

Lattice structures for bisimilar Probabilistic Automata

Johann Schuster

University of the Federal Armed Forces Munich
Neubiberg, Germany

johann.schuster@unibw.de

Markus Siegle

University of the Federal Armed Forces Munich
Neubiberg, Germany

markus.siegle@unibw.de

Abstract The paper shows that there is a deep structure on certain sets of bisimilar Probabilistic Automata (PA). The key prerequisite for these structures is a notion of compactness of PA. It is shown that compact bisimilar PA form lattices. These results are then used in order to establish normal forms (in the sense of [4]) not only for finite automata, but also for infinite automata, as long as they are compact.

1 Introduction

Probabilistic automata (PA) [11, 12, 8] are a powerful and popular modelling formalism, since they allow one to reason about the behaviour of systems which feature both randomness and nondeterminism. For probabilistic automata (PA), several notions of simulations and bisimulations have been defined in the literature [11, 10]. Bisimulation relations – in general, but also in particular for probabilistic automata – are employed for characterising equivalent behaviour. Thus they may serve as the basis for checking whether two systems are equivalent in some sense. As a straightforward consequence, bisimulation relations are also very valuable for reducing the size of a system, by replacing it with an equivalent but smaller one. Hereby the goal is to find the smallest possible bisimilar system, i.e. the minimal one. We concentrate on two notions of bisimulation for PA: strong probabilistic bisimulation and weak probabilistic bisimulation. For automata with finite sets of states and transitions, both are known to be decidable in polynomial time [2, 6].

Recently, the question of calculating minimal canonical forms (i.e. normal forms) of probabilistic automata has been tackled (again, for automata with finite sets of states and transitions) [4], where it turned out that this problem can also be solved in polynomial time. In the present paper, we go one step further: We show that finiteness is not required for defining minimal canonical forms. We point out that also for automata with countably infinite (“countable” for short) state space, countable set of actions, and possibly uncountable set of transitions, there are notions of normal forms. However, we show that, in contrast to the finite case, normal forms do not always exist. It rather turns out that an auxiliary condition, namely compactness [3], is crucial for the existence of normal forms.

In the context of PA, the distributions reached by probabilistic schedulers form convex sets. The first ones to use this observation were Segala and Cattani who developed a decision algorithm out of this fact [2]. In this paper, we combine the strongly geometric ideas of [2] with the ideas of compact automata of [3]. This enables us to extend the recent results of [4] to a class of PA with countable state space. In this way we may use geometrical ideas to show that normal forms of PA arise naturally as some “generating points” – in a strong or weak sense – of convex sets. However, we also show (by a counterexample) that in general one cannot expect normal forms for arbitrary PA with countable state space. For normal forms wrt. strong bisimulation, we may directly use a classical result from functional analysis, the theorem of Krein-Milman [7], that directly extends the ideas of [2] to the case of countable state spaces. For weak bisimulation, we extend the results of [4] by adding some compactness assumptions.

This paper is, to the best of our knowledge, the first approach to derive lattice structures on bisimilar objects. Technically, we work with partially ordered sets of bisimilar PA on which lattice structures are established, such that the normal form corresponds to the bottom element of the lattice. We show that there are unique bottom elements in the lattices we define. We work on quotients of PA, but do not address the question of how to find such quotients for arbitrary automata. For the countably infinite case, ideas from [1, 5] could be used to calculate quotients, but we do not investigate this further. However, even if it is hard to find quotients for infinite state systems, we feel that the lattice structures we establish are interesting just from an abstract point of view.

The paper is organised as follows: Sec. 2 recalls some basic facts on preorders, lattices and probabilistic automata, compact sets and extreme points. It also defines the notion of compact PA which is essential for our paper. Operations on bisimilar quotients and sets of quotients (intersection, union, rescaling) are defined in Sec. 3. The core of the paper consists of the results on lattice structures given in Sec. 4. Some illustrating examples are provided in Sec. 5, and Sec. 6 concludes the paper.

2 Preliminaries

Definition 1. *The disjoint union of two sets S_1 and S_2 is defined as $S_1 \dot{\cup} S_2 := \cup_{i \in \{1,2\}} \{(x, i) | x \in S_i\}$.*

The disjoint union is defined up to isomorphism, which implies commutativity $S_1 \dot{\cup} S_2 = S_2 \dot{\cup} S_1$. There are canonical embeddings $S_1 \rightarrow S_1 \dot{\cup} S_2, x \mapsto (x, 1)$ and $S_2 \rightarrow S_1 \dot{\cup} S_2, x \mapsto (x, 2)$.

2.1 Partial orders and lattices

Definition 2. *A partial order is a binary relation \leq over a set S which is antisymmetric, transitive, and either reflexive or irreflexive, i.e., for all a, b , and c in S , we have that:*

- *If $a \leq b$ and $b \leq a$ then $a = b$ (antisymmetry).*
- *If $a \leq b$ and $b \leq c$ then $a \leq c$ (transitivity).*
- *Either: $a \leq a$ (reflexivity) for all $a \in S$, or: $a \not\leq a$ (irreflexivity) for all $a \in S$.*

A set with a partial order is called a partially ordered set (also called a poset). For a poset we write $a < b$ iff $a \leq b$ and $a \neq b$ (and similar $a > b$).

Definition 3. *Let (S, \leq) be a poset, and let a and b be two elements in S . An element c of S is the meet (or greatest lower bound or infimum) of a and b , if the following two conditions are satisfied:*

- *$c \leq a$ and $c \leq b$ (i.e., c is a lower bound of a and b).*
- *For any $d \in S$, such that $d \leq a$ and $d \leq b$, we have $d \leq c$ (i.e., c is greater than or equal to any other lower bound of a and b).*

An element c of S is the join (or least upper bound or supremum) of a and b , if the following two conditions are satisfied:

- *$a \leq c$ and $b \leq c$ (i.e., c is a upper bound of a and b).*
- *For any $d \in S$, such that $a \leq d$ and $b \leq d$, we have $c \leq d$ (i.e., c is smaller than or equal to any other upper bound of a and b).*

Remark 1. *If there is a meet (join) of a and b , then indeed it is unique, since if both c and c' are greatest lower bounds (least upper bounds) of a and b , then $c \leq c'$ and $c' \leq c$, whence indeed $c' = c$.*

Definition 4. A poset (L, \leq) is a lattice if it satisfies the following two axioms.

- (Existence of binary meets) For any two elements a and b of L , the set $\{a, b\}$ has a meet: $a \wedge b$ (also known as the greatest lower bound, or the infimum).
- (Existence of binary joins) For any two elements a and b of L , the set $\{a, b\}$ has a join: $a \vee b$ (also known as the least upper bound, or the supremum).

A lattice is called bounded, if it has a least and a greatest element, i.e. elements l, g such that for all $x \in L$: $x \leq g$ and $l \leq x$. We will also write $\perp(L) = l$, $\top(L) = g$. A lattice is called complete, if meet and join exist for all subsets $A \subseteq L$.

