The Quest for Minimal Quotients for Probabilistic Automata

Christian Eisentraut[1], Holger Hermanns[1], Johann Schuster[2], Andrea Turrini[1], and Lijun Zhang[3,1]

[1] Department of Computer Science, Saarland University, Germany
[2] Department of Computer Science, University of the Federal Armed Forces Munich, Germany
[3] DTU Informatics, Technical University of Denmark, Denmark

Abstract. One of the prevailing ideas in applied concurrency theory and verification is the concept of automata minimization with respect to strong or weak bisimilarity. The minimal automata can be seen as canonical representations of the behaviour modulo the bisimilarity considered. Together with congruence results wrt. process algebraic operators, this can be exploited to alleviate the notorious state space explosion problem. In this paper, we aim at identifying minimal automata and canonical representations for concurrent probabilistic models. We present minimality and canonicity results for probabilistic automata wrt. strong and weak bisimilarity, together with polynomial time minimization algorithms.

1 Introduction

Markov decision processes (*MDP*s) are models appearing in areas such as operations research, automated planning, and decision support systems. In the concurrent systems context, they arise in the form of probabilistic automata (*PA*s) [17]. *PA*s form the backbone model of successful model checkers such as PRISM [12] enabling the analysis of randomised concurrent systems. Despite the remarkable versatility of this approach, its power is limited by the state space explosion problem, and several approaches are being pursued to alleviate that problem.

In related fields, a favourable strategy is to minimize the system – or components thereof – to the quotient under bisimilarity. This can speed up the overall model analysis or turn a too large problem into a tractable one [2, 4, 9]. Both, strong and weak bisimilarity are used in practice, with weaker relations leading to greater reduction. However, this approach has never been explored in the context of *MDP*s or probabilistic automata. One reason is that thus far no effective decision algorithm was at hand for weak bisimilarity on *PA*s. A polynomial time algorithm has been proposed only recently [10]. But that algorithm is a decision algorithm, not a minimization algorithm. This paper therefore focusses on a seemingly tiny problem: Does there exist a *unique minimal* representative of a given probabilistic automaton with respect to weak bisimilarity? And can we compute it? In fact, this turns out to be an intricate problem. We nevertheless arrive at a polynomial time algorithm.

Notably, minimality with respect to the number of states of a probabilistic automaton does not imply minimality with respect to the number of transitions. And further

minimization is possible with respect to transition fanouts, the latter referring to the number of target states of a transition with non-zero probability. The minimality of an automaton thus needs to be considered with respect to all the three characteristics: number of states, of transitions and of transitions' fanouts.

We consider our results as a breakthrough with wide ranging consequences. Since weak probabilistic bisimilarity enjoys congruence properties for parallel composition and hiding on *PA*s, compositional minimization approaches can now be carried out efficiently. And because *PA*s comprise *MDP*s, we think it is not far fetched to imagine fruitful applications in areas such as operations research, automated planning, and decision support systems.

As a byproduct, our results provide tailored algorithms for strong probabilistic bisimilarity on *PA*s and strong and weak bisimilarity on labelled transition systems.

Organization. After the preliminaries in Section 2, we recall the bisimulation relations in Section 3 and we introduce the preorders between automata in Section 4. Then we present automaton reductions in Section 5 which will be used to compute the normal forms defined in Section 6. We conclude the paper in Section 7 with some remarks.

2 Preliminaries

Sets, Relations and Distributions: Given sets X, Y, and Z and relations $\mathcal{R} \subseteq X \times Y$ and $\mathcal{S} \subseteq Y \times Z$, we denote by $\mathcal{R} \circ \mathcal{S}$ the relation $\mathcal{R} \circ \mathcal{S} \subseteq X \times Z$ such that $\mathcal{R} \circ \mathcal{S} = \{ (x, z) \mid \exists y \in Y. x \mathcal{R} y \wedge y \mathcal{S} z \}$.

For a set X, we denote by $\mathrm{SubDisc}(X)$ the set of discrete sub-probability distributions over X. Given $\rho \in \mathrm{SubDisc}(X)$, we denote by $|\rho|$ the size $\rho(X) = \sum_{s \in X} \rho(s)$ of a distribution. We call a distribution ρ *full*, or simply a *probability* distribution, if $|\rho| = 1$. The set of all discrete probability distributions over X is denoted by $\mathrm{Disc}(X)$. Given $\rho \in \mathrm{SubDisc}(X)$, we denote by $\mathrm{Supp}(\rho)$ the set $\{ x \in X \mid \rho(x) > 0 \}$, by $\rho(\bot)$ the value $1 - \rho(X)$ where $\bot \notin X$, by δ_x the *Dirac* distribution such that $\rho(x) = 1$ for $x \in X \cup \{\bot\}$ where $\delta_\bot$ represents the empty distribution such that $\rho(X) = 0$. For a constant $c \geq 0$, we denote by $c\rho$ the distribution defined by $(c\rho)(x) = c \cdot \rho(x)$ if $c|\rho| \leq 1$. Further, for $\rho \in \mathrm{Disc}(X)$ and $x \in X$ such that $\rho(x) < 1$, we denote by $\rho \backslash x$ the *rescaled* distribution such that $(\rho \backslash x)(y) = \frac{\rho(y)}{1-\rho(x)}$ if $y \neq x$, 0 otherwise. We define the distribution $\rho = \rho_1 \oplus \rho_2$ by $\rho(s) = \rho_1(s) + \rho_2(s)$ provided $|\rho| \leq 1$, and conversely we say ρ can be split into ρ_1 and ρ_2. Since $\oplus$ is associative and commutative, we may use the notation $\bigoplus$ for arbitrary finite sums.

The lifting $\mathcal{L}(\mathcal{R}) \subseteq \mathrm{Disc}(X) \times \mathrm{Disc}(X)$ [13] of an equivalence relation $\mathcal{R}$ on X is defined as: for $\rho_1, \rho_2 \in \mathrm{Disc}(X)$, $\rho_1 \ \mathcal{L}(\mathcal{R}) \ \rho_2$ if and only if for each $\mathcal{C} \in X/_\mathcal{R}$, $\rho_1(\mathcal{C}) = \rho_2(\mathcal{C})$, where $X/_\mathcal{R} = \{ [x]_\mathcal{R} \mid x \in X \}$ and $[x]_\mathcal{R} = \{ x' \in X \mid x' \ \mathcal{R} \ x \}$.

Models: A probabilistic automaton (*PA*) $\mathcal{A}$ is a tuple $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$, where S is a countable set of *states*, $\bar{s} \in S$ is the *start* state, Σ is a countable set of *actions*, and $\mathcal{T} \subseteq S \times \Sigma \times \mathrm{Disc}(S)$ is a *transition relation*. In this work we consider only finite PAs, i.e., automata such that S and $\mathcal{T}$ are finite.

An example of *PA* is sketched in Figure 1(a), the precise probabilities are left unspecified, and Dirac transitions directly connect states. The set Σ is partitioned into two

sets $H = \{\tau\}$ and E of internal (hidden) and external actions, respectively; we refer to $\bar{s}$ also as the *initial* state and we let s,t,u,v, and their variants with indices range over S and a, b range over Σ.

A transition $tr = (s, a, \nu) \in \mathcal{T}$, also denoted by $s \xrightarrow{a} \nu$, is said to *leave* from state s, to be *labelled* by a, and to *lead* to ν, also denoted by ν_{tr}. We denote by $src(tr)$ the *source* state s, by $act(tr)$ the *action* a, and by $trg(tr)$ the *target* distribution ν. We also say that s enables action a, that action a is enabled from s, and that (s, a, ν) is enabled from s. Finally, we denote by $\mathcal{T}(s)$ the set of transitions enabled from s, i.e., $\mathcal{T}(s) = \{\, tr \in \mathcal{T} \mid src(tr) = s \,\}$, and similarly for $a \in \Sigma$, by $\mathcal{T}(a)$ the set of transitions with action a, i.e., $\mathcal{T}(a) = \{\, tr \in \mathcal{T} \mid act(tr) = a \,\}$.

Given a state s, an action a and a countable set of indices I, we say that there exists a *combined transition* $s \xrightarrow{a}_{\mathrm{c}} \nu$ if there exist a family of transitions $\{(s, a, \nu_i) \in \mathcal{T}\}_{i \in I}$ and a family $\{c_i \in \mathbb{R}_{\geq 0}\}_{i \in I}$ such that $\sum_{i \in I} c_i = 1$ and $\nu = \bigoplus_{i \in I} c_i \nu_i$.

We call a *PA* $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ a Labelled Transition System (*LTS*), if $(s, a, \mu) \in \mathcal{T}$ implies $\mu = \delta_t$ for some $t \in S$.

