POWVER

Technical Report 2018-12

Title: **The Quest for Minimal Quotients for Probabilistic and Markov Automata**

Authors: Christian Eisentraut, Holger Hermanns, Johann Schuster, Andrea Turrini, Lijun Zhang

Report Number: 2018-12

ERC Project: Power to the People. Verified.

ERC Project ID: 695614

Funded Under: H2020-EU.1.1. – EXCELLENT SCIENCE

Host Institution: Universität des Saarlandes, Dependable Systems and Software
Saarland Informatics Campus

Published In: Inf. Comput. 262 (1)

This report contains an author-generated version of a publication in Inf. Comput. 262 (1).

Please cite this publication as follows:

Christian Eisentraut, Holger Hermanns, Johann Schuster, Andrea Turrini, Lijun Zhang.
The Quest for Minimal Quotients for Probabilistic and Markov Automata.
Information and Computation, Volume 262, Part 1, October 2018, 162-186.

POWER TO THE PEOPLE.
VERIFIED.
erc

The Quest for Minimal Quotients for Probabilistic and Markov Automata

Christian Eisentraut[a], Holger Hermanns[a], Johann Schuster[b], Andrea Turrini[c,*], Lijun Zhang[c,d]

[a] *Saarland University – Computer Science, Saarbrücken, Germany*
[b] *University of the Federal Armed Forces Munich, Neubiberg, Germany*
[c] *State Key Laboratory of Computer Science, Institute of Software, CAS, Beijing, China*
[d] *University of Chinese Academy of Sciences, Beijing, China*

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 and Markov automata modulo strong and weak probabilistic bisimilarity, together with the corresponding minimization algorithms. We also consider weak distribution bisimilarity, originally proposed for Markov automata. For this relation, the quest for minimality does not have a unique answer, since fanout minimality clashes with state and transition minimality. We present an SMT approach to enumerate fanout-minimal models.

Keywords: Probabilistic automata, Markov automata, Weak probabilistic bisimulation, Minimal quotient, Decision algorithm

1. Introduction

Markov decision processes (*MDPs*) 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 (*PAs*) [38]. *PAs* form the backbone model of successful model checkers such as PRISM [31] and IscasMc [25] 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.

*Corresponding author, email: `turrini@ios.ac.cn`

In related fields, a favorable 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 [3, 8, 27]. 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 *MDPs* or *PAs*. One reason is that thus far no effective decision algorithm was at hand for weak bisimilarity on *PAs*. A polynomial time algorithm has been proposed only recently [41] in the form of a decision algorithm, not a minimization algorithm. The algorithm proposed in [41] follows the classical *partition refinement* approach [7, 30, 34], which computes as byproduct the bisimulation relation and that can be used as starting point for the construction of the quotient. This paper therefore focuses on a seemingly tiny problem: does there exist a *unique minimal* representative of a given probabilistic automaton with respect to weak bisimilarity? Can we compute it? In fact, this turns out to be an intricate problem. We nevertheless arrive at polynomial time algorithms.

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

These results however do not carry over to a setting where weak probabilistic bisimilarity is based on distributions. This generalization, first presented on Markov automata (*MAs*) [17], has more challenging algorithmic implications [14, 37] and these challenges are echoed in the minimization context considered here. It turns out that for distribution bisimilarity, minimality with respect to fanout might conflict with minimality with respect to states and transitions. We provide a thorough discussion of the principal phenomena for distribution bisimilarity on both PA and MA, and develop an SMT approach to enumerate fanout minimal models.

Since weak probabilistic bisimilarity enjoys congruence properties for parallel composition and hiding on *PAs*, the results in this paper enable compositional minimization approaches to be carried out efficiently. Moreover, because *PAs* comprise *MDPs*, 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 *PAs* and strong and weak bisimilarity on labelled transition systems.

The paper is an extended version of the conference paper [15]. All discussions related to distribution bisimilarity and to Markov automata are original and unpublished contributions.

Organization of the paper. 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. In Section 7 we extend the results of the previous sections to the distribution-based bisimulations. We show in Section 8 that for distribution-based bisimulations the fanout minimality conflicts with state and transition minimality. Then, in Section 9 we discuss how the obtained results carry over to the genuine Markov automata setting. We conclude the paper in Section 10 with some remarks.

2. Preliminaries

Sets, Relations, and Distributions. Given a set X, we denote by $\mathcal{P}(X)$ the power set of X, i.e., $\mathcal{P}(X) = \{\, C \mid C \subseteq X \,\}$.

Given a relation $\mathcal{R} \subseteq X \times X$, we say that $\mathcal{R}$ is a preorder if it is reflexive and transitive. We say that $\mathcal{R}$ is an equivalence relation if it is a symmetric preorder. Given an equivalence relation $\mathcal{R}$ on X, we denote by $X/_{\mathcal{R}}$ the set of equivalence classes induced by $\mathcal{R}$ and, for $x \in X$, by $[x]_{\mathcal{R}}$ the class $\mathcal{C} \in X/_{\mathcal{R}}$ such that $x \in \mathcal{C}$.

Given three sets X, Y, and Z and two relations $\mathcal{R} \subseteq X \times Y$ and $\mathcal{S} \subseteq Y \times Z$, we denote by $\mathcal{R} \circ \mathcal{S}$ the relation $\mathcal{R} \circ \mathcal{S} = \{\, (x, z) \mid \exists y \in Y. x \mathrel{\mathcal{R}} y \wedge y \mathrel{\mathcal{S}} z \,\}$.