We will use the following simple but elementary lemma:

Lemma 1 (descending chain condition (DCC)). Let (S, \leq) be a poset, then the following statements are equivalent:

- Every nonempty subset $A \subseteq S$ contains an element minimal in A
- S contains no infinite descending chain $a_0 > a_1 > a_2 > \dots$

2.2 Probabilistic Automata

First we define the notion of discrete subdistribution and related terms and notations:

Definition 5 ((Sub-)distributions). Let S be a countable set. A mapping $\mu : S \rightarrow [0, 1]$ is called (discrete) subdistribution, if $\sum_{s \in S} \mu(s) \leq 1$. As usual we write $\mu(S')$ for $\sum_{s \in S'} \mu(s)$. The support of μ is defined as $\text{Supp}(\mu) := \{s \in S \mid \mu(s) > 0\}$. The empty subdistribution μ_\emptyset is defined by $\text{Supp}(\mu_\emptyset) = \emptyset$. The size of μ is defined as $|\mu| := \mu(S)$. A subdistribution μ is called distribution if $|\mu| = 1$. The sets $\text{Dist}(S)$ and $\text{Subdist}(S)$ denote distributions and subdistributions defined over the set S . Let $\Delta_s \in \text{Dist}(S)$ denote the Dirac distribution on s , i.e. $\Delta_s(s) = 1$. For two subdistributions μ, μ' the sum $\mu'' := \mu \oplus \mu'$ is defined as $\mu''(s) := \mu(s) + \mu'(s)$ (as long as $|\mu''| \leq 1$). As long as $c \cdot |\mu| \leq 1$, we denote by $c\mu$ the subdistribution defined by $(c\mu)(s) := c \cdot \mu(s)$. For a subdistribution μ and a state $s \in \text{Supp}(\mu)$ we define $\mu - s$ by

$$(\mu - s)(t) = \begin{cases} \mu(t) & \text{for } t \neq s \\ 0 & \text{for } t = s \end{cases}$$

Definition 6 (Lifting of relations on states to distributions). Whenever there is an equivalence relation $R \subseteq S \times S$, we may lift it to $\text{Dist}(S) \times \text{Dist}(S)$ in the following way. For $\mu, \gamma \in \text{Dist}(S)$ we write $\mu L(R) \gamma$ (or simply, by abuse of notation, $\mu R \gamma$) if and only if for each $C \in S/R$: $\mu(C) = \gamma(C)$.

Definition 7 (cf. [2]). A probabilistic automaton (PA) P is a tuple (S, Act, T, s_0) , where S is a countable set of states, $s_0 \in S$ is the initial state, Act is a countable set of actions ($\text{Act} = H \cup E$, H hidden actions, E external actions) and $T \subseteq S \times \text{Act} \times \text{Dist}(S)$ is a transition relation (can be uncountable). Whenever $(s, a, \mu) \in T$ we also write $s \xrightarrow{a} \mu$.

In this paper we restrict ourselves to the case where $H = \{\tau\}$, i.e. $E = \text{Act} \setminus \{\tau\}$. Note that by the countability of S it is clear that every distribution over S has at most countable support.

2.2.1 Weak transitions

In the following we use the definitions and terminology of [9], but we leave out the definitions for labelled transition systems. The only major difference is that we do not assume *finite branching*, i.e. for each state s the set $\{(a, \mu) \in Act \times Dist(S) \mid s \xrightarrow{a} \mu\}$ does not have to be finite. Given a transition $tr = (s, a, \mu)$, we denote s by $source(tr)$ and μ by μ_{tr} . An execution fragment of a PA $P = (S, Act, T, s_0)$ is a finite or infinite sequence $\alpha = q_0 a_1 q_1 a_2 q_2 \dots$ of alternating states and actions, starting with a state and, if the sequence is finite, ending in a state, where each $(q_i, a_{i+1}, \mu_{i+1}) \in T$ and $\mu_{i+1}(q_{i+1}) > 0$. State q_0 , the first state of α , is denoted by $fstate(\alpha)$. If α is a finite sequence, then the last state of α is denoted by $lstate(\alpha)$. An *execution* of P is an execution fragment (of P) where $q_0 = s_0$. We let $frags(P)$ denote the set of execution fragments of P and $frags^*(P)$ the set of finite execution fragments of P . Similarly, we let $execs(P)$ denote the set of executions of P and $execs^*(P)$ the set of finite executions. Execution fragment α is a *prefix* of execution fragment α' , denoted $\alpha \leq \alpha'$, if sequence α is a prefix of sequence α' .

The *trace* of an execution fragment α , written $trace(\alpha)$, is the sequence of actions obtained by restricting α to the set of external actions, i.e. $Act \setminus \{\tau\}$. For a set E of executions of a PA P , $traces(E)$ is the set of traces of the executions in E . We say that β is a trace of a PA P if there is an execution α of P with $trace(\alpha) = \beta$. Let $traces(P)$ denote the set of traces of P .

A *scheduler* for a PA P is a function $\sigma : frags^*(P) \rightarrow SubDisc(T)$ such that $tr \in supp(\sigma(\alpha))$ implies that $source(tr) = lstate(\alpha)$. This means that the image $\sigma(\alpha)$ is a *discrete* subdistribution over transitions. The defect of the subdistribution, i.e. $1 - |\sigma(\alpha)|$ is used for stopping in the current state. In other words, a scheduler is the entity that resolves nondeterminism in a probabilistic automaton by choosing randomly either to stop or to perform one of the transitions that are enabled from the current state. A scheduler σ is said to be *deterministic* if for each finite execution fragment α either $\sigma(\alpha)(T) = 0$ or $\sigma(\alpha) = \Delta(tr)$ (Dirac measure for tr) for some $tr \in T$. A scheduler is called *memoryless*, if it depends only on the last state of its argument, that is, for each pair α_1, α_2 of finite execution fragments, if $lstate(\alpha_1) = lstate(\alpha_2)$, then $\sigma(\alpha_1) = \sigma(\alpha_2)$.

A scheduler σ and a discrete initial probability measure $\mu_0 \in Dist(S)$ induce a measure ε on the sigma-field generated by cones of execution fragments as follows. If α is a finite execution fragment, then the *cone* of α is defined by $C_\alpha = \{\alpha' \in frags(P) \mid \alpha \leq \alpha'\}$. The measure ε of a cone C_α is defined recursively: If $\alpha = s$ for some $s \in S$ we define $\varepsilon(C_\alpha) = \mu_0(s)$. If α is of the form $\alpha' a' s'$ it is defined by the equation

$$\varepsilon(C_\alpha) = \varepsilon(C_{\alpha'}) \cdot \sum_{tr \in T(a')} \sigma(\alpha')(tr) \mu_{tr}(s'),$$

where $T(a')$ denotes the set of transitions of T that are labelled by a' . Standard measure theoretical arguments ensure that ε is well defined. We call the measure ε a probabilistic execution fragment of P , and we say that ε is generated by σ and μ_0 .

Consider a probabilistic execution fragment ε of a PA P , with first state s , i.e. $\mu_0 = \Delta(s)$, that assigns probability 1 to the set of all finite execution fragments α with trace $trace(\alpha) = \beta$ for some $\beta \in (Act \setminus \{\tau\})^*$. Let μ be the discrete measure defined by $\mu(s') = \varepsilon(\{\alpha \mid lstate(\alpha) = s'\})$. Then $s \xrightarrow{\beta}_C \mu$ is a *weak combined transition* of P . We call ε a *representation* of $s \xrightarrow{\beta}_C \mu$. If $s \xrightarrow{\beta}_C \mu$ is induced by a deterministic scheduler, we also write $s \xrightarrow{\beta} \mu$. In case $trace(\alpha)$ is empty we write $s \xrightarrow{\tau}_C \mu$.

Let $\{s \xrightarrow{a} \mu_i\}_{i \in I}$ be a collection of transitions of a PA P , and let $\{c_i\}_{i \in I}$ be a collection of probabilities such that $\sum_{i \in I} c_i = 1$. Then the triple $(s, a, \sum_{i \in I} c_i \mu_i)$ is called a (*strong*) *combined transition* of P and we write $s \xrightarrow{a}_C \sum_{i \in I} c_i \mu_i$.