Weak Transitions: An *execution fragment* α of a *PA* $\mathcal{A}$ is a finite or infinite sequence of alternating states and actions $\alpha = s_0 a_1 s_1 a_2 s_2 \ldots$ starting from a state $first(\alpha) = s_0$ and, if the sequence is finite, ending with a state $last(\alpha)$, such that for each $i > 0$ there exists $(s_{i-1}, a_i, \nu_i) \in \mathcal{T}$ such that $\nu_i(s_i) > 0$. The *length* of α, denoted by $|\alpha|$, is the number of occurrences of actions in α. If α is infinite, then $|\alpha| = \infty$. Denote by $frags(\mathcal{A})$ the set of execution fragments of $\mathcal{A}$ and by $frags^*(\mathcal{A})$ the set of finite execution fragments of $\mathcal{A}$. An execution fragment α is a *prefix* of an execution fragment α', denoted by $\alpha \leqslant \alpha'$, if the sequence α is a prefix of the sequence α'. The *trace* of α, denoted by $trace(\alpha)$, is the sub-sequence of external actions of α; we denote by ε the empty trace. Similarly, we define $trace(a) = a$ for $a \in E$ and $trace(\tau) = \varepsilon$.

Given a *PA* $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$, the *reachable fragment* of $\mathcal{A}$ is the *PA* $RF(\mathcal{A}) = (S', \bar{s}, \Sigma, \mathcal{T}')$ where $S' = \{\, s \in S \mid \exists \alpha \in frags^*(\mathcal{A}).first(\alpha) = \bar{s} \wedge last(\alpha) = s \,\}$ and $\mathcal{T}' = \{\, (s, a, \nu) \in \mathcal{T} \mid s \in S' \,\}$.

A *scheduler* for a *PA* $\mathcal{A}$ is a function $\sigma \colon frags^*(\mathcal{A}) \to \mathrm{SubDisc}(\mathcal{T})$ such that for each finite execution fragment α, $\sigma(\alpha) \in \mathrm{SubDisc}(\mathcal{T}(last(\alpha)))$. A scheduler is *Dirac* if it assigns a Dirac distribution to each execution fragment and it is *determinate* if for each pair of execution fragments α, α', $trace(\alpha) = trace(\alpha')$ and $last(\alpha) = last(\alpha')$ imply that $\sigma(\alpha) = \sigma(\alpha')$. It is worthwhile to note that a determinate scheduler satisfies $\sigma(\alpha) = \sigma(last(\alpha))$ when $trace(\alpha) = \varepsilon$.

Given a scheduler σ and a finite execution fragment α, the distribution $\sigma(\alpha)$ describes how transitions are chosen to move on from $last(\alpha)$. A scheduler σ and a state s induce a probability distribution $\nu_{\sigma,s}$ over execution fragments as follows. The basic measurable events are the cones of finite execution fragments, where the cone of a finite execution fragment α, denoted by C_α, is the set $\{\, \alpha' \in frags(\mathcal{A}) \mid \alpha \leqslant \alpha' \,\}$. The probability $\nu_{\sigma,s}$ of a cone C_α is defined recursively as follows:

$$\nu_{\sigma,s}(C_\alpha) = \begin{cases} 0 & \text{if } \alpha = t \text{ for a state } t \neq s, \\ 1 & \text{if } \alpha = s, \\ \nu_{\sigma,s}(C_{\alpha'}) \cdot \sum_{tr \in \mathcal{T}(a)} \sigma(\alpha')(tr) \cdot \nu_{tr}(t) & \text{if } \alpha = \alpha' a t. \end{cases}$$

Standard measure theoretical arguments ensure that $\nu_{\sigma,s}$ extends uniquely to the σ-field generated by cones. We call the measure $\nu_{\sigma,s}$ a *probabilistic execution fragment* of $\mathcal{A}$ and we say that it is generated by σ from s. Given a finite execution fragment α, we define $\nu_{\sigma,s}(\alpha)$ as $\nu_{\sigma,s}(\alpha) = \nu_{\sigma,s}(C_\alpha) \cdot \sigma(\alpha)(\bot)$, where $\sigma(\alpha)(\bot)$ is the probability of choosing no transitions, i.e., of terminating the computation after α has occurred.

We say that there is a *weak combined transition* from $s \in S$ to $\nu \in \mathrm{Disc}(S)$ labelled by $a \in \Sigma$ that is induced by σ, denoted by $s \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \nu$, if there exists a scheduler σ such that the following holds for the induced probabilistic execution fragment $\nu_{\sigma,s}$:

1. $\nu_{\sigma,s}(\mathit{frags}^*(\mathcal{A})) = 1$;
2. for each $\alpha \in \mathit{frags}^*(\mathcal{A})$, if $\nu_{\sigma,s}(\alpha) > 0$ then $\mathit{trace}(\alpha) = \mathit{trace}(a)$;
3. for each state t, $\nu_{\sigma,s}(\{\, \alpha \in \mathit{frags}^*(\mathcal{A}) \mid \mathit{last}(\alpha) = t \,\}) = \nu(t)$.

We say that there is a *weak transition* from $s \in S$ to $\nu \in \mathrm{Disc}(S)$ labelled by $a \in \Sigma$ that is induced by σ, denoted by $s \stackrel{a}{\Longrightarrow} \nu$, if there exists a Dirac scheduler σ inducing $s \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \nu$.
We say that there is a *weak hyper transition* from $\rho \in \mathrm{Disc}(S)$ to $\nu \in \mathrm{Disc}(S)$ labelled by $a \in \Sigma$, denoted by $\rho \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \nu$, if there exists a family of weak combined transitions $\{s \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \nu_s\}_{s \in \mathrm{Supp}(\rho)}$ such that $\nu = \bigoplus_{s \in \mathrm{Supp}(\rho)} \rho(s) \cdot \nu_s$.

Given two weak hyper transitions, it is known that their concatenation is still a weak hyper transition, provided that one of the two weak hyper transitions is labelled by τ.

Lemma 1 (cf. [14, Prop. 3.6]). *Given a PA $\mathcal{A}$ and an action a, if there exist two weak hyper transitions $\nu_1 \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \nu_2$ and $\nu_2 \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_3$ (or $\nu_1 \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_2$ and $\nu_2 \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \nu_3$), then there exists the weak hyper transition $\nu_1 \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \nu_3$.*

In the remainder of the paper we make use of this lemma without mentioning it further. The following technical lemma allows us to decompose a weak hyper transition $\mu \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \mu'$ into several weak hyper transitions $\mu_i \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \mu_i'$. This can be seen as an extension of the family of weak combined transitions to a family of generic weak hyper transitions.

Lemma 2 (cf. [7, Lemmas 5 and 6]). *Let $\mu, \mu' \in \mathrm{Disc}(S)$ and $k \in \mathbb{N}$. $\mu \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \mu'$ iff $\mu = \mu_1 \oplus \cdots \oplus \mu_k$ for subdistributions $\mu_1, \ldots, \mu_k$ and for each $i = 1, \ldots, k$ a distribution μ_i' exists, such that $\mu_i \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \mu_i'$ and $\mu' = \bigoplus_{i=1,\ldots,k} \mu_i'$.*

We will often lift mappings defined on a set of states S to mappings over distributions $\mathrm{Disc}(S)$ in a generic way.

Definition 1 (Lifting of Functions). *Given arbitrary sets S and M, and $\mu \in \mathrm{Disc}(S)$, we* lift *a mapping $b\colon S \to M$ to $b\colon \mathrm{Disc}(S) \to \mathrm{Disc}(M)$ by defining $(b(\mu))(m) = \sum_{s \in b^{-1}(m)} \mu(s)$ for each $m \in M$.*

3 Bisimulations

In the following, we define strong and weak (probabilistic) bisimulations. Let $\rightsquigarrow \in \{\longrightarrow, \longrightarrow_{\mathrm{c}}, \Longrightarrow, \Longrightarrow_{\mathrm{c}}\}$.

Definition 2 (Generic Bisimulation). *Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ be a PA. An equivalence relation $\mathcal{R} \subseteq S \times S$ is a $\rightsquigarrow$-bisimulation if for every action $a \in \Sigma$, distribution $\mu \in \mathrm{Disc}(S)$, and states $s, s' \in S$, with $s \ \mathcal{R} \ s'$ it holds that $s \xrightarrow{a} \mu$ implies $s' \overset{a}{\rightsquigarrow} \gamma$ for some γ and $\mu \ \mathcal{L}(\mathcal{R}) \ \gamma$.*

We denote by $\asymp_{\rightsquigarrow}$ the union of all $\rightsquigarrow$-bisimulations. Two *PA*s $\mathcal{A}$, $\mathcal{A}'$ are $\rightsquigarrow$-bisimilar, written $\mathcal{A} \asymp_{\rightsquigarrow} \mathcal{A}'$ if their initial states are bisimilar in the direct sum of the two automata. We recover the standard characterization for strong and weak bisimilarities from this definition as follows:

1. Strong Bisimilarity for *LTS*, denoted $\sim_S$, is $\asymp_{\longrightarrow}$.
2. Strong Probabilistic Bisimilarity for *PA*, denoted $\sim$, is $\asymp_{\longrightarrow_c}$.
3. Weak Bisimilarity for *LTS*, denoted $\approx_S$, is $\asymp_{\Longrightarrow}$.
4. Weak Probabilistic Bisimilarity for *PA*, denoted $\approx$, is $\asymp_{\Longrightarrow_c}$.