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

Given 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 $\rho(\perp)$ the value $1 - \rho(X)$ where $\perp \notin X$, by δ_x the *Dirac* distribution such that $\rho(y) = 1$ if $y = x$, 0 otherwise, and by $\delta_\perp$ the *empty* distribution such that $|\delta_\perp| = 0$. Given $\rho \in \mathrm{SubDisc}(X)$, we may also write $\rho = \{\, (x, p_x) \mid x \in X \,\}$ where p_x is the probability $\rho(x)$ of x; we usually omit the pairs (x, p_x) where $p_x = 0$. For a constant $c \geq 0$, we denote by $c \cdot \rho$ the distribution defined by $(c \cdot \rho)(x) = c \cdot \rho(x)$ provided $c \cdot |\rho| \leq 1$. Further, for $\rho \in \mathrm{SubDisc}(X)$ 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. For $\rho \in \mathrm{SubDisc}(X)$ and $x \in X$, we denote by $\rho - x$ the distribution such that $(\rho - x)(y) = \rho(y)$ 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. Given a countable set of indices I, we say that ρ is a *convex combination* of a family of distributions $\{\rho_i \in \mathrm{SubDisc}(X)\}_{i \in I}$ if there exists a family $\{c_i \in \mathbb{R}_{\geq 0}\}_{i \in I}$ such that $\sum_{i \in I} c_i = 1$ and $\rho = \bigoplus_{i \in I} c_i \cdot \rho_i$. Finally, for $\rho \in \mathrm{SubDisc}(X)$, $x \in \mathrm{Supp}(\rho)$ and $y \notin \mathrm{Supp}(\rho)$, we denote by $\rho[y/x]$ the distribution such that $\rho[y/x](z) = \rho(x)$ if $z = y$, $\rho(z)$ otherwise.

The lifting $\mathcal{L}(\mathcal{R})$ [29] of an equivalence relation $\mathcal{R}$ on X is an equivalence relation $\mathcal{L}(\mathcal{R}) \subseteq \mathrm{Disc}(X) \times \mathrm{Disc}(X)$ defined as: for $\rho_1, \rho_2 \in \mathrm{Disc}(X)$, $\rho_1 \mathrel{\mathcal{L}(\mathcal{R})} \rho_2$ if and only if for each $\mathcal{C} \in X/_{\mathcal{R}}$, $\rho_1(\mathcal{C}) = \rho_2(\mathcal{C})$.

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

Definition 1 (Lifting of Functions)**.** Given arbitrary sets X and Y, and $\rho \in \mathrm{Disc}(X)$, we *lift* a mapping $b \colon X \to Y$ to $b \colon \mathrm{Disc}(X) \to \mathrm{Disc}(Y)$ by defining $(b(\rho))(y) = \sum_{x \in b^{-1}(y)} \rho(x)$ for each $y \in Y$.

Models. A Markov automaton (*MA*) [12, 16, 17] is a tuple $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T}_I, \mathcal{T}_M)$, where S is a countable set of *states*, $\bar{s} \in S$ is the *start* (or *initial*) state, Σ is a countable set of *actions*, $\mathcal{T}_I \subseteq S \times \Sigma \times \mathrm{Disc}(S)$ is an *interactive transition relation*, and $\mathcal{T}_M \subseteq S \times \mathbb{R}_{\geq 0} \times S$ is a *Markovian transition relation.*

We call an *MA* $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T}_I, \mathcal{T}_M)$ a Probabilistic Automaton (*PA*) [38, 39] if $\mathcal{T}_M = \emptyset$ and we call a *PA* $\mathcal{A}$ a Labelled Transition System (*LTS*) if $(s, a, \mu) \in \mathcal{T}_I$ implies $\mu = \delta_t$ for some $t \in S$. For *PAs* and *LTSs*, we may omit the empty Markovian transition relation from $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T}_I, \emptyset)$, i.e., we may simply write $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T}_I)$; in this work we consider only finite automata, i.e., automata such that S, $\mathcal{T}_I$, and $\mathcal{T}_M$ are finite. Note that we can not simply require S and Σ to be finite, since this does not guarantee that $\mathcal{T}_I$ is finite: in fact, for $S = \{s_0, s_1\}$ and $\Sigma = \{\tau\}$, we can have $\mathcal{T}_I = \{\, (s_0, \tau, \{(s_0, p), (s_1, 1-p)\}) \mid p \in [0,1] \,\}$ which is uncountable.

The set Σ is partitioned into two sets $H = \{\tau\}$ and E of internal (hidden) and external actions, respectively; we let s, t, u, v, and their variants with indices range over S and a, b range over Σ. For *MAs* (but *not* for *PAs*), states can be partitioned into stable and instable states: given $s \in S$, we call s *stable*, denoted by $s{\downarrow}$, if there is no $\mu \in \mathrm{Disc}(S)$ such that $(s, \tau, \mu) \in \mathcal{T}_I$; we call s *instable* otherwise. Intuitively, s is a stable state if the maximal progress assumption allows time to progress, as governed by Markovian transitions; s is instead an instable state if time progress is blocked by the presence of a τ transition, which forces the system to take an interactive transition immediately.