2.3 Bisimulations

Definition 8 (Strong probabilistic bisimulation [12]). *An equivalence relation R on the set of states S of a PA $P = (S, Act, T, s_0)$ is called strong probabilistic bisimulation if and only if $x R y$ implies for all $a \in Act$: $(x \xrightarrow{a} \mu)$ implies $(y \xrightarrow{a}_C \mu')$ with $\mu(C) = \mu'(C)$ for all $C \in S/R$. Two PA are called strongly bisimilar if their initial states are related by a strong probabilistic bisimulation relation on the direct sum of their states.*

Definition 9 (Weak probabilistic bisimulation [12]). *An equivalence relation R on the set of states S of a PA $P = (S, Act, T, s_0)$ is called weak probabilistic bisimulation if and only if $x R y$ implies for all $a \in Act$: $(x \xrightarrow{a} \mu)$ implies $(y \xrightarrow{a}_C \mu')$ with $\mu(C) = \mu'(C)$ for all $C \in S/R$. Two PA are called weakly bisimilar if their initial states are related by a weak probabilistic bisimulation relation on the direct sum of their states.*

Note that (as we always use combined transitions), in the following we generally omit the word “probabilistic” for our bisimulation relations, even if in the sense of [12] we speak about strong probabilistic and weak probabilistic bisimulations. In the sequel we denote, as usual, by \sim a strong bisimulation relation and by \approx a weak bisimulation relation.

It has been shown in [12] that decision of bisimilarities is closely related to convex sets of reachable distributions, i.e. $S_{\sim}(s, a) := \{\mu \in Dist(S) \mid s \xrightarrow{a}_C \mu\} / \sim$ and $S_{\approx}(s, a) := \{\mu \in Dist(S) \mid s \xrightarrow{a}_C \mu\} / \approx$. Those sets are considered modulo the bisimilarities known and splitters are constructed out of them. For details we refer to [12]. The important point is that those sets can be defined for infinite state systems as well. The only difference is that they can no longer be regarded as subsets of \mathbb{R}^n then.

It is clear by definition that the unreachable parts of an automaton do not play any role for bisimilarity, which motivates the following definition.

Definition 10 (Reachable states). *Let $P = (S, Act, T, s_0)$ be a PA, $S' \subseteq S$ its set of reachable states, i.e. those states that can be reached with non-zero probability by a scheduler starting from s_0 . Let $T' := T \upharpoonright_{S' \times Act \times Dist(S')}$ be the restriction of the transition relation to S' . We define $r(P) := (S', Act, T', s_0)$ and call it the reachable fraction of P .*

2.4 Isomorphic & quotient automata

We want to be able to identify automata that are basically the same, only having different names of their states. The following definition formalises this.

Definition 11 (Isomorphic automata). *Let $P = (S, Act, T, s_0)$ and S' a set with $|S'| = |S|$. We call a bijective mapping between sets $\iota : S \rightarrow S'$ a (set-)isomorphism. Via this isomorphism we may push forward distributions on S to distributions on S' by the mapping $\mu \circ \iota^{-1}$, as the following diagram shows.*

$$\begin{array}{ccc}
 S & \xrightarrow{\iota} & S' \\
 \downarrow \mu & \swarrow \mu \circ \iota^{-1} & \\
 [0, 1] & &
 \end{array}$$

That in turn means that we may push forward transitions between states in the set S to transitions between states in the set S' using our isomorphism by the following mapping

$$\begin{array}{ccc}
 \iota : S \times Act \times Dist(S) & \rightarrow & S' \times Act \times Dist(S') \\
 (s, a, \mu) & \mapsto & (\iota(s), a, \mu \circ \iota^{-1})
 \end{array}$$

By this, we may define the automaton $\iota(P) := (\iota(S), Act, \iota(T), \iota(s_0))$ where we, as usual, denote for mappings $f : A \rightarrow B$ by $f(A) := \{f(a) | a \in A\} \subseteq B$ the image of A under f . Two PA $P = (S, Act, T, s_0)$, $P' = (S', Act, T', s'_0)$ are called isomorphic, if there exists an isomorphism $\iota : S \rightarrow S'$ such that $P' = \iota(P)$. For isomorphic automata we write $P \equiv_{iso} P'$.

We would like to stress that we require the *same* set of actions for two automata to be isomorphic. The reason for this is that also for bisimulations the *same* set of actions is considered.

Lemma 2. *The relation \equiv_{iso} is an equivalence relation.*

Proof. Follows by direct verification from the properties of isomorphisms □

Whenever an equivalence relation R is given (R will be chosen to be \sim or \approx), it is common to look at a special automaton that is defined over equivalence classes of states. The following definition formalises this.

Definition 12 (Quotient automaton). *Let $P = (S, Act, T, s_0)$ be a PA and R an equivalence relation over S . We write P/R to denote the quotient automaton of P wrt. R , that is*

$$P/R = (S/R, Act, T/R, [s_0]_R)$$

with $T/R \subseteq S/R \times Act \times Dist(S/R)$ such that $([s]_R, a, \mu) \in T/R$ if and only if there exists a state $s' \in [s]_R$ such that $(s', a, \mu') \in T$ and $\forall [t]_R \in S/R : \mu([t]_R) = \sum_{t' \in [t]_R} \mu'(t')$. We call an automaton a quotient wrt. R if it holds that $P \equiv_{iso} P/R$.

2.5 Compact automata

One key property when searching for lattice structures on PA is compactness. Compactness definitions already have been introduced in [3] for the alternating model.

Definition 13 (adapted from Def. 9 in [3]). *Let $P = (S, Act, T, s_0)$ be a PA. The function d on $Dist(S) \times Dist(S)$ is defined by*

$$d(\mu_1, \mu_2) := \sup_{A \subseteq S} |\mu_1(A) - \mu_2(A)|$$

Even if in [3] it is mentioned without a proof, that function d defined above is really a metric, we give an explicit proof here.

Lemma 3. *The function d in Def. 13 is a metric on $Dist(S)$.*

Proof. It is clear by definition that d is non-negative and symmetric. Identity of indiscernibles is worth a thought: For general measures we only know when $d(\mu_1, \mu_2) = 0$ that μ_1 and μ_2 coincide up to a zero set. As $Dist(S)$ is a set of *discrete* probability measures, the values $\{(s, \mu(s)) | s \in S\}$ completely determine μ . Therefore $d(\mu_1, \mu_2) = 0 \Rightarrow \mu_1 = \mu_2$. Also the triangle inequality holds: $d(\mu_1, \mu_3) = \sup_{A \subseteq S} |\mu_1(A) - \mu_3(A)| = \sup_{A \subseteq S} |\mu_1(A) - \mu_2(A) + \mu_2(A) - \mu_3(A)| \leq \sup_{A \subseteq S} (|\mu_1(A) - \mu_2(A)| + |\mu_2(A) - \mu_3(A)|) \leq \sup_{A \subseteq S} |\mu_1(A) - \mu_2(A)| + \sup_{A \subseteq S} |\mu_2(A) - \mu_3(A)| = d(\mu_1, \mu_2) + d(\mu_2, \mu_3)$ □

Corollary 1. *Let $P = (S, Act, T, s_0)$ be a PA and fix an equivalence relation $R \in \{\sim, \approx\}$ on S , $s \in S$ and $a \in Act$. Then for every pair $(s, a) \in S \times Act$ there is a metric space $(S_R(s, a), d)$.*

Definition 14 (adapted from Def. 10 in [3]). *Let $P = (S, Act, T, s_0)$ be a PA and consider some fixed equivalence relation $R \in \{\sim, \approx\}$ on S , $s \in S$ and $a \in Act$. We say that state s is *a-compact* wrt. R , if the set $S_R(s, a)$ is compact under the metric d . P is called *compact* wrt. R , if $\forall s \in S \forall a \in Act : s$ is *a-compact* wrt. R . We omit “wrt. R ” if the context is clear.*

The definition given by Desharnais et al. can be seen also in the view of products of metric spaces. We recall the definition of a product of metric spaces:

Definition 15. *If (M_i, d_i) , $i \in \mathbb{N}$ are metric spaces, we define a metric on the countably infinite Cartesian product $\prod_{i \in \mathbb{N}} M_i$ by*

$$d(x, y) = \sum_{i=1}^{\infty} \frac{1}{2^i} \frac{d_i(x_i, y_i)}{1 + d_i(x_i, y_i)}.$$

It is well-known that this construction leads to the metric space $(\prod_{i \in \mathbb{N}} M_i, d)$.