For the rest of the paper, we let the symbol $\asymp$ range over $\{\sim, \sim_S, \approx, \approx_S\}$. The relations $\sim_S$ and $\approx_S$ coincide with the respective notions of strong and weak bisimilarity on *LTS* [15]. The same holds for the probabilistic bisimilarities $\sim$ and $\approx$ on *PA*s [18]. In the sequel we assume that bisimilarities are only applied to suitable automata, for example, if we write $\mathcal{A} \sim_S \mathcal{A}'$, we implicitly assume $\mathcal{A}, \mathcal{A}' \in$ *LTS*.

4 Preorders

The size of an automaton is usually expressed in terms of the size of the set of states $|S|$ and the size of the transition relation $|\mathcal{T}|$ of the automaton. The complexity of algorithms working on probabilistic automata often depends exactly on those two metrics. A less commonly considered metric is the number of target states of a transition reached with a probability greater than zero. Especially in practical applications it is known that the first two of these metrics – the number of states and transitions of an automaton – can be drastically reduced while preserving its behaviour wrt. some notion of bisimilarity. In contrast, the last metric is usually considered a constant, but in some cases it can be reduced as well. We will formalize these three metrics by means of three preorder relations, thus allowing us to define the notion of *minimal automata* up to bisimilarity.

To capture the last of the three metrics, we introduce the following definition.

Definition 3 (Transition Fanout). *For a distribution $\mu \in Dist(S)$ we define $\|\mu\| = |Supp(\mu)|$. For a set of transitions T we define $\|T\| = \sum_{(s,a,\mu)\in T} \|\mu\|$.*

Definition 4 (Size Preorders). *Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ and $\mathcal{A}' = (S', \bar{s}', \Sigma', \mathcal{T}')$ be two PAs, and let $\asymp$ be a notion of bisimilarity. We write*

- *$\mathcal{A} \overset{\preceq^{|S|}}{\asymp} \mathcal{A}'$ if $\mathcal{A} \asymp \mathcal{A}'$ and $|S| \leq |S'|$,*
- *$\mathcal{A} \overset{\preceq^{|\mathcal{T}|}}{\asymp} \mathcal{A}'$ if $\mathcal{A} \asymp \mathcal{A}'$ and $|\mathcal{T}| \leq |\mathcal{T}'|$, and*
- *$\mathcal{A} \overset{\preceq^{\|\mathcal{T}\|}}{\asymp} \mathcal{A}'$ if $\mathcal{A} \asymp \mathcal{A}'$ and $\|\mathcal{T}\| \leq \|\mathcal{T}'\|$.*

We let from now on $\preceq$ range over $\overset{\preceq^{|S|}}{\asymp}$, $\overset{\preceq^{|\mathcal{T}|}}{\asymp}$, and $\overset{\preceq^{\|\mathcal{T}\|}}{\asymp}$ for $\asymp \in \{\sim, \sim_S, \approx, \approx_S\}$, if not mentioned otherwise. It is easy to verify that these relations are preorders.

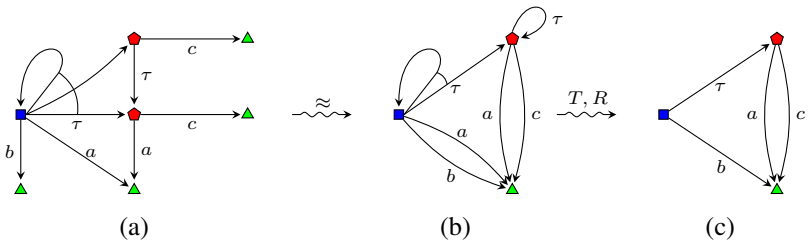

Fig. 1. (a) Example PA, (b) Quotient reduction. (c) Rescaling of convex-transitive reduction.

Definition 5 ($\preceq$-Minimal Automata). *We call a PA $\mathcal{A}$* $\preceq$-minimal*, if whenever $\mathcal{A}' \preceq \mathcal{A}$ for some PA $\mathcal{A}'$, then also $\mathcal{A} \preceq \mathcal{A}'$.*

Lemma 3 (Existence of $\preceq$-Minimal Automata). *For every PA $\mathcal{A}$ there exists a PA $\mathcal{A}'$ such that $\mathcal{A}' \asymp \mathcal{A}$ and $\mathcal{A}'$ is $\preceq$-minimal.*

For each of the preorders considered, the proof of this lemma exploits that for every automaton the respective metric is a finite natural number and at least 0.

As each relation $\preceq$ is a preorder, minimal automata are not necessarily unique. For example, two $\overset{\sim}{\preceq}^{|S|}$-minimal automata $\mathcal{A}$ and $\mathcal{A}'$ with $\mathcal{A} \asymp \mathcal{A}'$ may differ in the underlying set of states, and/or transitions. This will be investigated in Section 6.

5 Reductions

In this section, we introduce and formalize several methods to reduce the size of an automaton. Except for the first method, quotient reduction, the methods are especially tailored towards one or two distinct notions of bisimilarity. We summarize the properties of the reductions at the end of this section. We will further show that each reduction can be computed in polynomial time.

5.1 Quotient Reduction

Definition 6 (Quotient Automaton). *Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ be a PA and $\mathcal{P}(S) = \{\, C \mid C \subseteq S \,\}$. Given an equivalence relation $\asymp$ on S, we define the* quotient PA $[\mathcal{A}]_{\asymp}$ with respect to $\asymp$ *as the reachable fragment of the PA $(S/_{\asymp}, [\bar{s}]_{\asymp}, \Sigma, [\mathcal{T}]_{\asymp})$ where (i) the equivalence class mapping $[\,\cdot\,]_{\asymp} : S \to \mathcal{P}(S)$ is defined for every $s \in S$ as $[s]_{\asymp} = \{\, s' \mid s' \asymp s \,\}$, (ii) $S/_{\asymp} = \{\, [s]_{\asymp} \mid s \in S \,\}$, and (iii) $[\mathcal{T}]_{\asymp} = \{\, ([s]_{\asymp}, a, [\mu]_{\asymp}) \mid (s, a, \mu) \in \mathcal{T} \,\}$.*

Note that $[\mu]_{\asymp}$ means lifting the quotient mapping on states $[\cdot]_{\asymp}$ to distributions according to Definition 1.

Definition 7 (Quotient Reduction). *We write $\mathcal{A} \overset{\asymp}{\hookrightarrow} \mathcal{A}'$ if $\mathcal{A}' = [\mathcal{A}]_{\asymp}$.*

Fig. 1(b) shows the result of applying Def. 7 to weak bisimilarity and the PA in Fig. 1(a).

5.2 Convex Reduction

In essence, strong probabilistic bisimilarity $\sim$ enhances standard bisimilarity by the idea that the observable behaviour of a system is closed under convex combinations of transitions. Using this fact, we minimize the number of transitions in a *PA* by replacing the transitions of each state by a unique and *minimal* set of generating transitions.

Definition 8. *Let $P = \{p_1, \ldots, p_n \in \mathbb{R}^k\}$ be a finite set of points in $\mathbb{R}^k$. We call $\mathit{CHull}(P) = \{\, c \in \mathbb{R}^k \mid \exists c_1, \ldots, c_n \geq 0 : \sum_{i=1}^n c_i = 1$ and $c = \sum_{i=1}^n c_i \cdot p_i \,\}$ the convex hull of P.*

C is a *finitely generated convex set*, if $C = \mathit{CHull}(P)$ for a finite set $P \subseteq \mathbb{R}^k$. The following lemma guarantees the optimality of our approach with respect to $\precsim^{|\mathcal{T}|}$.