Mapping MAs to PAs. In this work we are interested in the bisimulation relations defined on Markov automata; as we will see later, such relations are essentially defined on the underlying *PA* once the Markovian transitions are considered as external interactive transitions. We now recall the mapping from *MAs* to *PAs* proposed in [12, 16, 17, 37] that maps the Markovian transitions from a state to newly added interactive transitions encoding both rates and relative probability of reaching the successor states. More precisely, given an *MA* $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T}_I, \mathcal{T}_M)$ and $s, s' \in S$, define

$$\begin{aligned} \mathit{rate}(s, s') &= \sum_{(s,\lambda,s') \in \mathcal{T}_M} \lambda, \\ \mathit{rate}(s) &= \sum_{t \in S} \mathit{rate}(s, t), \text{ and} \\ \lambda(s) &= \begin{cases} \{\, (t, \frac{\mathit{rate}(s,t)}{\mathit{rate}(s)}) \mid t \in S \,\} & \text{if } \mathit{rate}(s) \neq 0, \\ \delta_s & \text{otherwise.} \end{cases} \end{aligned}$$

Then, $\mathcal{A}$ is mapped by $\mathcal{E} \colon \mathit{MA} \to \mathit{PA}$ to the *PA* $\mathcal{A} \mapsto \mathcal{A}' = (S', \bar{s}', \Sigma', \mathcal{T}_I{}')$ whose components are $S' = S$, $\bar{s}' = \bar{s}$, $\Sigma' = \Sigma \cup \{\, \chi_{\mathit{rate}(s)} \mid s \in S \,\}$, and $\mathcal{T}_I{}' = \mathcal{T}_I \cup \{\, (s, \chi_{\mathit{rate}(s)}, \lambda(s)) \mid s \in S \wedge s{\downarrow} \,\}$, under the assumption that $\Sigma \cap \{\, \chi_{\mathit{rate}(s)} \mid s \in S \,\} = \emptyset$. Note that each new action in $\{\, \chi_{\mathit{rate}(s)} \mid s \in S \,\}$ is an external action. It is worth mentioning that $\mathcal{E}$ implements the maximal progress assumption, that is, only stable states will be able to expose timed behaviour. As a result, the image of $\mathcal{E}$ may contain unreachable fragments.

In the remainder of the paper we mainly work with *PAs* and then omit the subscript I from the interactive transition relation $\mathcal{T}_I$.

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

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

Weak Transitions of PAs. An *execution fragment* α of a *PA* $\mathcal{A}$ is a finite or infinite sequence of alternating states and actions $\alpha = s_0 a_1 s_1 a_2 s_2 \dots$ starting from a state $\mathit{first}(\alpha) = s_0$ and, if the sequence is finite, ending with a state $\mathit{last}(\alpha)$, such that for each $i > 0$ there exists $(s_{i-1}, a_i, \mu_i) \in \mathcal{T}$ such that $\mu_i(s_i) > 0$. The *length* of α, denoted by $\mathit{len}(\alpha)$, is the number of occurrences of actions in α. If α is infinite, then $\mathit{len}(\alpha) = \infty$. Denote by $\mathit{frags}(\mathcal{A})$ the set of execution fragments of $\mathcal{A}$ and by $\mathit{frags}^*(\mathcal{A})$ the set of finite execution fragments of $\mathcal{A}$. An execution fragment 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 $\mathit{trace}(\alpha)$, is the sub-sequence of external actions of α; we denote by ε the empty trace. We extend the definition of trace to actions as follows: $\mathit{trace}(a) = a$ for $a \in E$ and $\mathit{trace}(\tau) = \varepsilon$. For instance, for $a \in E$, $\mathit{trace}(s_0 a s_1) = \mathit{trace}(s_0 \tau s_1 \tau \dots \tau s_{i-1} a s_i \tau \dots \tau s_n) = a = \mathit{trace}(a)$ and $\mathit{trace}(s_0) = \mathit{trace}(s_0 \tau s_1 \tau \dots \tau s_n) = \varepsilon = \mathit{trace}(\tau)$.

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

A *scheduler* for a *PA* $\mathcal{A}$ is a function $\sigma \colon \mathit{frags}^*(\mathcal{A}) \to \mathrm{SubDisc}(\mathcal{T})$ such that for each finite execution fragment α, $\mathrm{Supp}(\sigma(\alpha)) \subseteq \mathcal{T}(\mathit{last}(\alpha), \cdot)$. A scheduler is called *Dirac* if it assigns a Dirac distribution to each finite execution fragment and it is called *determinate* [7] if for each pair of finite execution fragments α and α', if $\mathit{trace}(\alpha) = \mathit{trace}(\alpha')$ and $\mathit{last}(\alpha) = \mathit{last}(\alpha')$, then $\sigma(\alpha) = \sigma(\alpha')$. It is worthwhile to note that a determinate scheduler satisfies $\sigma(\alpha) = \sigma(\mathit{last}(\alpha))$ whenever $\mathit{trace}(\alpha) = \varepsilon$.

Given a scheduler σ and a finite execution fragment α, the distribution $\sigma(\alpha)$ describes how transitions are chosen to move on from $\mathit{last}(\alpha)$. A scheduler σ and a state s induce a probability measure $\mu_{\sigma,s}$ 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 $C_\alpha = \{\, \alpha' \in \mathit{frags}(\mathcal{A}) \mid \alpha \leqslant \alpha' \,\}$. The probability $\mu_{\sigma,s}$ of a cone C_α is defined recursively as follows:

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

Standard measure theoretical arguments ensure that $\mu_{\sigma,s}$ extends uniquely to the σ-field generated by cones. We call the measure $\mu_{\sigma,s}$ 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 $\mu_{\sigma,s}(\alpha)$ as $\mu_{\sigma,s}(\alpha) = \mu_{\sigma,s}(C_\alpha) \cdot \sigma(\alpha)(\bot)$, 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 $\mu \in \mathrm{Disc}(S)$ labelled by $a \in \Sigma$ that is induced by σ, denoted by $s \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \mu$, if there exists a scheduler σ such that the following holds for the induced probabilistic execution fragment $\mu_{\sigma,s}$:

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

We say that there is a *weak transition* from $s \in S$ to $\mu \in \mathrm{Disc}(S)$ labelled by $a \in \Sigma$ that is induced by σ, denoted by $s \stackrel{a}{\Longrightarrow} \mu$, if there exists a Dirac scheduler σ inducing $s \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \mu$.