This metric metrizes the product topology - under this metric all projection functions are continuous.

We are now able to relate the metric space $(\prod_{(s,a) \in S \times Act} S_R(s, a), d)$ to the definition of compactness given in Def. 14.

Lemma 4 (Reformulation of Def. 14). *Let $P = (S, Act, T, s_0)$ be a quotient of a PA wrt. some equivalence relation R . The compactness of P defined in Def. 14 is equivalent to the compactness of the metric space $(\prod_{(s,a) \in S \times Act} S_R(s, a), d)$.*

Proof. Assume that Def. 14 holds. Tychonoff's theorem (or equivalently the axiom of choice) ensures that the product of compact spaces is compact again, so the product space in Def. 15 is compact.

The projections $\prod_{(s,a) \in S \times Act} S_R(s, a) \rightarrow S_R(s, a)$ are continuous for every pair (s, a) . Continuous mappings between metric spaces map compact sets to compact sets. So Def. 14 also holds. \square

The question about compactness is only relevant in the case of infinite automata. As long as PA are constructed out of finite sets (both states and transitions), they are compact anyway. This is due to the following classical theorem:

Theorem 1 (Heine-Borel). *For a subset of $M \subseteq \mathbb{R}^n$ (\mathbb{R}^n the metric space of all n -Tuples with Euclidean metric) the following statements are equivalent:*

- M is bounded and closed
- every open cover of M has a finite subcover, that is, M is compact

An important theorem for compact sets is the Krein-Milman theorem (see below). It only works for locally convex vector spaces. We use as a basic fact, that the metric space $(l^\infty, d(x, y) = \sup_{i \in \mathbb{N}} |x_i - y_i|)$ of bounded sequences in \mathbb{R} is locally convex (using the seminorms $p_i(\{x_n\}_n) = |x_i|$, $i \in \mathbb{N}$). Sequences may be used to characterise distributions in the following way. Assume that the states are ordered by the natural numbers. Define the mapping $f : (S_R(s, a), d) \rightarrow (l^\infty, d)$, $\mu \mapsto (\mu(x_1), \mu(x_2), \dots)$. We show that this mapping is continuous, which is straight-forward, as singletons are also subsets. $\sup_{A \subseteq S} |\mu_1(A) - \mu_2(A)| < \varepsilon \Rightarrow \sup_{a \in S} |\mu_1(a) - \mu_2(a)| < \varepsilon$. As continuous images of compact sets are compact, we know that $f(S_R(s, a))$ is compact. By definition of convex schedulers it is also clear that $f(S_R(s, a))$ is convex. The product of $\prod_{(s,a): S_R(s,a) \neq \emptyset} f(S_R(s, a))$ is also compact (Tychonoff's Theorem – accepting the axiom of choice) and convex as a product of convex sets. These sets will be used later to apply the Krein-Milman theorem.


Definition 16 (Extreme points). *An extreme point of a convex set A is a point $x \in A$ with the property that if $x = cy + (1 - c)z$ with $y, z \in A$ and $c \in [0, 1]$, then $y = x$ or $z = x$. $\mathbb{E}(A)$ will denote the set of extreme points of A .*

Now the Krein-Milman theorem says that a compact convex subset of a locally convex vector space is the convex hull of its extreme points:

Theorem 2 (Krein-Milman [7]). *Let A be a compact convex subset of a locally convex vector space X , then $A = CHull(\mathbb{E}(A))$.*

3 Operations on bisimilar quotient automata

Definition 17 (Set operations on bisimilar quotients). Let $P = (S, \text{Act}, T, s_0)$ and $P' = (S', \text{Act}, T', s'_0)$ be PA. Let $S^\circ := S \dot{\cup} S'$ be the disjoint union of state spaces and consider a fixed bisimulation relation $R \in \{\sim, \approx\}$, $R \subseteq S^\circ \times S^\circ$. Let P and P' be bisimilar, i.e. $s_0 R s'_0$. In the following, we assume that P and P' are quotients wrt. R . The intersection of S and S' (and the union) is performed in S°/R by means of the canonical mappings




With this diagram in mind we may define

$$S' \cap_R S = \{[x] \in S^\circ/R \mid \exists s \in S : s \in [x] \wedge \exists s' \in S' : s' \in [x]\}$$

$$S' \cup_R S = \{[x] \in S^\circ/R \mid \exists s \in S : s \in [x] \vee \exists s' \in S' : s' \in [x]\}$$

Using these canonical mappings, similar operations can be defined on the sets of transitions. Clearly $T \subseteq S \times \text{Act} \times \text{Dist}(S)$ and $T' \subseteq S' \times \text{Act} \times \text{Dist}(S')$. The set operations are performed in the set $S^\circ/R \times \text{Act} \times \text{Dist}(S^\circ/R)$ by means of the corresponding canonical mappings.



We thus may define (using the abbreviation $\mathfrak{T} = S^\circ/R \times \text{Act} \times \text{Dist}(S^\circ/R)$):

$$T' \cap_R T = \{([x], a, [\mu]) \in \mathfrak{T} \mid (\exists (s, a, \gamma) \in T : s \in [x] \wedge \gamma \in [\mu]) \wedge (\exists (s', a, \gamma') \in T' : s' \in [x] \wedge \gamma' \in [\mu])\}$$

$$T' \cup_R T = \{([x], a, [\mu]) \in \mathfrak{T} \mid (\exists (s, a, \gamma) \in T : s \in [x] \wedge \gamma \in [\mu]) \vee (\exists (s', a, \gamma') \in T' : s' \in [x] \wedge \gamma' \in [\mu])\}$$

Finally we note that s_0 has to be mapped to S°/R . Summing up, we define the intersection quotient wrt. R as

$$P \cap P' := (S' \cap_R S, \text{Act}, T \cap_R T', [s_0]_R)$$

and, similarly, the union quotient wrt. R as

$$P \cup P' := (S' \cup_R S, \text{Act}, T \cup_R T', [s_0]_R)$$

With the same mappings one may define $P \subseteq P'$ if and only if $S \subseteq S'$ (in S°/R) and $T \subseteq T'$ in $S^\circ/R \times \text{Act} \times \text{Dist}(S^\circ/R)$.

As we target on bisimilar automata in this paper, we would like to stress that only the reachable fraction of automata is relevant. Our quotient definition and the above set operations are defined for the general case, where unreachable states may occur, but of course they apply also to the case where only reachable states are present.

Lemma 5. Let $R \in \{\sim, \approx\}$ and $P = (S, \text{Act}, T, s_0)$, $P' = (S', \text{Act}, T', s'_0)$ be bisimilar quotients of PA wrt. R . It holds that $P \cap P' \equiv_{iso} P' \cap P$ and $P \cup P' \equiv_{iso} P' \cup P$.

Proof. Clear by identifying the different direct sums according to R \square

As it will turn out that τ -loops leading back to the same state (as part of a distribution) can disturb the lattice property in the case of weak bisimulation, we give the following definition.