Lemma 4 (cf. [3, Sec. 2]). *Every finitely generated convex set C has a unique minimal set of generators $\mathit{Gen}(C)$ such that $C = \mathit{CHull}(\mathit{Gen}(C))$.*

Definition 9 (Convex Reduction). *Let $\mathcal{A}$ be a PA. We write $\mathcal{A} \overset{C}{\rightsquigarrow} \mathcal{A}'$ if the automaton $\mathcal{A}'$ differs from $\mathcal{A}$ only by replacing the set $\mathcal{T}$ by the set $\mathcal{T}'$, where*

$$(s, a, \gamma) \in \mathcal{T}' \text{ if and only if } \gamma \in \mathit{Gen}(\mathit{CHull}(\{\, \mu \mid (s, a, \mu) \in \mathcal{T} \,\})).$$

5.3 Convex-Transitive Reduction

Just like strong probabilistic bisimilarity, weak probabilistic bisimilarity embodies the idea that the observable behaviour of a system is closed under convex combinations. Yet, this has to be interpreted for weak transitions. Finding a minimal set of generators turns out to be harder in this setting, as the behaviour of each state s no longer only depends on (convex combinations of) single step transitions leaving s. Instead, reachable distributions are now characterized by arbitrarily complex schedulers and their convex combinations. This convex set is known to be finitely generated [3].

We take inspiration from the standard approach followed in transitive reduction of order relations. Intuitively, this is the opposite of the transitive closure operation, and is achieved by removing transitions that can be reconstructed from other transitions by transitivity. In this spirit, we propose a simple algorithm that iteratively removes transitions, as long as their target distribution can also be reached by a weak combination of other transitions. Similar to transitive reduction on order relations, this reduction algorithm has polynomial complexity.

We will later show that this reduction leads to a minimal result with respect to $\precapprox^{|\mathcal{T}|}$, if applied on a model that a priori has been subjected to a quotient reduction.

Definition 10 (Convex-Transition Reduction Preorder).
Given the PAs $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ and $\mathcal{A}' = (S', \bar{s}', \Sigma', \mathcal{T}')$, we write $\mathcal{A} \subseteq_{\mathcal{T}} \mathcal{A}'$ if and only if $\mathcal{T} \subseteq \mathcal{T}'$, $S = S'$, $\Sigma = \Sigma'$, $\bar{s} = \bar{s}'$, and for each transition $(s, a, \mu) \in \mathcal{T}'$ there exists a weak combined transition $s \overset{a}{\Longrightarrow}_{\mathrm{c}} \mu$ in $\mathcal{A}$.

Definition 11 ($\subseteq_{\mathcal{T}}$-Minimal Automata). *We call a PA $\mathcal{A}$ $\subseteq_{\mathcal{T}}$-*minimal*, if whenever $\mathcal{A}' \subseteq_{\mathcal{T}} \mathcal{A}$ for some PA $\mathcal{A}'$, then also $\mathcal{A} \subseteq_{\mathcal{T}} \mathcal{A}'$.*

Lemma 5 (Existence of $\subseteq_{\mathcal{T}}$-Minimal Automata). *For every PA $\mathcal{A}$ there exists a PA $\mathcal{A}'$ such that $\mathcal{A}' \approx \mathcal{A}$ and $\mathcal{A}'$ is $\subseteq_{\mathcal{T}}$-minimal.*

Definition 12 (Convex Transitive Reduction). *Let $\mathcal{A}$ be a PA. We write $\mathcal{A} \overset{T}{\rightsquigarrow} \mathcal{A}'$ if $\mathcal{A}' \subseteq_{\mathcal{T}} \mathcal{A}$ and $\mathcal{A}'$ is $\subseteq_{\mathcal{T}}$-minimal.*

Notably, this reduction relation is non-deterministic, i.e., non-functional. But, as we will show in Section 6, it is unique up to isomorphism ($=_{iso}$), if applied to a quotient reduced automaton. The overall result will therefore be unique up to isomorphism. As a special case, this reduction can be applied to non-probabilistic transition systems (*LTS*s), where it then coincides with transitive reduction of order relations. For this it is irrelevant that this reduction allows to combine transitions, as long as we work on a quotient reduced system, because in that system bisimilar states have been collapsed into a single representative. Thus, a Dirac transition to a single state can only be matched by a Dirac transition to precisely that state. In the *LTS* setting, $\overset{T}{\rightsquigarrow}$ preserves $\approx_S$, and in fact is a necessary step to arrive at the transition minimal quotient. As a side note, though this must have been considered in the context of tools exploiting weak bisimilarity [5, 8], we are not aware of a publication mentioning this point.

5.4 Rescaling

A particular fine point of weak probabilistic bisimilarities [1] is related to internal transitions that induce a nonzero chance of residing inside the class. If looking at the quotient, this concerns any internal transition (s, τ, μ) that contains the source state s with nontrivial probability, i.e., $0 < \mu(s) < 1$. For those transitions, we can renormalise the probability of all other states in the support set by $1 - \mu(s)$ without breaking weak bisimilarity. In other words, such μ can be replaced by the rescaled distribution $\mu \backslash s$.

Definition 13 (Rescaling). *Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ be a PA. We write $\mathcal{A} \overset{R}{\rightsquigarrow} \mathcal{A}'$ if $\mathcal{A}' = (S, \bar{s}, \Sigma, \mathcal{T}')$ such that for each $(s, a, \mu') \in \mathcal{T}'$, either $a \in E$ and $(s, a, \mu') \in \mathcal{T}$, or $a \in H$ and there exists $(s, \tau, \mu) \in \mathcal{T}$ such that $\mu(s) < 1$ and $\mu' = \mu \backslash s$.*

As it will turn out, this reduction is the final step to obtain minimality with respect to $\overset{\|\mathcal{T}\|}{\approx}$ if applied a posteriori to the other two reductions, $\overset{\approx}{\rightsquigarrow}$ and $\overset{T}{\rightsquigarrow}$. Figure 1(c) depicts the result of applying this sequence of reductions on the *PA* in Figure 1(a). Figure 1(b) shows the automaton after it has been subjected to quotient reduction only.

5.5 Properties of Reductions

We summarize preservation and computability properties of the reduction relations.

Lemma 6 (Preservation of Bisimilarities).

1. *$\mathcal{A} \overset{\asymp}{\rightsquigarrow} \mathcal{A}'$ implies $\mathcal{A} \asymp \mathcal{A}'$ for each $\mathcal{A}, \mathcal{A}'$ and $\asymp \in \{\sim, \sim_S, \approx, \approx_S\}$.*
2. *$\mathcal{A} \overset{C}{\rightsquigarrow} \mathcal{A}'$ implies $\mathcal{A} \sim \mathcal{A}'$ for each $\mathcal{A}, \mathcal{A}' \in$ PA.*
3. *$\mathcal{A} \overset{T}{\rightsquigarrow} \mathcal{A}'$ implies $\mathcal{A} \asymp \mathcal{A}'$ for each $\mathcal{A}, \mathcal{A}'$ and $\asymp \in \{\approx_S, \approx\}$.*

4. $\mathcal{A} \overset{R}{\rightsquigarrow} \mathcal{A}'$ *implies* $\mathcal{A} \approx \mathcal{A}'$ *for each* $\mathcal{A}, \mathcal{A}' \in$ PA.

Proof. **Proof for $\overset{\asymp}{\rightsquigarrow}$, $\overset{C}{\rightsquigarrow}$ and $\overset{T}{\rightsquigarrow}$:** The result follows almost immediately from the definitions of the reductions.
Proof for $\overset{R}{\rightsquigarrow}$: Since by definition of $\overset{R}{\rightsquigarrow}$, $\mathcal{A}$ and $\mathcal{A}'$ have the same set of states, we use ν to refer to distributions from both $\mathcal{A}$ and $\mathcal{A}'$; we still use s' to remark that we consider the state s in $\mathcal{A}'$.