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

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. [32, Proposition 3.6])**.** *Given a PA $\mathcal{A}$ and an action a, if there exist two weak hyper transitions $\mu_1 \xRightarrow{a}_{\mathrm{c}} \mu_2$ and $\mu_2 \xRightarrow{\tau}_{\mathrm{c}} \mu_3$ (or $\mu_1 \xRightarrow{\tau}_{\mathrm{c}} \mu_2$ and $\mu_2 \xRightarrow{a}_{\mathrm{c}} \mu_3$), then there exists the weak hyper transition $\mu_1 \xRightarrow{a}_{\mathrm{c}} \mu_3$.*

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 $\mu \xRightarrow{a}_{\mathrm{c}} \mu'$ into several weak hyper transitions $\mu_i \xRightarrow{a}_{\mathrm{c}} \mu'_i$. 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. [18, Lemmas 5 and 6])**.** *For $\mu, \mu' \in \mathrm{Disc}(S)$, $\mu \xRightarrow{a}_{\mathrm{c}} \mu'$ if and only if there exist a finite set of indexes I and two families of subdistributions $\{\mu_i \in \mathrm{SubDisc}(S)\}_{i \in I}$ and $\{\mu'_i \in \mathrm{SubDisc}(S)\}_{i \in I}$ such that $\mu = \bigoplus_{i \in I} \mu_i$, $\mu' = \bigoplus_{i \in I} \mu'_i$, and $\mu_i \xRightarrow{a}_{\mathrm{c}} \mu'_i$ for each $i \in I$.*

3. Bisimulations

3.1. State-based Bisimulations

In the following, we define state-based strong and weak (probabilistic) bisimulation relations. These relations are defined on the states of an automaton, and follow the classical idea that for each pair of related states, each transition available from one of the states has to be mimicked by the other while preserving the relation, i.e., by reaching related states.

Let $\rightsquigarrow \in \{\longrightarrow, \longrightarrow_{\mathrm{c}}, \Longrightarrow, \Longrightarrow_{\mathrm{c}}\}$.

Definition 2 (Generic State-based 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 $s \mathrel{\mathcal{R}} s'$, it holds that $s \xrightarrow{a} \mu$ implies $s' \overset{a}{\rightsquigarrow} \mu'$ for some $\mu' \in \mathrm{Disc}(S)$ such that $\mu \mathrel{\mathcal{L}(\mathcal{R})} \mu'$.

We denote by $\asymp_{\rightsquigarrow}$ the union of all $\rightsquigarrow$-bisimulations. Two *PAs* $\mathcal{A}$, $\mathcal{A}'$ are $\rightsquigarrow$-bisimilar, written $\mathcal{A} \asymp_{\rightsquigarrow} \mathcal{A}'$ if their initial states are bisimilar in the direct sum of the two *PAs*. Here, by direct sum of $\mathcal{A}_1$ and $\mathcal{A}_2$ we refer to the *PA* $\mathcal{A}$ whose set of states S is the disjoint union $S_1 \uplus S_2$ and whose set of transitions $\mathcal{T}$ is the union of $\mathcal{T}_1$ and $\mathcal{T}_2$ (cf. [39, Section 4.2]).

We recover the standard characterization for strong and weak bisimilarities from this definition as follows:

1. Strong Bisimilarity for *LTSs*, denoted $\sim_{\scriptscriptstyle LTS}$, respectively, is $\asymp_{\longrightarrow}$.

2. Strong Probabilistic Bisimilarity for *PAs*, denoted $\sim$, is $\asymp_{\longrightarrow_{\mathrm{c}}}$.

3. Weak Bisimilarity for *LTSs*, denoted $\approx_{\scriptscriptstyle LTS}$, is $\asymp_{\Longrightarrow}$.

(a) (b)

Figure 1: (a) An example of *PA*. (b) An example of *LTS*.

4. Weak Probabilistic Bisimilarity for *PAs*, denoted $\approx$, is $\asymp_{\Longrightarrow_c}$.

The relations $\sim_{LTS}$ and $\approx_{LTS}$ coincide with the respective notions of strong and weak bisimilarity on *LTS* [33]. The same holds for the probabilistic bisimilarities $\sim$ and $\approx$ on *PAs* [40]. In the sequel we assume that bisimilarities are only applied to suitable automata.