Definition 18 (Rescaled Automata). *An automaton $P = (S, Act, T, s_0)$ is called rescaled, if for all $(s, \tau, \mu) \in T$ it holds either that $\mu(s) = 0$ or $\mu(s) = 1$.*

Note that the rescaling as defined in Def. 18 is only one possibility (incidentally the one leading to the smallest transition fanout), but different rescalings, where each τ transition would loop back to its source state with fixed probability $0 \leq p < 1$ (or probability 1 for “loops”) would also work.

Remark 2. *Let $P = (S, Act, T, s_0)$ be a PA. We split its set of transitions into the sets $T_{-\tau} := \{(s, a, \mu) \in T \mid a \neq \tau\}$, $T_{\Delta} = \{(s, \tau, \mu) \in T \mid \mu = \Delta_s\}$ and $T_{-\Delta} = \{(s, \tau, \mu) \in T \mid \mu \neq \Delta_s\}$. Now we define the function $res : T_{-\Delta} \rightarrow S \times Act \times Dist(S)$ by $res(s, \tau, \mu) := (s, \tau, \frac{1}{1-v(s)}(v - s))$ which is well-defined, as we always have $v(s) \neq 1$. With*

$$T' := T_{-\tau} \cup T_{\Delta} \cup res(T_{-\Delta})$$

we may define the rescaled automaton $P^{res} := (S, Act, T', s_0)$. When using randomised schedulers, it is a basic fact, that $P \approx P^{res}$.

3.1 Sets of quotients

Definition 19 (Sets of quotients). *Let P be an automaton and \mathcal{PA} be the set of all PA. Define the set*

$$\Omega_{\sim}(P) := \{A \in \mathcal{PA} \mid A \text{ quotient wrt. } \sim, A \sim P\}$$

and the set

$$\Omega_{\approx}(P) := \{A \in \mathcal{PA} \mid A \text{ quotient wrt. } \approx, A \approx P, A \text{ rescaled}\}$$

The reachable fractions can also be considered:

$$\Omega_{\sim}^*(P) := \{A \in \mathcal{PA} \mid A \text{ quotient wrt. } \sim, A \sim P, A = r(A)\}$$

and the set

$$\Omega_{\approx}^*(P) := \{A \in \mathcal{PA} \mid A \text{ quotient wrt. } \approx, A \approx P, A \text{ rescaled}, A = r(A)\}$$

Of course, the automata in the set $\Omega_{\sim}^*(P)$ (in the set $\Omega_{\approx}^*(P)$) all have the same number of states.

Lemma 6. *It holds that $\Omega_{\sim}^*(P) \subseteq \Omega_{\sim}(P)$ and $\Omega_{\approx}^*(P) \subseteq \Omega_{\approx}(P)$.*

Remark 3. *As usual for equivalence relations, we may consider quotients wrt. isomorphism, i.e. $\Omega_{\sim}(P)/\equiv_{iso}$ and $\Omega_{\approx}(P)/\equiv_{iso}$. Without loss of generality we may identify the states of every automaton in these quotients by a subset of the natural numbers \mathbb{N} (using 1 for the initial states). For the sets $\Omega_{\sim}(P)$ and $\Omega_{\approx}(P)$ such an enumeration cannot be performed, as there are uncountably many unreachable parts (this is shown in detail by Lemma 13).*

For the rest of the paper we now assume that the states in $\Omega_{\sim}^*(P)/\equiv_{iso}$ and $\Omega_{\approx}^*(P)/\equiv_{iso}$ are consecutive natural numbers, starting with 1 for the initial state.

The rest of this section is devoted to show that for compact P there are well-defined operations

$$\cap : \Omega_{\sim}(P)/\equiv_{iso} \times \Omega_{\sim}(P)/\equiv_{iso} \rightarrow \Omega_{\sim}(P)/\equiv_{iso}$$

$$\cap : \Omega_{\approx}(P)/\equiv_{iso} \times \Omega_{\approx}(P)/\equiv_{iso} \rightarrow \Omega_{\approx}(P)/\equiv_{iso}$$

Similar operations exist for \cup , but their existence and well-definedness is clear. Note that it is a priori not clear that the result is again in $\Omega_{\sim}(P)/\equiv_{iso}$ (or $\Omega_{\approx}(P)/\equiv_{iso}$, respectively), with other words that the results are again bisimilar to P . This will be shown in the following.

Remark 4 (Counterexample). *We show that there is in general no well-defined operation $\cap : \mathfrak{Q}_{\sim}(P)/\equiv_{iso} \times \mathfrak{Q}_{\sim}(P)/\equiv_{iso} \rightarrow \mathfrak{Q}_{\sim}(P)/\equiv_{iso}$, that means the compactness of P is crucial. We use the (non-compact) automata*

$$P := (\{1, 2, 3\}, \{\tau, a, b\}, \{(1, \tau, \frac{1}{n}\Delta_2 \oplus (1 - \frac{1}{n})\Delta_3)\}_{n \in \mathbb{N}, n \geq 2} \cup \{(2, a, \Delta_2), (3, b, \Delta_3)\}, 1)$$

and

$$P' := (\{1, 2, 3\}, \{\tau, a, b\}, \{(1, \tau, \frac{e}{n}\Delta_2 \oplus (1 - \frac{e}{n})\Delta_3)\}_{n \in \mathbb{N}, n \geq 6} \\ \cup \{(1, \tau, \frac{1}{2}\Delta_2 \oplus \frac{1}{2}\Delta_3), (2, a, \Delta_2), (3, b, \Delta_3)\}, 1).$$

It is easy to see that $P \sim P'$ and $P, P' \in \mathfrak{Q}_{\sim}(P)/\equiv_{iso}$. It is clear that the intersection of both automata is only

$$P^\cap := (\{1, 2, 3\}, \{\tau, a, b\}, \{(1, \tau, \frac{1}{2}\Delta_2 \oplus \frac{1}{2}\Delta_3), (2, a, \Delta_2), (3, b, \Delta_3)\}, 1),$$

as e is not a rational number. So the result of the intersection is clearly not bisimilar, i.e. not in $\mathfrak{Q}_{\sim}(P)/\equiv_{iso}$. The reason for this is, that both automata allow for a limiting distribution $(1, \tau, \Delta_3)$, but both don't reach this distribution. As the ways how this limit is reached (i.e. the sequences) are different, the intersection no longer leads to a bisimilar result. Note that there are many other examples (e.g. every irrational root may be taken instead of e).

Thus we have shown:

Lemma 7. *In general $(\mathfrak{Q}_{\sim}(P)/\equiv_{iso}, \subseteq)$ and $(\mathfrak{Q}_{\sim}^*(P)/\equiv_{iso}, \subseteq)$ are not lattices.*

4 Lattice structures

For the rest of this paper we consider a subset of PA where $(\mathfrak{Q}_R(P)/\equiv_{iso}, \subseteq)$ and $(\mathfrak{Q}_R^*(P)/\equiv_{iso}, \subseteq)$, $R \in \{\sim, \approx\}$ are lattices.

Lemma 8 (compact sets of quotients). *Let $R \in \{\sim, \approx\}$. The PAs contained in $\mathfrak{Q}_R(P)/\equiv_{iso}$ are either all compact, or all non-compact. Thus the property of compactness is well-defined for $\mathfrak{Q}_R(P)/\equiv_{iso}$.*

Proof. Let P be a PA, $R \in \{\sim, \approx\}$ and $P' \in \mathfrak{Q}_R(P)/\equiv_{iso}$. P' is compact if and only if P is compact, as the sets $S_R(s, a)$ are unique up to isomorphism. \square

In the following we will only consider compact PA and therefore also compact sets of quotients. In metric spaces, compactness is equivalent to sequence compactness.