Let $\mathcal{I}$ be the equivalence relation on $S \uplus S'$ whose set of classes is $\{ \{s, s'\} \mid s \in S \}$, i.e., we relate each state s with its primed counterpart in $\mathcal{A}'$. $\mathcal{I}$ is a weak probabilistic bisimulation for $\mathcal{A}$ and $\mathcal{A}'$: let $s \ \mathcal{I} \ t$ and $s \xrightarrow{a} \nu$; if $s = t$, then also t enables the transition $t \xrightarrow{a} \nu$ and $\nu \ \mathcal{L}(\mathcal{I}) \ \nu$. Suppose that $s \neq t$; if $a \in E$, then also t enables the transition $t \xrightarrow{a} \nu$, thus $\nu \ \mathcal{L}(\mathcal{I}) \ \nu$. Now, consider $a \in H$: if $s \in S$ and $t \in S'$, i.e., $t = s'$, then t is able to match such transition by the weak combined transition $t \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu$ as induced by the scheduler σ such that $\sigma(t)(\bot) = \nu(s)$, $\sigma(t)(tr) = 1 - \nu(s)$, and $\sigma(\alpha)(\bot) = 1$ for each finite execution fragment $\alpha \neq t$, where $tr = (t, \tau, \nu \backslash s)$. Note that this applies also when $\nu = \delta_s$ as the resulting scheduler assigns $\sigma(t)(\bot) = \nu(s) = 1$ so the induced weak combined transition is $t \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_t$ and $\delta_s \ \mathcal{L}(\mathcal{I}) \ \delta_t$. Otherwise, if $s \in S'$ and $t \in S$, i.e., $s = t'$, then $s \xrightarrow{a} \nu$ is actually a transition $s \xrightarrow{a} \rho \backslash s$ that t is able to match by the weak combined transition $t \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu$ as induced by the determinate scheduler σ such that $\sigma(\alpha)(tr') = 1$ for each $\alpha \in \mathit{frags}^*(\mathcal{A})$ with $\mathit{last}(\alpha) = t$, and $\sigma(\alpha)(\bot) = 1$ for each finite execution fragment α with $\mathit{last}(\alpha) \neq t$ where $tr' = (t, \tau, \rho)$. □

Proposition 1 (Computability of Reductions). *For every PA* $\mathcal{A}$*, a PA* $\mathcal{A}'$ *can be found in polynomial time, such that* $\mathcal{A} \rightsquigarrow \mathcal{A}'$ *for* $\rightsquigarrow \in \{\overset{\asymp}{\rightsquigarrow}, \overset{C}{\rightsquigarrow}, \overset{T}{\rightsquigarrow}, \overset{R}{\rightsquigarrow}\}$.

Proof (outline). The result for $\overset{\asymp}{\rightsquigarrow}$ follows by the corresponding polynomial decision procedures [3, 8, 10, 11, 16] and reachability analysis; $\overset{C}{\rightsquigarrow}$ requires for each state and each enabled action to solve $\mathcal{O}(|\mathcal{T}|)$ linear programming problems (cf. [3, Sec. 6]) in order to find the set of generators of reachable distributions; $\overset{R}{\rightsquigarrow}$ can be obtained by computing for each transition $s \xrightarrow{\tau} \nu$ the distribution $\nu \backslash s$ that requires at most $\mathcal{O}(|S|)$ operations; finally, $\overset{T}{\rightsquigarrow}$ can be computed by iteratively refining $\mathcal{A}$ by removing one transition obtaining $\mathcal{A}'$ and deciding whether $\mathcal{A} \approx \mathcal{A}'$. Since this is polynomial [10] and the check is performed at most $|\mathcal{T}|$ times, computing $\overset{T}{\rightsquigarrow}$ is polynomial. □

6 Normal Forms

We are now concerned with minimality and uniqueness properties induced by the reduction operations with respect to the metrics discussed. To discuss uniqueness, it is convenient to introduce normal forms as means to canonically represent automata in such a way that two automata are equivalent if and only if they have identical normal forms. Or better, if and only if the normal forms are identical up to isomorphism (structural identity). Two PAs $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ and $\mathcal{A}' = (S', \bar{s}', \Sigma', \mathcal{T}')$ are *isomorphic*, denoted by $\mathcal{A} =_{iso} \mathcal{A}'$, if $\Sigma = \Sigma'$ and there is a bijective mapping $b \colon S \to S'$ such that $b(\bar{s}) = \bar{s}'$ and $(s, a, \mu) \in \mathcal{T}$ if and only if $(b(s), a, b(\mu)) \in \mathcal{T}'$.

Definition 14 (Normal Form). *Given an equivalence relation $\asymp$ over PAs, we call NF: PA $\to$ PA a* normal form*, if it satisfies for every PA $\mathcal{A}$*

- *NF($\mathcal{A}$) $\asymp \mathcal{A}$, and*
- *for every PA $\mathcal{A}'$ it holds that $\mathcal{A} \asymp \mathcal{A}'$ if and only if NF($\mathcal{A}$) $=_{iso}$ NF($\mathcal{A}'$).*

It is natural to strive for normal forms that are distinguished in a certain sense. Not surprisingly, we will strive for normal forms that are distinguished as being the smallest possible representation of the behaviour they represent. In the following, we call a total and functional subset of a binary relation $r \subseteq PA \times PA$ a *function in* r. Note that every function in r is a mapping $PA \to PA$.

Definition 15 (Normal Form Instances).

- *Let $NF_{\sim_S} = \overset{\sim}{\underset{S}{\hookrightarrow}}$.*
- *Let $NF_{\approx_S}$ be an arbitrary function in $\overset{\approx}{\underset{S}{\hookrightarrow}} \circ \overset{T}{\hookrightarrow}$.*
- *Let $NF_{\sim} = \overset{\sim}{\hookrightarrow} \circ \overset{C}{\hookrightarrow}$.*
- *Let $NF_{\approx}$ be an arbitrary function in $\overset{\approx}{\hookrightarrow} \circ \overset{T}{\hookrightarrow} \circ \overset{R}{\hookrightarrow}$.*

Theorem 1. *Let $\asymp \in \{\sim, \sim_S, \approx, \approx_S\}$.*

1. ***Minimality:** $NF_{\asymp}(\mathcal{A})$ is $\underset{\asymp}{\preceq}^{|S|}$, $\underset{\asymp}{\preceq}^{|\mathcal{T}|}$, and $\underset{\asymp}{\preceq}^{\|\mathcal{T}\|}$-minimal for every $\mathcal{A} \in PA$.*
2. ***Uniqueness of minimals:** If $\mathcal{A}$ and $\mathcal{A}'$ are $\underset{\asymp}{\preceq}^{|S|}$, $\underset{\asymp}{\preceq}^{|\mathcal{T}|}$, and $\underset{\asymp}{\preceq}^{\|\mathcal{T}\|}$-minimal automata, and $\mathcal{A} \asymp \mathcal{A}'$, then also $\mathcal{A} =_{iso} \mathcal{A}'$,*
3. ***Normal forms:** $NF_{\asymp}$ is uniquely defined up to $=_{iso}$, and is a normal form.*

It is straightforward to check that all normal forms $NF_{\asymp}$ above are indeed mappings. Furthermore, by Lemma 6, it follows that in each of the cases $NF_{\asymp}(\mathcal{A}) \asymp \mathcal{A}$.

The remainder of this section is devoted to the proof of Theorem 1. We begin with a lemma that we use often.

Lemma 7 (Preservation of Minimality). *Let $\preceq \in \{\underset{\asymp}{\preceq}^{|S|}, \underset{\asymp}{\preceq}^{|\mathcal{T}|}, \underset{\asymp}{\preceq}^{\|\mathcal{T}\|}, \subseteq_{\mathcal{T}}\}$. If $\mathcal{A} =_{iso} \mathcal{A}'$ and $\mathcal{A}$ is $\preceq$-minimal, then $\mathcal{A}'$ is $\preceq$-minimal, too.*

For each normal form, the proof will refer to the following crucial, but already folklore insight, that the quotient automaton is minimal with respect to the number of states.

Lemma 8 (State Minimality of Quotient Automata). *For every $\mathcal{A} \in PA$, the automaton $\mathcal{A}'$ with $\mathcal{A} \overset{\asymp}{\hookrightarrow} \mathcal{A}'$ is $\underset{\asymp}{\preceq}^{|S|}$-minimal.*

Next, we show that $\underset{\asymp}{\preceq}^{|S|}$and $\underset{\asymp}{\preceq}^{|\mathcal{T}|}$-minimality can be achieved at the same time in one automaton. For bisimilarities on *LTS*s, this is enough to conclude also $\underset{\asymp}{\preceq}^{\|\mathcal{T}\|}$-minimality, as this always coincides with $\underset{\asymp}{\preceq}^{|\mathcal{T}|}$-minimality here (as all transition have the form (s, a, δ_t)).