Example 1. As an example of *PAs*, consider the *PA* shown in Figure 1(a), where $p, q, r \in \mathbb{R}_{>0}$ are such that $p + q + r = 1$, and we omit the probability 1 from all other transitions. It is easy to observe that all states with the same shape and color are indeed weak probabilistic bisimilar: the states depicted as $\blacktriangle$ are the only states not enabling any transition, so they satisfy the weak probabilistic bisimulation definition; for the same motivation, they are also strong probabilistic bisimilar. Note that they can not be strong or weak probabilistic bisimilar to any other state since the latter enables at least one transition labelled with an external action which can not be matched by $\blacktriangle$. The two states $\pentagon$ are weak probabilistic bisimilar since both of them enable a c-transition reaching with probability 1 the class of $\blacktriangle$; the a-transition $\pentagon \xrightarrow{a} \delta_{\blacktriangle}$ enabled by the bottom $\pentagon$ state can be matched by the top $\pentagon$ state by first performing the transition $\pentagon \xrightarrow{\tau} \delta_{\pentagon}$ and then the transition $\pentagon \xrightarrow{a} \delta_{\blacktriangle}$, which results in the weak combined transition $\pentagon \stackrel{a}{\Longrightarrow}_c \delta_{\blacktriangle}$ for which clearly $\delta_{\blacktriangle} \mathrel{\mathcal{L}(\mathcal{R})} \delta_{\blacktriangle}$ holds. Note that the top and bottom $\pentagon$ states are not strong probabilistic bisimilar since the top $\pentagon$ state is not able to match the transition $\pentagon \xrightarrow{a} \delta_{\blacktriangle}$ enabled by the bottom $\pentagon$ state and the bottom $\pentagon$ state is not able to match the transition $\pentagon \xrightarrow{\tau} \delta_{\pentagon}$ enabled by the top $\pentagon$ state. The only state $\blacksquare$ can not be strong or weak probabilistic bisimilar to any other state since it is the only one enabling a b-transition $\blacksquare \xrightarrow{b} \delta_{\blacktriangle}$ and no other state s is able to perform a weak combined transition $s \stackrel{b}{\Longrightarrow}_c \delta_{\blacktriangle}$.

An example of *LTS* is given in Figure 1(b). As for the *PA* shown in Figure 1(a), all states with the same shape and color are weak bisimilar, with a similar motivation. They are also strong bisimilar except for the $\pentagon$ states, which are again distinguished by the transitions $\pentagon \xrightarrow{a} \delta_{\blacktriangle}$ and $\pentagon \xrightarrow{\tau} \delta_{\pentagon}$. ■

3.2. Distribution-based Bisimulations

In the setting of Markov automata, a weaker notion of bisimulation has been introduced recently [11, 12, 16, 17]. It is defined on distributions in-

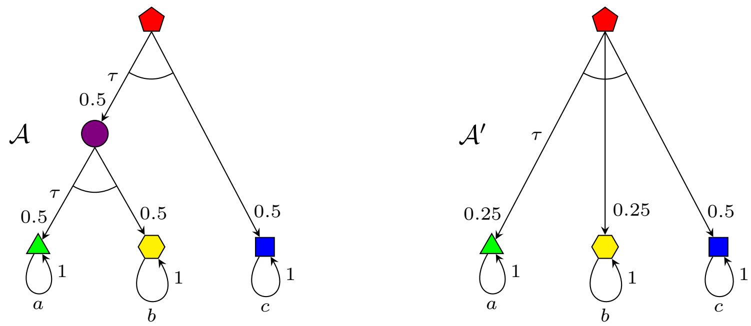

Figure 2: Two weak distribution bisimilar *PAs*

stead of states. This permits to equate *MAs* that are not weak probabilistic bisimilar despite intuitively having indistinguishable behaviour. The resulting bisimulation is coarser than weak probabilistic bisimulation but otherwise enjoys the same properties, especially wrt. compositionality. As in [14] we refer to it as *weak distribution bisimulation*, and follow the notation adopted in [16, 17]. Apart from notational differences, the relation essentially coincides with the weak bisimulation introduced in [11, 12].

Definition 3 (Weak distribution bisimulation (cf. [16, 17]))**.** Let $\mathcal{A}$ be a *PA*. An equivalence relation $\mathcal{R} \subseteq \mathrm{SubDisc}(S) \times \mathrm{SubDisc}(S)$ is a *weak distribution bisimulation* if for each pair of sub-distributions $(\mu_1, \mu_2) \in \mathcal{R}$ it holds that

1. $|\mu_1| = |\mu_2|$ and
2. for each $t \in \mathrm{Supp}(\mu_1)$ and each $a \in \Sigma$, there exist $\mu_2^{\rightarrow}, \mu_2^{\times} \in \mathrm{SubDisc}(S)$ such that
   (a) $\mu_2 \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \mu_2^{\rightarrow} \oplus \mu_2^{\times}$,
   (b) $(\mu_1(t) \cdot \delta_t) \mathrel{\mathcal{R}} \mu_2^{\rightarrow}$ and $(\mu_1 - t) \mathrel{\mathcal{R}} \mu_2^{\times}$, and
   (c) whenever $t \xrightarrow{a} \mu_1'$, then there exists μ_2' such that $\mu_2^{\rightarrow} \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \mu_2'$ and $(\mu_1(t) \cdot \mu_1') \mathrel{\mathcal{R}} \mu_2'$.

Two distributions μ_1, μ_2 are called *weak distribution bisimilar* (with respect to some *PA* $\mathcal{A}$), written $\mu_1 \approx \mu_2$, if the pair (μ_1, μ_2) is contained in a weak distribution bisimulation relation (with respect to $\mathcal{A}$).

Two *PA* are called weak distribution bisimilar if the Dirac distributions of their initial states are weak distribution bisimilar in the direct sum of the *PAs*, i.e., $\delta_{\bar{s}_1} \approx \delta_{\bar{s}_2}$.

We denote by $\approx_\delta$ the induced equivalence relation on states, i.e., $\approx_\delta = \{ (u, v) \in S \times S \mid \delta_u \approx \delta_v \}$.