Lemma 9. *$(\mathfrak{Q}_{\sim}(P)/\equiv_{iso}, \subseteq)$ is a poset. If P is compact, it is even a lattice with union (intersection) of automata as join (meet) operations. Intersections are also possible over arbitrary sets of bisimilar automata.*

Proof. By Krein-Milman Theorem (Theorem 2) we know that for countable state spaces the sets $S_{\sim}(s, a)$ (seen as subsets of l^∞) have a unique set of extreme points $\mathbb{E}(S_{\sim}(s, a))$. This set is included in every bisimilar automaton, therefore in every intersection. This is the reason why the intersection of all automata in a set of bisimilar automata still leads to a bisimilar automaton. Transitions leading to some distribution in $S_{\sim}(s, a)$ that is not extreme do not change the bisimilarity. This is the reason why the union of two bisimilar automata leads to a bisimilar automaton. Finally observe that the unreachable fraction of states and transitions does not play any role for bisimilarity. \square

In the strong case, a minimal set of transitions leading to extreme points can be chosen from the strong transitions emanating from s . In the weak case this does not have to be the case, as an extreme point might be reached transitively via some intermediate distributions.

Astonishingly, there's a similar statement also for weak bisimilarity, when only rescaled automata are considered. The interesting point is that the arguments of the proof of Lemma 9 do not apply one-to-one, as the extreme points can now also be weak transitions. Therefore it is (with the arguments of the proof for strong bisimulation) not clear that the intersection leads again to a weakly bisimilar PA.

Lemma 10. $(\Omega_{\approx}(P)/\equiv_{iso}, \subseteq)$ is a poset. When $\Omega_{\approx}(P)/\equiv_{iso}$ is compact, it is even a lattice with union (intersection) of automata as join (meet) operations. Intersections are also possible over arbitrary sets of bisimilar automata.

Proof. The meet operation for *finite* PA is justified by Lemma 12 in [4] and the observation that the unreachable fraction of states and transitions does not play any role for bisimilarity. The join operation is clear by definition of weak bisimilarity and the same observation.

For the infinite case assume for the rest of the proof that there are two quotients P_1 and P_2 , $P_1 \approx P_2$ and we identify the (countable) reachable state spaces S_1 and S_2 as in Def. 17 by S°/R (note that this induces a bijection).

Now fix one state $s \in S^\circ/\approx$. By assumption it must have two sets of emanating a -transitions – one in P_1 , one in P_2 . During the proof we use the notation $S_R^i(s, a)$, $i \in \{1, 2\}$ to denote the sets of reachable distributions of automaton 1 and 2, respectively. We have to show that already the intersection of both sets is enough to generate $S_{\approx}(s, a)$. The main difference to the strong case is that we only know that $S_{\approx}^1(s, a) = S_{\approx}^2(s, a)$ for all $a \in Act$, but this does not necessarily imply that also $S_{\approx}^1(s, a) = S_{\approx}^2(s, a)$. According to Def. 19 both automata P_1 and P_2 are rescaled. We consider the one-step-transitions $T_a^i = \{\mu \in Dist(S^\circ/R) \mid (s, a, \mu) \in T_i\}$. Assume further that the identical transitions (modulo bijection) already have been identified (in $S^\circ/R \times Act \times Dist(S^\circ/R)$) and are denoted $T_a^\cap = T_a^1 \cap T_a^2$ and let $T_a^{-\cap i} = T_a^i \setminus T_a^\cap$. We will show that the identical transitions are enough to cover all the behaviour, i.e. all transitions in $T_a^{-\cap i}$, $i \in \{1, 2\}$, can be omitted.

We will first consider the case $a = \tau$. We fix an arbitrary $\mu \in T_\tau^{-\cap 1}$ for which we show that it can be generated by transitions from T_τ^\cap . It is clear that both sets T_τ^1 and T_τ^2 generate the set $S_{\approx}(s, \tau)$ (in the sense that all weak rescaled transitions must start with a combination of these transitions as first step).

With the construction from the proof of Lemma 12 in [4] we get a weak transition¹ in $T_1 \setminus \{(s, \tau, \mu)\}$ such that $s \Rightarrow_C \mu$. Note that the construction in [4] has been given for *finite* state spaces, but all constructions there are also possible in the countable compact case due to sequential compactness. We split this transition in its first-step-probability $\mu^{(1)}$, such that $s \rightarrow_C \mu^{(1)} \Rightarrow_C \mu$. We see that without loss of generality we may reduce the set T_1 to the set $T_1 \setminus \{(s, \tau, \mu)\}$ without losing bisimilarity. With this construction we proceed as follows: Pick now some other transition (s, τ, μ') with $\mu' \in T_\tau^{-\cap 1} \setminus \{\mu\}$ that is used somewhere in the weak transition $s \rightarrow_C \mu^{(1)} \Rightarrow_C \mu$. Now, again by the construction in the proof of Lemma 12 in [4], substitute all strong τ transitions from s to μ' by weak transitions in $T_1 \setminus \{(s, \tau, \mu), (s, \tau, \mu')\}$ such that $s \rightarrow_C \mu^{(2)} \Rightarrow_C \mu^{(1)} \Rightarrow_C \mu$ (possibly $\mu^{(2)} = \mu^{(1)}$ when we substituted some transition not taken in the first step). Continuing in the same way yields a sequence of combined one-step-transitions leading to the series of distributions $(\mu^{(i)})_{i \in \mathbb{N}}$.

Now compactness comes into play: This sequence must have limit points in $S_{\approx}(s, \tau)$. If there is one limit point already in T_τ^\cap , by construction there will be a scheduler that uses only transitions from T_τ^\cap

¹The proof in [4] shows that a transition $s \xrightarrow{\tau} \mu$ – there denoted $s \xrightarrow{\tau} v_s$ – is redundant, as long as it is not in the intersection T_τ^\cap , by constructing the above mentioned transition. The only case where it is not redundant leads to a contradiction to the quotient property.

when leaving s . If there is a limit point μ^* in $T_\tau^{-\cap 1}$, there is a non-trivial loop $\mu^* \Rightarrow_c \gamma^* \Rightarrow_c \mu^*$ for some $\gamma^* \in T_\tau^{-\cap 2}$ which would contradict the quotient property as it would render at least two states bisimilar. Therefore we conclude that this case cannot happen and μ can be omitted from $T_\tau^{-\cap 1}$ without losing bisimilarity.

As μ was chosen arbitrarily, this construction can be done for all elements in $T_\tau^{-\cap 1}$ – even if they might be uncountably many – (and analogously for $T_\tau^{-\cap 2}$). Therefore all transitions in $T_\tau^{-\cap 1}$ (and analogously for $T_\tau^{-\cap 2}$) may be omitted without losing bisimilarity.

For the case $a \neq \tau$ we may omit those strong a -transitions that can be mimicked by a weak a transition (after leaving out this strong transition). Observe that there cannot be cyclic relations “ s uses the a transition of t (with probability 1)” and vice versa, because then s and t would be bisimilar, which is a contradiction to the quotient property.

For the general case of an arbitrary intersection (i.e. possibly more than two automata) observe the following: Let A be an arbitrary set of weakly bisimilar quotients. We have to show that the automaton $P_0 = \bigcap_{P \in A} P$ is still bisimilar. Pick an arbitrary automaton $P \in A$ we have to show that all transitions of P that are not in P_0 are redundant. We may find for every such transition tr a bisimilar automaton P_{tr} that does not have such a transition (for otherwise the transition would have to be in the intersection). By the above procedure for two automata we show that the transition tr is redundant. This can be done for all other transitions not in P_0 . \square

Remark 5. *The crucial point in the proof of Lemma 10 is compactness. For general quotients a sequence of redundant distributions may not have a limit in $S_{\approx}(s, \tau)$ (This would be e.g. the distribution $(1, \tau, \Delta_3)$ in Remark 4). Sequential compactness ensures that there exists such a limit in $S_{\approx}(s, \tau)$.*

As we have shown that the intersection of arbitrary sets of bisimilar automata is again bisimilar by Lemma 9 and 10, the following theorem is immediate. It extends Theorem 1 of [4] to the compact automata case.