Lemma 9 (Compatibility of $\underset{\asymp}{\preceq}^{|S|}$and $\underset{\asymp}{\preceq}^{|\mathcal{T}|}$-minimality). *For every PA $\mathcal{A}$ there exists a PA $\mathcal{A}'$ with $\mathcal{A}' \asymp \mathcal{A}$, which is $\underset{\asymp}{\preceq}^{|S|}$and $\underset{\asymp}{\preceq}^{|\mathcal{T}|}$-minimal.*

Proof. By Lemma 3, there exists a PA $\mathcal{A}$ that is $\underset{\asymp}{\preceq}^{|\mathcal{T}|}$-minimal. Consider $[\mathcal{A}]_{\asymp}$. From Definition 6 it is clear that for every transition of $[\mathcal{A}]_{\asymp}$ there exists a transition in $\mathcal{A}$. Thus, $[\mathcal{A}]_{\asymp} \underset{\asymp}{\preceq}^{|\mathcal{T}|} \mathcal{A}$, and hence, $[\mathcal{A}]_{\asymp}$ must also be $\underset{\asymp}{\preceq}^{|\mathcal{T}|}$-minimal. Furthermore, by Lemma 8, $[\mathcal{A}]_{\asymp}$ must also be $\underset{\asymp}{\preceq}^{|S|}$-minimal, and finally with Lemma 6 $\mathcal{A} \asymp \mathcal{A}'$ follows. □

Strong Bisimilarities

Lemma 10 (Canonicity of Normal Form). *Let $\asymp \in \{\sim_S, \sim\}$, $\mathcal{A} \in$ PA, and $\mathcal{A}' = NF_{\asymp}(\mathcal{A})$. For every $\precsim^{|S|}_{\asymp}$ and $\precsim^{|\mathcal{T}|}_{\asymp}$-minimal PA $\mathcal{A}_m$ with $\mathcal{A}_m \asymp \mathcal{A}$, also $\mathcal{A}_m =_{iso} \mathcal{A}'$.*

Proof. We skip the proof for $\asymp = \sim_S$ and proceed with the more complicated case of $\asymp = \sim$. Let $\mathcal{A}_m$ be a $\precsim^{|S|}_{\asymp}$and $\precsim^{|\mathcal{T}|}_{\asymp}$-minimal automaton. Recall that $NF_{\sim} = \overset{\sim}{\rightsquigarrow} \circ \overset{C}{\rightsquigarrow}$. As applying $\overset{\sim}{\rightsquigarrow}$ to $\mathcal{A}$ leads to a $\precsim^{|S|}_{\sim}$-minimal automaton according to Lemma 8, and $\overset{C}{\rightsquigarrow}$ obviously does not alter the number of states, $NF_{\sim}(\mathcal{A}) = \mathcal{A}'$ is $\precsim^{|S|}_{\sim}$-minimal, and thus $|S_m| = |S'|$, as $\mathcal{A}_m$ is $\precsim^{|S|}_{\sim}$-minimal by assumption.

Since $\mathcal{A}' \sim \mathcal{A}$ and $\mathcal{A} \sim \mathcal{A}_m$, we have $\mathcal{A}' \sim \mathcal{A}_m$. We will now argue that $b = \sim \cap (S' \times S_m)$ is in fact a suitable mapping to establish $\mathcal{A}' =_{iso} \mathcal{A}_m$. We start by showing that b is functional, injective and surjective. Assume b is not injective. Then there must exist states $s_1, s_2 \in S'$ and $t \in S_m$, such that $b(s_1) = t$ and $b(s_2) = t$. But this implies $s_1 \sim t$ and $s_2 \sim t$. By transitivity, this implies $s_1 \sim s_2$, contradicting Lemma 8. Functionality can be shown similarly. We skip the details. If b is not surjective, this would immediately contradict the assumption that $\mathcal{A}_m$ is $\precsim^{|S|}_{\sim}$-minimal, since then any state $t \in \mathcal{A}_m$ for which no $s \in S'$ exists, such that $b(s) = t$ could be removed without violating $\mathcal{A}' \sim \mathcal{A}_m$.

Most of the other conditions that have to be checked to show that b is an isomorphism are straightforward, except for the condition

$$(s, a, \mu) \in \mathcal{T} \quad \text{if and only if} \quad (b(s), a, b(\mu)) \in \mathcal{T}'. \tag{$\star$}$$

The set of combined transitions any state s of $\mathcal{A}'$ can do must equal the set of combined transitions that $b(s)$ can do as $s \sim b(s)$. By reduction $\overset{C}{\rightsquigarrow}$, the set of transitions leaving s must be minimal, according to Lemma 4, and must also be unique. As the transitions of $b(s)$ are minimal by assumption, the uniqueness of the minimal set of generators guarantees Condition ($\star$). □

For $\sim_S$ and $\sim$, Theorem 1 now follows almost immediately by Lemma 9, Lemma 10 and Lemma 6. For $\sim_S$, we in addition need the observation that $\mathcal{A}$ is $\precsim^{\|\mathcal{T}\|}_{\asymp}$-minimal if and only if it is $\precsim^{|\mathcal{T}|}_{\asymp}$-minimal, as we remarked before Lemma 9. For $\sim$, the same observation holds, but follows from the uniqueness of the minimal set of generators (Lemma 4).

Weak Bisimilarities The following two lemmas are the weak counterparts to Lemma 10.

Lemma 11. *Let $\mathcal{A}$ be a PA and $\mathcal{A}' = NF_{\approx_S}(\mathcal{A})$. Let $\mathcal{A}_m$ be a $\precsim^{|S|}_{\approx_S}$and $\precsim^{|\mathcal{T}|}_{\approx_S}$-minimal PA satisfying $\mathcal{A}_m \approx_S \mathcal{A}$. Then $\mathcal{A}' =_{iso} \mathcal{A}_m$.*

We skip the proof of this lemma, as it is similar to, but simpler than the proof of the following lemma. Theorem 1 can then be proven in complete analogy to the proof for $\sim_S$.

It is instructive to note that in the following lemma, we need to apply the reduction $\overset{R}{\rightsquigarrow}$ to arrive at an uniqueness result. Only applying $\overset{\approx}{\rightsquigarrow}$ followed by $\overset{T}{\rightsquigarrow}$ will still lead to $\precsim^{|S|}_{\asymp}$and $\precsim^{|\mathcal{T}|}_{\asymp}$-minimal automata, but they will not agree up to $=_{iso}$, in full generality. Different to Lemmas 11 and 10, the following lemma is slightly more general.

Lemma 12. *Let $\mathcal{A}$ be a $\precsim^{|S|}_{\approx}$-minimal PA, $\mathcal{A} \overset{T}{\rightsquigarrow} \circ \overset{R}{\rightsquigarrow} \mathcal{A}'$, and $\mathcal{A}'_m$ be a $\precsim^{|S|}_{\approx}$and $\precsim^{|\mathcal{T}|}_{\approx}$-minimal PA satisfying $\mathcal{A}'_m \approx \mathcal{A}$. Now let $\mathcal{A}'_m \overset{R}{\rightsquigarrow} \mathcal{A}_m$ for some $\mathcal{A}_m$. Then $\mathcal{A}' =_{iso} \mathcal{A}_m$.*

Proof. Let $\mathcal{A}_m$ and $\mathcal{A}'$ be chosen as in the claim. We then proceed similarly as in the proof of Lemma 10 to show that $b = \approx \cap (S_m \times S')$ is a bijection. Then we will be able to establish that b is a suitable mapping to establish $\mathcal{A}_m =_{iso} \mathcal{A}'$.

Assume, to derive a contradiction, that b is not an isomorphism. Since b is a bijection between S_m and S' (note that all automata in this lemma must be $\precsim^{|S|}_{\approx}$-minimal), in order to have $\mathcal{A}_m \neq_{iso} \mathcal{A}'$ there must exist $s \in S_m$, $t \in S'$ with $s \approx t$ (i.e., $b(s) = t$), and (i) either a transition $s \xrightarrow{a}\!\!\!\bullet\; \nu_s \in \mathcal{T}_m$ but there does not exist $t \xrightarrow{a}\!\!\!\bullet\; \nu_t \in \mathcal{T}'$ such that $\nu_s \; \mathcal{L}(\approx) \; \nu_t$, i.e., there does not exist a transition $t \xrightarrow{a}\!\!\!\bullet\; \nu_t \in \mathcal{T}'$ such that $\nu_t = b(\nu_s)$, or (ii) a transition $t \xrightarrow{a}\!\!\!\bullet\; \nu_t \in \mathcal{T}'$ but there does not exist $s \xrightarrow{a}\!\!\!\bullet\; \nu_s \in \mathcal{T}_m$ such that $\nu_s \; \mathcal{L}(\approx) \; \nu_t$. We proceed with the proof of (i).

Note that this cannot be caused by two transitions with $\nu_t \neq b(\nu_s)$ but $b(\nu_s \backslash s) = \nu_t \backslash t$, since both automata are rescaled. However, since $s \approx t$, it follows that there exists $t \overset{a}{\Longrightarrow}_{\mathrm{c}} \nu_t$ such that $\nu_s \; \mathcal{L}(\approx) \; \nu_t$. Now, there are two cases: either $a \in E$, or $a \in H$. We provide the detailed proof for $a = \tau$ whose schematic proof idea is depicted below; the case $a \neq \tau$ is similar.

