probabilistic systems, the strictly alternating model of Hansson, the alternating model of Philippou, Lee, and Sokolsky, and the non-alternating model of Segala, and we have considered the notions of bisimulations defined on such models. We have introduced three typologies of bisimulation, the nondeterministic, divided and mixed typologies, that capture the ideas behind the existing definitions, and we have compared the relations within models and across models via embeddings.

We have observed that the divided typology, used in the strictly alternating model, is too strict when moving to weak bisimulation and coincides with the mixed typology in the other cases, and thus we have discarded it. The other two typologies turn out to be incomparable when moving to weak bisimulations, and it is not clear which one is better. We do prefer the mixed typology, however, because the nondeterministic typology does not introduce anything new, in the sense that it uses transformations essentially to view the probabilistic states of the alternating models just as a technical artifact to represent a non-alternating automaton.

We believe that the results of this paper confirm the idea that, although there are several proposals of models and of bisimulation relations, all relations are essentially the same and everything can be understood in a uniform way by working in the general model of probabilistic automata. Our objective is to investigate whether the same idea applies to branching bisimulation.

References

[1] C. Baier and H. Hermanns, "Weak bisimulation for fully probabilistic processes," CTIT, University of Twente, Tech. Rep. 99/12, 1999.

[2] C. Baier, H. Hermanns, J-P.Katoen, and V. Wolf, "Comparative branching-time semantics for markov chains (extended abstract)," *Proc. CONCUR*, lncs 2761, pp. 492–507, 2003.

[3] E. Bandini and R. Segala, "Axiomatizations for probabilistic bisimulation," *Proc.* 28^{th} *ICALP*, lncs 2076, pp. 370–381, 2001.

[4] Y. Deng, "Axiomatisations and types for probabilistic and mobile processes" Ph.D. th., Ec. de Mines, 2005.

[5] R. v. Glabbeek, S. Smolka, B. Steffen, and C. Tofts, "Reactive, generative, and stratified models of probabilistic processes" *Proc.* 5^{th} *LICS*, pp. 130–141, 1990.

[6] H. Hansson, "Time and probability in formal design of distributed systems," Ph.D. th., Uppsala, 1991.

[7] K. Larsen and A. Skou, "Bisimulation through probabilistic testing," *Proc.* 16^{th} *POPL*, pp. 344–352, 1989.

[8] A. Philippou, I. Lee, and O. Sokolsky, "Weak bisimulation for probabilistic systems," *Proc. CONCUR*, lncs 1877, pp. 334–349, 2000.

[9] M. Rabin, "Probabilistic automata," *Information and Control*, vol. 6, pp. 230–245, 1963.

[10] R. Segala, "Modeling and verification of randomized distributed real-time systems," Ph.D. th., MIT, 1995.

[11] R. Segala and N. Lynch, "Probabilistic simulations for probabilistic processes," *Proc. CONCUR*, lncs 836, pp. 481–496, 1994.

[12] A. Sokolova and E. d. Vink, "Probabilistic automata: system types, parallel composition and comparison," *Validation of Stochastic Systems*, lncs 2925, 2004.

[13] M. Vardi, "Automatic verification of probabilistic concurrent finite-state programs," *Proc. 26th FOCS*, pp. 327–338, 1985.