Theorem 3. *Let P be compact wrt. R . Then $(\mathfrak{Q}_R(P)/\equiv_{iso}, \subseteq)$ has a unique minimal element.*

Definition 20. *Let P be compact wrt. R , $R \in \{\sim, \approx\}$. Then there is a well-defined mapping $\mathcal{N} : PA \rightarrow PA/\equiv_{iso}$, given by $P \mapsto \perp(\mathfrak{Q}_R(P)/\equiv_{iso})$ that assign to every PA P the minimal element in $\mathfrak{Q}_R(P)/\equiv_{iso}$. The mapping to $\mathfrak{Q}_{\sim}(P)/\equiv_{iso}$ is just quotienting, while the mapping to $\mathfrak{Q}_{\approx}(P)/\equiv_{iso}$ is quotienting followed by rescaling. The mapping \mathcal{N} (for R fixed) is called normal form.*

This definition corresponds to the normal form definition given in [4].

Corollary 2. *The notations for normal forms defined in Def. 20 are normal forms in the sense of [4].*

As the bisimilarity only considers the reachable fraction of the state space, the following corollary is immediate.

Corollary 3. *Lemma 9, Lemma 10 and Theorem 3 also hold for $\mathfrak{Q}_{\sim}^*(P)/\equiv_{iso}$ instead of $\mathfrak{Q}_{\sim}(P)/\equiv_{iso}$ and $\mathfrak{Q}_{\approx}^*(P)/\equiv_{iso}$ instead of $\mathfrak{Q}_{\approx}(P)/\equiv_{iso}$.*

Some of the lattices we have constructed above have the appealing property of being bounded.

Lemma 11. *The lattices $(\mathfrak{Q}_{\sim}^*(P)/\equiv_{iso}, \subseteq)$ and $(\mathfrak{Q}_{\approx}^*(P)/\equiv_{iso}, \subseteq)$ are bounded.*

Proof. The lower bound is clearly given by the normal forms defined above. So it remains to show that there is also an upper bound. Let $P = (S, Act, T, s_0)$, for $(\mathfrak{Q}_{\sim}^*(P)/\equiv_{iso}, \subseteq)$ and $(\mathfrak{Q}_{\approx}^*(P)/\equiv_{iso}, \subseteq)$ we know that it is sufficient to consider $(\mathfrak{Q}_{\sim}^*(r(P))/\equiv_{iso}, \subseteq)$ and $(\mathfrak{Q}_{\approx}^*(r(P))/\equiv_{iso}, \subseteq)$. As S is countable, also the set of reachable states will be countable. As there is no restriction on T wrt. finiteness, the union over all transition relations in $\mathfrak{Q}_R^*(P)/\equiv_{iso}$ (identified by R) will be a valid transition relation (i.e. \bigcup_R of all transition relations of quotients in the set $\mathfrak{Q}_R^*(P)/\equiv_{iso}$). \square

Lemma 12. *For a finite automaton P ,*

1. *the maximal elements in $(\Omega^*(P)/\equiv_{iso}, \subseteq)$ and $(\Omega_{\approx}^*(P)/\equiv_{iso}, \subseteq)$ are not necessarily finite.*
2. *the minimal elements in $(\Omega^*(P)/\equiv_{iso}, \subseteq)$ and $(\Omega_{\approx}^*(P)/\equiv_{iso}, \subseteq)$ are finite.*

Proof. Assume that $P = (S, Act, T, s_0)$ is given where $S = \{A, B, C\}$. Assume further that transitions (A, τ, Δ_B) and (A, τ, Δ_C) exist. Then every $(A, \tau, c\Delta_B \oplus (1-c)\Delta_C)$ would also be a bisimilar transition for every $c \in (0, 1)$. Thus, the number of such distributions is uncountably infinite. The union of all those transitions would therefore lead to a non-finite automaton. \square

Corollary 4. *For a finite automaton P , the minimal elements in $(\Omega_{\sim}(P)/\equiv_{iso}, \subseteq)$ and $(\Omega_{\approx}(P)/\equiv_{iso}, \subseteq)$ are finite.*

Lemma 13. *The lattices $(\Omega_{\sim}(P)/\equiv_{iso}, \subseteq)$ and $(\Omega_{\approx}(P)/\equiv_{iso}, \subseteq)$ are unbounded whenever there is at least one action $a \in Act$, $a \neq \tau$.*

Proof. Let $P = (S, Act, T, s_0)$ and $a \in Act$, $a \neq \tau$. We can construct two unreachable non-bisimilar states s_1, s_2 by using the transition $s_1 \xrightarrow{a} \Delta_{s_1}$ which is not possible by s_2 (possibly $s_1 \sim s$ or $s_2 \sim s$ for some $s \in S$, then we would use the corresponding state(s) from S , not adding them by a disjunct union in the following construction). So without loss of generality we may assume that $P = (S \cup \{s_1, s_2\}, Act, T \cup \{s_1 \xrightarrow{a} \Delta_{s_1}\})$. The further construction goes in two steps. The first step is to construct a countable set of ‘fresh’ distinguished unreachable states $U = \mathbb{N}$ in the following way. use a single transition $i \xrightarrow{\tau} \frac{1}{i}\Delta_{s_0} \oplus (1 - \frac{1}{i})\Delta_{s_1}$ for every $i \in U$ (again, if $i \sim s$ for some $s \in S$, we may assume that $s \in U$, but will not add s again to the state space). So we reach the bisimilar automaton $P' = (S \dot{\cup} U, Act, T \cup \{i \xrightarrow{\tau} \frac{1}{i}\Delta_{s_0} \oplus (1 - \frac{1}{i})\Delta_{s_1}\}_{i \in U}, s_0)$. It is easy to see that all states in U are not bisimilar. For the second step notice that the powerset of a countable is uncountable. For every such subset $A \subseteq \mathbb{N}$ we may choose a distribution μ_A where $\mu_A(s) > 0$ if $s \in A$, 0 otherwise. Now, we can construct (uncountably many) PA P'_A bisimilar to P where we add one additional unreachable state s_A with the transition $s_A \xrightarrow{\tau} \mu_A$. Each of those unreachable states would need a separate state in the maximal element of our lattice (leaving alone the states that are bisimilar to one state of S by chance), meaning that there would be an uncountable number of states in the maximal element. This is a contradiction to the countability of the state space. The weak case follows similarly. \square

5 Examples

The first example is given in Fig. 1a. We start with the two strongly bisimilar automata at the top of the figure. The intersection of both is the automaton at the bottom of the figure.

The next example is given in Fig. 1b. We start with the two weakly bisimilar automata at the top of the figure. Bisimilarity essentially stems from the following facts: Firstly, the transition $(1, \tau, \Delta_3)$ of the PA on the right can be mimicked by the left automaton by a Dirac determinate scheduler that in state 1 always chooses $(1, \tau, \frac{1}{2}\Delta_2 \oplus \frac{1}{2}\Delta_3)$ and in state 2 always $(2, \tau, \Delta_3)$ and stops in state 3. Secondly, the transition $(1, \tau, \frac{1}{4}\Delta_2 \oplus \frac{3}{4}\Delta_3)$ of the left automaton can be mimicked by the right automaton by choosing each of the τ -transitions emanating from state 1 with probability $\frac{1}{2}$. The intersection of both – which is the canonical form – is the automaton at the bottom of the figure.