Example 2. As an example of weak distribution bisimilar *PAs*, consider the two *PAs* $\mathcal{A}$ and $\mathcal{A}'$ shown in Figure 2 (cf. [17, Figure 3]). In the remainder of the example, as notation we use μ to denote a distribution in $\mathcal{A}$ and μ' to denote its corresponding counterpart in $\mathcal{A}'$. The two *PAs* $\mathcal{A}$ and $\mathcal{A}'$ are bisimilar according to the reflexive, symmetric, and transitive closure of the relation

$$\begin{aligned}\mathcal{R} = \quad & \{\, (\delta_\bigstar, \delta'_\bigstar) \mid \bigstar \in \{\pentagon, \blacktriangle, \hexagon, \blacksquare\}\,\} \\ & \cup \{(\mu_{\bullet\blacksquare}, \mu'_{\blacktriangle\hexagon\blacksquare})\} \\ & \cup \{(\delta_{\bullet}, \mu'_{\blacktriangle\hexagon}), (\mu_{\blacktriangle\hexagon}, \mu'_{\blacktriangle\hexagon})\} \\ & \cup \{(\mu_{\blacktriangle\hexagon\blacksquare}, \mu'_{\blacktriangle\hexagon\blacksquare}), (\mu_{\blacktriangle\hexagon}, \mu'_{\blacktriangle\hexagon}), (\mu_{\blacktriangle\blacksquare}, \mu'_{\blacktriangle\blacksquare}), (\mu_{\hexagon\blacksquare}, \mu'_{\hexagon\blacksquare})\} \\ & \cup \{(\delta_\perp, \delta_\perp)\}\end{aligned}$$

where $\mu_{\bullet\blacksquare} = \{(\bullet, 0.5), (\blacksquare, 0.5)\}$, $\mu_{\blacktriangle\hexagon\blacksquare} = \{(\blacktriangle, 0.25), (\hexagon, 0.25), (\blacksquare, 0.5)\}$, $\mu_{\blacktriangle\hexagon} = \{(\blacktriangle, 0.5), (\hexagon, 0.5)\}$, $\mu_{\blacktriangle\blacksquare} = \{(\blacktriangle, 0.25), (\blacksquare, 0.5)\}$, and $\mu_{\hexagon\blacksquare} = \{(\hexagon, 0.25), (\blacksquare, 0.5)\}$. Moreover, we add to $\mathcal{R}$ all pairs $(c \cdot \mu, c \cdot \rho)$ with $c \in \mathbb{R}_{\geq 0}$ that are required, provided that $(\mu, \rho) \in \mathcal{R}$ (cf. [18, Lemma 7]); for instance, the pair $(0.5 \cdot \delta_{\bullet}, 0.5 \cdot \mu_{\blacktriangle\hexagon})$ is one of the added pairs.

The only interesting pair is $(\mu_{\bullet\blacksquare}, \mu_{\blacktriangle\hexagon\blacksquare})$; all other pairs are straightforward. For the pair $(\mu_{\bullet\blacksquare}, \mu_{\blacktriangle\hexagon\blacksquare})$ obviously it holds $|\mu_{\bullet\blacksquare}| = 1 = |\mu_{\blacktriangle\hexagon\blacksquare}|$. For the state $t = \bullet$, we can take $\mu_2^{\rightarrow} = 0.5 \cdot \mu'_{\blacktriangle\hexagon}$ and $\mu_2^{\times} = 0.5 \cdot \delta'_{\blacksquare}$; clearly $\mu_{\blacktriangle\hexagon\blacksquare} \xRightarrow{\tau}_{\mathrm{c}} (\mu_2^{\rightarrow} \oplus \mu_2^{\times}) = \mu_{\blacktriangle\hexagon\blacksquare}$, $0.5 \cdot \delta_{\bullet} = (\mu_{\bullet\blacksquare}(\bullet) \cdot \delta_{\bullet})\ \mathcal{R}\ \mu_2^{\rightarrow} = 0.5 \cdot \mu'_{\blacktriangle\hexagon}$, and $0.5 \cdot \delta_{\blacksquare} = (\mu_{\bullet\blacksquare} - \bullet)\ \mathcal{R}\ \mu_2^{\times} = 0.5 \cdot \delta'_{\blacksquare}$. As μ'_2 for matching the only transition $\bullet \xrightarrow{\tau} \mu_{\blacktriangle\hexagon}$, we can take $\mu'_2 = 0.5 \cdot \mu'_{\blacktriangle\hexagon}$, which trivially satisfies $0.5\mu'_{\blacktriangle\hexagon} = \mu_2^{\rightarrow} \xRightarrow{\tau}_{\mathrm{c}} \mu'_2 = 0.5 \cdot \mu'_{\blacktriangle\hexagon}$ and $0.5 \cdot \mu_{\blacktriangle\hexagon} = (\mu_{\bullet\blacksquare}(\bullet) \cdot \mu_{\blacktriangle\hexagon})\ \mathcal{R}\ \mu'_2 = 0.5 \cdot \mu'_{\blacktriangle\hexagon}$. For the state $t = \blacksquare$, we can just take $\mu_2^{\rightarrow} = 0.5 \cdot \delta'_{\blacksquare}$ and $\mu_2^{\times} = 0.5 \cdot \mu'_{\blacktriangle\hexagon}$ and the remaining checks are trivial. ■

4. Preorders

The size of an automaton is usually expressed in terms of the size $|S|$ of the set of states S and the size $|\mathcal{T}|$ of the transition relation $\mathcal{T}$ of the automaton. The complexity of algorithms working on *PAs* 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, i.e., the size of its support. Especially in practical applications it is known that the first two of these metrics – the number of states and transitions of the automaton – can be drastically reduced while preserving its behaviour with respect to some notion of bisimilarity. In contrast, the last metric is usually considered a constant. Nevertheless 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 4 (Transition Fanout)**.** Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ be a *PA*. For a distribution $\mu \in \mathrm{Disc}(S)$ we define $\|\mu\| = |\mathrm{Supp}(\mu)|$. For a set of transitions $T \subseteq \mathcal{T}$ we define $\|T\| = \sum_{(s,a,\mu) \in T} \|\mu\|$.