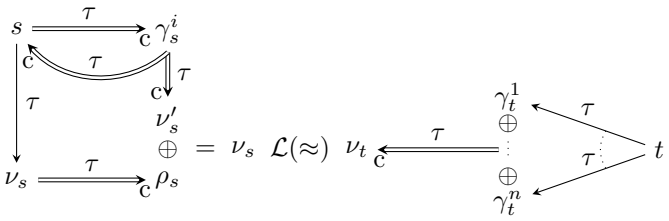

Let σ_t be the scheduler inducing $t \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_t$ and $t \xrightarrow{\tau} \gamma_t^1, \ldots, t \xrightarrow{\tau} \gamma_t^n$ be all transitions such that $\sigma_t(t)(t \xrightarrow{\tau} \gamma_t^i) > 0$ and $\gamma_t^i \not\mathcal{L}(\approx) \; \nu_s$, that is, $t \xrightarrow{\tau} \gamma_t^i$ is a transition used in the first step of the weak combined transition $t \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_t$; it is immediate to see that $(\bigoplus_{i=1}^n \gamma_t^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_t$. Since $s \approx t$, it follows that there exists γ_s^i for each $1 \leq i \leq n$ such that $s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \gamma_s^i$ and $\gamma_s^i \; \mathcal{L}(\approx) \; \gamma_t^i$. Furthermore, $(\bigoplus_{i=1}^n \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_s$, as $(\bigoplus_{i=1}^n \gamma_t^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_t$ and $\nu_t = b(\nu_s)$.

Now, consider a generic γ_s^j; there are two cases depending on whether $s \xrightarrow{\tau} \nu_s$ is used to reach ν_s. If it is not used by any of the γ_s^i, then there exists the weak combined transition $s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^n \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_s$ that does not involve $s \xrightarrow{\tau} \nu_s$, hence $s \xrightarrow{\tau} \nu_s$ can be omitted. This contradicts the $\precsim^{|\mathcal{T}|}_{\approx}$-minimality of $\mathcal{A}_m$.

So, suppose that $s \xrightarrow{\tau} \nu_s$ is used in order to reach ν_s. Since $(\bigoplus_{i=1}^n \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_s$, we may split this hyper-transition into two parts according to Lemma 2, depending on whether $s \xrightarrow{\tau} \nu_s$ is chosen by the scheduler with non-zero probability: $(\bigoplus_{i=1}^n \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \nu'_s$ with weight $c_1 \geq 0$ that does not involve $s \xrightarrow{\tau} \nu_s$, and $(\bigoplus_{i=1}^n \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s$ with weight $c_2 > 0$ that involves $s \xrightarrow{\tau} \nu_s$ such that $c_1 + c_2 = 1$ and there exists ρ_s such that ($s \xrightarrow{\tau} \nu_s$ and) $\nu_s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \rho_s$ and $\nu_s = (c_1 \nu'_s \oplus c_2 \rho_s)$. Note that we use ρ_s instead of ν_s since it may be that, in order to reach distribution equivalent to ν_s, we

have to adjust probabilities by performing more steps. Now, consider the convex combination of the two weak combined transitions $Tr_1 = s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^{n} \gamma_s^i) \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_s'$ and $Tr_2 = s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^{n} \gamma_s^i) \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s \stackrel{\tau}{\longrightarrow} \nu_s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \rho_s$, with weights c_1 and c_2, respectively. Since $(c_1\nu_s' \oplus c_2\rho_s) = \nu_s$, we have that such convex combination corresponds to the weak transition $s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_s$, so we can replace the transition $s \stackrel{\tau}{\longrightarrow} \nu_s$ by the weak combined transition $Tr = c_1 \cdot Tr_1 \oplus c_2 \cdot Tr_2$ with $\nu_s = c_1\nu_s' \oplus c_2\rho_s$. Since $s \stackrel{\tau}{\longrightarrow} \nu_s$ still occurs in $Tr_2 = s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s \stackrel{\tau}{\longrightarrow} \nu_s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \rho_s$, we can recursively replace it by the same weak combined transition Tr, hence, after k replacements, we have that $\nu_s = c_1\nu_s' \oplus c_2c_1\nu_s' \oplus c_2^2c_1\nu_s' \oplus \cdots \oplus c_2^k\rho_s = (\bigoplus_{l=0}^{k-1} c_1c_2^l\nu_s') \oplus c_2^k\rho_s$, that is, $(\bigoplus_{l=0}^{k-1}(1-c_2)c_2^l\nu_s') \oplus c_2^k\rho_s$. If we tend k to infinite, since $c_2 < 1$, we derive that $\nu_s = \nu_s'$, therefore there exists the weak combined transition $s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^{n} \gamma_s^i) \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \nu_s$ that does not involve $s \stackrel{\tau}{\longrightarrow} \nu_s$, hence again $s \stackrel{\tau}{\longrightarrow} \nu_s$ can be omitted. This contradicts the $\precsim^{|\mathcal{T}|}_{\approx}$-minimality of $\mathcal{A}_m$. The proof of case (ii) is completely analogous, except that the contradictions will be derived with respect to $\subseteq_{\mathcal{T}}$, which is a result of the fact that $\mathcal{A}'$ has been reduced according to $\overset{T}{\rightsquigarrow}$.

As final note, consider the weight c_2 and suppose that $c_2 = 1$. Since $s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^{n} \gamma_s^i) \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s$ with $(\bigoplus_{i=1}^{n} \gamma_s^i) \not\in (\approx) \delta_s$, it follows that each state in the support of $\bigoplus_{i=1}^{n} \gamma_s^i$ is actually weak bisimilar to s as the states touched in the loop $s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^{n} \gamma_s^i) \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s$ form a strongly connected component. But this contradicts the $\precsim^{|S|}_{\approx}$-minimality of $\mathcal{A}_m$. □

Corollary 1. *Let $\mathcal{A}$ be a $\precsim^{|S|}_{\approx}$-minimal PA.*

$\mathcal{A}$ is $\subseteq_{\mathcal{T}}$-minimal if and only if it is $\precsim^{|\mathcal{T}|}_{\approx}$-minimal.

Proof. Let $\mathcal{A}$ be $\precsim^{|S|}_{\approx}$-minimal. For the first direction of the *if and only if*, note first that by Lemma 9, a *PA* $\mathcal{A}'_m$ must exist, which is minimal with respect to $\precsim^{|\mathcal{T}|}_{\approx}$and $\precsim^{|S|}_{\approx}$. Let $\mathcal{A}'_m \overset{R}{\rightsquigarrow} \mathcal{A}_m$. Clearly, $\mathcal{A}_m$ must be $\precsim^{|S|}_{\approx}$and $\precsim^{|\mathcal{T}|}_{\approx}$-minimal, too. As by assumption, $\mathcal{A}$ is $\subseteq_{\mathcal{T}}$-minimal, $\mathcal{A} \overset{T}{\rightsquigarrow} \mathcal{A}$. Let $\mathcal{A}'$ satisfy $\mathcal{A} \overset{R}{\rightsquigarrow} \mathcal{A}'$. We combine the two reductions and see that $\mathcal{A} \overset{T}{\rightsquigarrow} \circ \overset{R}{\rightsquigarrow} \mathcal{A}'$. This allows us to apply Lemma 12 to obtain $\mathcal{A}' =_{iso} \mathcal{A}_m$. As $\mathcal{A}' =_{iso} \mathcal{A}_m$ implies that both have the same number of transitions, also $\mathcal{A}'$ must be $\precsim^{|\mathcal{T}|}_{\approx}$-minimal. If we can now show that also $\mathcal{A}$ and $\mathcal{A}'$ have the same number of transitions, we are done. Assume the contrary to arrive at a contradiction. As $\mathcal{A} \overset{R}{\rightsquigarrow} \mathcal{A}'$, this is only possible if there are two transitions (s, τ, μ) and (s, τ, γ) in $\mathcal{A}$ such that $\mu \backslash s = \gamma \backslash s$. But then, one of them could have been removed without changing the combined weak transitions s can perform, contradicting the assumption that $\mathcal{A}$ is $\subseteq_{\mathcal{T}}$-minimal.

For the other direction, assume $\mathcal{A}$ is in addition $\precsim^{|\mathcal{T}|}_{\approx}$-minimal. As removing transitions from $\mathcal{A}$ would lead to an automaton that is smaller with respect to $\precsim^{|\mathcal{T}|}_{\approx}$, it must be the case that any such automaton $\mathcal{A}'$ does not satisfy $\mathcal{A}' \approx \mathcal{A}$, otherwise contradicting the assumption that $\mathcal{A}$ was $\precsim^{|\mathcal{T}|}_{\approx}$-minimal. But then it immediately follows that $\mathcal{A}$ is also $\subseteq_{\mathcal{T}}$-minimal. □

Lemma 13. *If $\mathcal{A}$ is $\precsim^{\|\mathcal{T}\|}_{\approx}$-minimal, then there also exists $\mathcal{A}'$, such that $\mathcal{A} \approx \mathcal{A}'$ and $\mathcal{A}'$ is $\precsim^{|S|}_{\approx}$, $\precsim^{|\mathcal{T}|}_{\approx}$, and $\precsim^{\|\mathcal{T}\|}_{\approx}$-minimal.*

Proof. We first show that for every $\precsim_{\approx}^{\|\mathcal{T}\|}$-minimal automaton $\mathcal{A}$ there is one that is also $\precsim_{\approx}^{|S|}$-minimal. As candidate, we take the unique automaton $\mathcal{A}'$ such that $\mathcal{A} \overset{\approx}{\rightsquigarrow} \mathcal{A}'$. From Definitions 6 and 7 it is clear that the transitions of $\mathcal{A}'$ can be surjectively mapped to transitions of $\mathcal{A}$, such that every transition of $\mathcal{A}'$ is smaller or equal with respect to $\|\cdot\|$ than its image transition in $\mathcal{A}$. Thus, minimality with respect to $\precsim_{\approx}^{\|\mathcal{T}\|}$ is preserved.