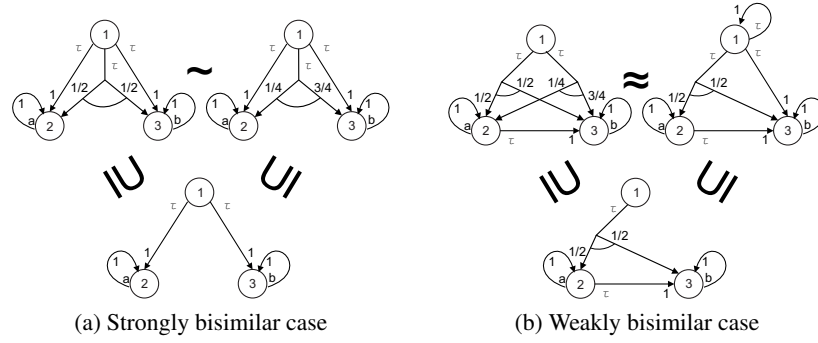


Figure 1: Intersection of bisimilar automata




Figure 2: Examples continued

The next example shows that it is essential to have elements of $\Omega_{\sim}(P)/\equiv_{iso}$, otherwise the intersection will not make sense. This example is given in Fig. 2a. It is clear that e.g. the distribution $\frac{3}{4}\Delta_2 \oplus \frac{1}{4}\Delta_3$ cannot be realised starting from state 1 by the right automaton, but it can be realised by the left automaton, so both cannot be bisimilar. Therefore, the intersection of both automata is not bisimilar to the left automaton.

5.1 A bounded example

Assume that we consider strong bisimulation and want to calculate $\Omega_{\sim}^*(P)/\equiv_{iso}$ for the automaton P given in Fig. 2b (middle). The least element (let's call it \perp) is shown on the left, the greatest element cannot be adequately depicted, as it has (uncountably) infinitely many transitions. The situation becomes clearer when (as suggested in [2]) considering distributions as points in \mathbb{R}^n . This is done in the lower line of the figure. We only show the distributions that are possible by Dirac determinate schedulers starting from state 1. As there are only two successor states, \mathbb{R}^2 suffices for our purpose. With this picture in mind it is clear that the greatest element is the PA given by $\top = (\{1, 2, 3\}, \{\tau, a, b\}, \{(2, a, \Delta_2), (3, b, \Delta_3)\} \cup \{(1, \tau, c\Delta_2 \oplus (1-c)\Delta_3) | c \in [0, 1]\}, 1)$. Clearly this is not a finite PA. Generalising from this automaton, for $\mathfrak{A} \subseteq [0, 1]$ we may define $P_{\mathfrak{A}} := (\{1, 2, 3\}, \{\tau, a, b\}, \{(2, a, \Delta_2), (3, b, \Delta_3)\} \cup \{(1, \tau, c\Delta_2 \oplus (1-c)\Delta_3) | c \in \mathfrak{A}\}, 1)$. Note that $\perp = P_{\{0,1\}}$, $\top = P_{[0,1]}$. But now it is clear how the set $\Omega_{\sim}^*(P)/\equiv_{iso}$ must look like: $\Omega_{\sim}^*(P)/\equiv_{iso} = \{P_{\mathfrak{A}} | \mathfrak{A} \subseteq [0, 1], 0 \in \mathfrak{A}, 1 \in \mathfrak{A}\}$ Leaving out 0 or 1 from \mathfrak{A} will break bisimilarity.

6 Conclusion

This paper extends the notion of normal forms introduced in [4] to the case of compact automata with countably infinite state space, countably infinite set of actions and possibly uncountably many transitions. We have justified the canonicity of normal forms by introducing them as the intersection of all bisimilar automata, not from an abstract point of view as in [4]. The structure presented is nice as a theoretical result, but (at least for the moment) there is no immediate practical applicability, as the hard part is to construct the sets $\Omega_{\sim}(P)$ and $\Omega_{\approx}(P)$ where an uncountably infinite number of automata would have to be constructed. The structure itself is particularly nice for teaching purposes and for better understanding the possible 'shapes' of infinite automata.

References

- [1] T. Brázdil, A. Kučera & O. Stražovský (2004): *Deciding Probabilistic Bisimilarity Over Infinite-State Probabilistic Systems*. In P. Gardner & N. Yoshida, editors: *CONCUR 2004, Lecture Notes in Computer Science* 3170, Springer Berlin Heidelberg, pp. 193–208, doi:10.1007/978-3-540-28644-8_13.
- [2] S. Cattani & R. Segala (2002): *Decision Algorithms for Probabilistic Bisimulation*. In: *CONCUR 2002, Lecture Notes in Computer Science* 2421, Springer, pp. 371–385, doi:10.1007/3-540-45694-5_25.
- [3] J. Desharnais, V. Gupta, R. Jagadeesan & P. Panangaden (2010): *Weak bisimulation is sound and complete for PCTL**. *Inf. Comput.* 208(2), pp. 203–219, doi:10.1016/j.ic.2009.11.002.
- [4] C. Eisentraut, H. Hermanns, J. Schuster, A. Turrini & L. Zhang (2013): *The Quest for Minimal Quotients for Probabilistic Automata*. In: *TACAS 2013, Lecture Notes in Computer Science* 7795, Springer, pp. 16–31, doi:10.1007/978-3-642-36742-7_2.
- [5] V. Forejt, P. Jancar, S. Kiefer & J. Worrell (2012): *Bisimilarity of Probabilistic Pushdown Automata*. In: *FSTTCS 2012, Leibniz International Proceedings in Informatics (LIPIcs)* 18, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, pp. 448–460, doi:10.4230/LIPIcs.FSTTCS.2012.448. Available at <http://drops.dagstuhl.de/opus/volltexte/2012/3880>.
- [6] H. Hermanns & A. Turrini (2012): *Deciding Probabilistic Automata Weak Bisimulation in Polynomial Time*. In: *FSTTCS 2012, Leibniz International Proceedings in Informatics (LIPIcs)* 18, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, pp. 435–447, doi:10.4230/LIPIcs.FSTTCS.2012.435. Available at <http://drops.dagstuhl.de/opus/volltexte/2012/3879/>.
- [7] M. Krein & D. Milman (1940): *On extreme points of regular convex sets*. *Studia Mathematica* 9(1), pp. 133–138. Available at <http://eudml.org/doc/219061>.
- [8] N. Lynch, R. Segala & F. Vaandrager (2003): *Compositionality for probabilistic automata*. In: *CONCUR 2003*, Springer, pp. 208–221, doi:10.1007/978-3-540-45187-7_14.
- [9] N. A. Lynch, R. Segala & F. W. Vaandrager (2007): *Observing Branching Structure through Probabilistic Contexts*. *SIAM J. Comput.* 37(4), pp. 977–1013, doi:10.1147/S0097539704446487.
- [10] A. Parma & R. Segala (2007): *Logical characterizations of bisimulations for discrete probabilistic systems*. In: *Foundations of Software Science and Computational Structures*, Springer, pp. 287–301, doi:10.1007/978-3-540-71389-0_21.
- [11] R. Segala (1995): *Modeling and Verification of Randomized Distributed Real-Time Systems*. Ph.D. thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology. Available at <http://profs.sci.univr.it/~segala/www/phd.html>.
- [12] R. Segala & N. A. Lynch (1995): *Probabilistic Simulations for Probabilistic Processes*. *Nord. J. Comput.* 2(2), pp. 250–273, doi:10.1007/BFb0015027.