Definition 5 (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 $\asymp \in \{\sim, \sim_{LTS}, \approx, \approx_{LTS}, \approxeq\}$ be a notion of bisimilarity. We write

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

In the remainder of the paper we let $\preceq$ range over $\precsim^{|S|}$, $\precsim^{|\mathcal{T}|}$, and $\precsim^{\|\mathcal{T}\|}$ where $\asymp \in \{\sim, \sim_{LTS}, \approx, \approx_{LTS}, \approxeq\}$, if not mentioned otherwise. It is easy to verify that these relations are preorders.

Definition 6 ($\preceq$-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 $\precsim^{|S|}$-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 Sections 6 and 7.

5. Reductions

In this section, we introduce and formalize several methods to reduce the size of an automaton with respect to state-based bisimulations. We postpone the reductions specific to distribution-based bisimulations to Section 7. 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 7 (Quotient Automaton)**.** Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ be a *PA*. 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 $[\mathcal{T}]_\asymp = \{\, ([s]_\asymp, a, [\mu]_\asymp) \mid (s, a, \mu) \in \mathcal{T} \,\}$.

Note that $[\mu]_\asymp$ results from lifting to distributions (cf. Definition 1) the quotient mapping on states $[\,\cdot\,]_\asymp \colon S \to \mathcal{P}(S)$ such that $s \mapsto [s]_\asymp$.

Definition 8 (Quotient Reduction)**.** We write $\mathcal{A} \overset{\asymp}{\leadsto} \mathcal{A}'$ if $\mathcal{A}' = [\mathcal{A}]_\asymp$.

Figure 3(b) shows the result of applying Definition 8 to the *PA* in Figure 3(a) and weak probabilistic bisimilarity, where bisimilar states have the same shape and color, as seen in Example 1.

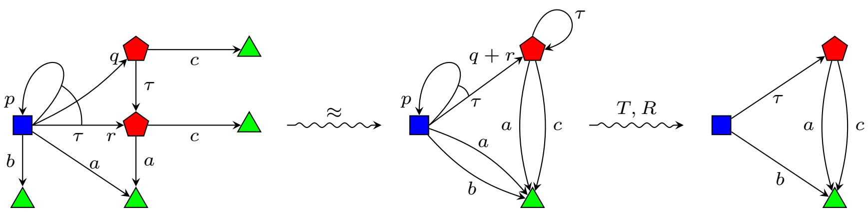

(a) (b) (c)

Figure 3: (a) Example PA (same as in Figure 1(a)). (b) Quotient reduction. (c) Rescaling of convex-transitive reduction.

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 9. Let $P = \{p_1, \dots, p_n \in \mathbb{R}^k\}$ be a finite set of points in $\mathbb{R}^k$. We call $\mathit{CHull}(P) = \{\, p \in \mathbb{R}^k \mid \exists c_1, \dots, c_n \in \mathbb{R}_{\geq 0} : \sum_{i=1}^{n} c_i = 1 \text{ and } p = \sum_{i=1}^{n} c_i \cdot p_i \,\}$ 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. [7, Section 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 10 (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

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

Note that, when $\mathcal{A}$ is an *LTS*, then the outcome $\mathcal{A}'$ of $\mathcal{A} \overset{C}{\rightsquigarrow} \mathcal{A}'$ is $\mathcal{A}$ itself, since for any given *LTS*, state s, and action a, it holds that $\mathit{Gen}(\mathit{CHull}(\{\, \mu \in \mathrm{Disc}(S) \mid (s, a, \mu) \in \mathcal{T} \,\})) = \{\, \mu \in \mathrm{Disc}(S) \mid (s, a, \mu) \in \mathcal{T} \,\}$; the inclusion $\subseteq$ is immediate, since $\mathit{Gen}(\mathit{CHull}(C)) \subseteq C$ holds for each set C. Regarding the inclusion $\supseteq$, for each *LTS* we have that each transition $(s, a, \mu) \in \mathcal{T}$ is actually of the form (s, a, δ_t) for some $t \in S$; since each Dirac distribution δ_t is representable as convex combination $\delta_t = \bigoplus_{i \in I} c_i \cdot \mu_i$ solely if $\mu_i = \delta_t$ for each $i \in I$, this implies that $\{\, \mu \in \mathrm{Disc}(S) \mid (s, a, \mu) \in \mathcal{T} \,\} \subseteq \mathit{Gen}(\mathit{CHull}(\{\, \mu \in \mathrm{Disc}(S) \mid (s, a, \mu) \in \mathcal{T} \,\}))$, as required.

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 [7].

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 11 (Convex-Transition Reduction Preorder)**.** Given two *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 $s \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \mu$ in $\mathcal{A}$.

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 $\mathcal{A} \stackrel{T}{\rightsquigarrow} \mathcal{A}'$ 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. However, 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 (*LTSs*), where it then coincides with transitive reduction of order relations. For this it is irrelevant that this reduction allows the combination of 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, $\stackrel{T}{\rightsquigarrow}$ preserves $\approx_{LTS}$, and in fact it is a necessary step to arrive at the transition minimal quotient. Minimalization concepts for *LTSs* were already introduced in [5, 6, 19, 20] while the concept of normal forms for *LTSs* has been introduced already in [6], and likely they have been considered in the context of tools exploiting weak bisimilarity [9, 23]. For the probabilistic minimizations/normal forms we are not aware of any publication.