Now we show that any $\mathcal{A}''$, which satisfies $\mathcal{A}' \overset{T}{\rightsquigarrow} \mathcal{A}''$ is in addition $\precsim_{\approx}^{|\mathcal{T}|}$-minimal. Clearly, the numbers of states of $\mathcal{A}'$ and $\mathcal{A}''$ are the same. Furthermore, the transitions of $\mathcal{A}''$ form a subset of the transitions of $\mathcal{A}'$. Thus, as $\mathcal{A}'$ is $\precsim_{\approx}^{\|\mathcal{T}\|}$-minimal, also $\mathcal{A}''$ must be $\precsim_{\approx}^{\|\mathcal{T}\|}$-minimal. By Definition 12, $\mathcal{A}''$ is minimal with respect to $\subseteq_{\mathcal{T}}$, and thus, by Corollary 1, also with respect to $\precsim_{\approx}^{|\mathcal{T}|}$. □

Corollary 2. *For every PA $\mathcal{A}$ there exists a PA $\mathcal{A}'$ with $\mathcal{A}' \approx \mathcal{A}$, which is $\precsim_{\approx}^{|S|}$, $\precsim_{\approx}^{|\mathcal{T}|}$ and $\precsim_{\approx}^{\|\mathcal{T}\|}$-minimal.*

Proof. Follows immediately from Lemma 3 and Lemma 13. □

Lemma 14 (Canonicity of Normal Form). *Let $\mathcal{A}' = NF_{\approx}(\mathcal{A})$. Let $\mathcal{A}_m$ be a $\precsim_{\approx}^{|S|}$, $\precsim_{\approx}^{|\mathcal{T}|}$, and $\precsim_{\approx}^{\|\mathcal{T}\|}$-minimal automaton satisfying $\mathcal{A}_m \approx \mathcal{A}$. Then $\mathcal{A}' =_{iso} \mathcal{A}_m$.*

Proof. By Corollary 2 we know that $\mathcal{A}_m$ exists such that $\mathcal{A}_m \approx \mathcal{A}$ and $\mathcal{A}_m$ is $\precsim_{\approx}^{|S|}$, $\precsim_{\approx}^{|\mathcal{T}|}$ and $\precsim_{\approx}^{\|\mathcal{T}\|}$-minimal. Furthermore, as $\mathcal{A}_m$ is $\precsim_{\approx}^{\|\mathcal{T}\|}$-minimal, it must hold $\mathcal{A}_m \overset{R}{\rightsquigarrow} \mathcal{A}_m$. Finally, as $\mathcal{A}' = NF_{\approx}(\mathcal{A})$, there must exist $\mathcal{A}''$ such that $\mathcal{A} \overset{\approx}{\rightsquigarrow} \mathcal{A}''$ and $\mathcal{A}'' \overset{T}{\rightsquigarrow} \circ \overset{R}{\rightsquigarrow} \mathcal{A}'$, and by the Definition of $\overset{\approx}{\rightsquigarrow}$ and Lemma 8, $\mathcal{A}''$ is $\precsim_{\approx}^{|S|}$-minimal. Thus, we may apply Lemma 12 to obtain our result. □

Theorem 1 now follows for $\approx$ with Corollary 2 and Lemma 14.

7 Conclusion

This paper has successfully answered the question how to compute the minimal, canonical representation of probabilistic automata under strong and weak bisimilarity, together with polynomial time minimization algorithms. Canonical forms have also appeared in axiomatic treatments of probabilistic calculi [6], but are obtained by adding transitions via saturation, so without aiming for minimality. Figure 2 summarizes what steps are needed to perform the minimization in labelled transition systems (left) and probabilistic automata (right). The triplets indicate minimality (✓) or non-minimality (×) with respect to $|S|$, then $|\mathcal{T}|$, then $\|\mathcal{T}\|$. For example, ✓✓× indicates that state and transition numbers are minimal, but transition fanout size can be non-minimal.

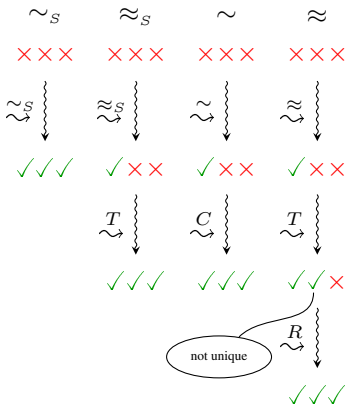

Fig. 2. Algorithmic steps in minimal quotient computation.

The algorithms we developed can be exploited in an effective compositional minimization strategy for *PA*s (or *MDP*s), because strong and weak bisimilarity are congruence relations for the standard process algebraic operators. With this, we see a rich spectrum of potential applications in operations research, automated planning, and in the decision support context.

Acknowledgements. This work is supported by DFG/NWO bilateral research programme ROCKS, by the DFG as part of the SFB/TR 14 AVACS, by the EU FP7 Programme under grant agreement no. 295261 (MEALS), 318490 (SENSATION), and TREsPASS (318003). Andrea Turrini is supported by the Cluster of Excellence "Multimodal Computing and Interaction" (MMCI), part of the German Excellence Initiative. Lijun Zhang is supported by IDEA4CPS and MT-LAB, a VKR Centre of Excellence.

References

1. C. Baier and H. Hermanns. Weak bisimulation for fully probabilistic processes. In *CAV*, pp. 119–130, 1997.
2. B. Barbot, T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Efficient CTMC model checking of linear real-time objectives. In *TACAS*, pp. 128–142, 2011.
3. S. Cattani and R. Segala. Decision algorithms for probabilistic bisimulation. In *CONCUR*, vol. 2421 of *LNCS*, pp. 371–385, 2002.
4. G. Chehaibar, H. Garavel, L. Mounier, N. Tawbi, and F. Zulian. Specification and verification of the PowerScale™ bus arbitration protocol: An industrial experiment with LOTOS. In *FORTE*, pp. 435–450, 1996.
5. P. Crouzen and F. Lang. Smart reduction. In *FASE*, vol. 6603 of *LNCS*, pp. 111–126, 2011.
6. Y. Deng and M. Hennessy. On the semantics of Markov automata. In *ICALP*, pp. 307–318, 2011.
7. C. Eisentraut, H. Hermanns, and L. Zhang. On probabilistic automata in continuous time. Reports of SFB/TR 14 AVACS 62, SFB/TR 14 AVACS, long version of LICS 342–351, 2010.
8. J.-C. Fernandez and L. Mounier. A tool set for deciding behavioral equivalences. In *CONCUR*, vol. 527 of *LNCS*, pp. 23–42, 1991.
9. H. Hermanns and J.-P. Katoen. Automated compositional Markov chain generation for a plain-old telephone system. *Science of Computer Programming*, 36(1):97–127, 2000.
10. H. Hermanns and A. Turrini. Deciding probabilistic automata weak bisimulation in polynomial time. In *FSTTCS*, pp. 435–447, 2012.
11. P. C. Kanellakis and S. A. Smolka. CCS expressions, finite state processes, and three problems of equivalence. In *PODC*, pp. 228–240, 1983.
12. M. Z. Kwiatkowska, G. Norman, and D. Parker. Prism 4.0: Verification of probabilistic real-time systems. In *CAV*, vol. 6806 of *LNCS*, pp. 585–591, 2011.
13. K. G. Larsen and A. Skou. Bisimulation through probabilistic testing (preliminary report). In *POPL*, pp. 344–352, 1989.
14. N. A. Lynch, R. Segala, and F. W. Vaandrager. Observing branching structure through probabilistic contexts. *SIAM Journal on Computing*, 37(4):977–1013, 2007.
15. R. Milner. *Communication and Concurrency*. Prentice-Hall International, 1989.
16. R. Paige and R. E. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987.
17. R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.
18. R. Segala and N. A. Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2(2):250–273, 1995.