5.4. Rescaling

A particular fine point of weak probabilistic bisimilarities [1] is related to internal transitions that induce a non-zero chance of residing inside the class. If looking at the quotient, this concerns any internal transition (s, τ, μ) that reaches 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 $a = \tau \in H$ 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 $\precapprox^{\|\mathcal{T}\|}$ if applied a posteriori to the other two reductions, $\overset{\approx}{\rightsquigarrow}$ and $\overset{T}{\rightsquigarrow}$. Figure 3(c) depicts the result of applying this sequence of reductions on the *PA* in Figure 3(a). Figure 3(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)**.** *Given two PAs $\mathcal{A}$ and $\mathcal{A}'$,*

1. $\mathcal{A} \overset{\asymp}{\rightsquigarrow} \mathcal{A}'$ *implies* $\mathcal{A} \asymp \mathcal{A}'$ *for each* $\asymp \in \{\sim, \sim_{\scriptscriptstyle LTS}, \approx, \approx_{\scriptscriptstyle LTS}\}$.
2. $\mathcal{A} \overset{C}{\rightsquigarrow} \mathcal{A}'$ *implies* $\mathcal{A} \asymp \mathcal{A}'$ *for each* $\asymp \in \{\sim, \sim_{\scriptscriptstyle LTS}, \approx, \approx_{\scriptscriptstyle LTS}\}$.
3. $\mathcal{A} \overset{T}{\rightsquigarrow} \mathcal{A}'$ *implies* $\mathcal{A} \asymp \mathcal{A}'$ *for each* $\asymp \in \{\approx_{\scriptscriptstyle LTS}, \approx\}$.
4. $\mathcal{A} \overset{R}{\rightsquigarrow} \mathcal{A}'$ *implies* $\mathcal{A} \approx \mathcal{A}'$.

Proof. The cases for $\overset{\asymp}{\rightsquigarrow}$, $\overset{C}{\rightsquigarrow}$, and $\overset{T}{\rightsquigarrow}$ follow almost immediately from the definitions of the reductions.

Now, consider $\overset{R}{\rightsquigarrow}$: since by definition of $\overset{R}{\rightsquigarrow}$, $\mathcal{A}$ and $\mathcal{A}'$ have the same set of states, we use μ 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 $S \uplus S'$ such that for each $s \in S$, $[s]_{\mathcal{I}} = \{s, s'\}$, i.e., each class corresponds to exactly four pairs in $\mathcal{I}$: (s, s), (s, s'), (s', s), and (s', s'), that is, we first relate each state s with its primed counterpart in $\mathcal{A}'$ and then take the reflexive and symmetric closure.

We claim that $\mathcal{I}$ is a weak probabilistic bisimulation for $\mathcal{A}$ and $\mathcal{A}'$: let $s \mathrel{\mathcal{I}} t$ and $s \xrightarrow{a} \mu$; if $s = t$, i.e., we are considering the pairs (s, s) or (s', s'), then also t enables the transition $t \xrightarrow{a} \mu$ and $\mu \mathrel{\mathcal{L}(\mathcal{I})} \mu$.

Suppose now that $s \neq t$, i.e., we are considering the pairs (s, s') or (s', s); if $a \in E$, then by definition of $\overset{R}{\rightsquigarrow}$ we have that also t enables the transition $t \xrightarrow{a} \mu$, thus $\mu \; \mathcal{L}(\mathcal{I}) \; \mu$ clearly holds.

Now, consider $a \in H$: if $s \in S$ and $t \in S'$, i.e., we are considering the pair (s, s'), then t is able to match such a transition by the weak combined transition $t \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \mu'$ as induced by the scheduler σ such that $\sigma(t)(\bot) = \mu(s)$, $\sigma(t)(tr) = 1 - \mu(s)$, and $\sigma(\alpha)(\bot) = 1$ for each finite execution fragment $\alpha \neq t$, where $tr = (t, \tau, \mu \backslash s)$. The resulting distribution μ' is such that $\mu'(s) = 1 \cdot \mu(s) + (1 - \mu(s)) \cdot (\mu \backslash s)(s) \cdot 1 = \mu(s) + (1 - \mu(s)) \cdot 0 = \mu(s)$ and, for $z \neq s$, $\mu'(z) = (1 - \mu(s)) \cdot (\mu \backslash s)(z) \cdot 1 = (1 - \mu(s)) \cdot \frac{\mu(z)}{(1-\mu(s))} = \mu(z)$, hence $\mu \; \mathcal{L}(\mathcal{I}) \; \mu'$, as required. Note that this applies also when $\mu = \delta_s$ as the resulting scheduler assigns $\sigma(t)(\bot) = \mu(s) = 1$ so the induced weak combined transition is $t \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_t$ and $\delta_s \; \mathcal{L}(\mathcal{I}) \; \delta_t$.

If $s \in S'$ and $t \in S$, i.e., we are considering the pair (s', s), then $s \xrightarrow{a} \mu$ is actually a transition $s \xrightarrow{\tau} \rho \backslash s$ that t is able to match by the weak combined transition $t \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \mu$ 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 $tr' = (t, \tau, \rho)$. The resulting distribution μ is such that $\mu(t) = \sum_{i=0}^{\infty} (\sigma(t)(tr') \cdot \rho(t))^i \cdot \sigma(t)(\bot) = \sum_{i=0}^{\infty} (1 \cdot \rho(t))^i \cdot 0 = 0 = (\rho \backslash t)(t)$, since σ being determinate implies that $\sigma(\alpha)(tr') = \sigma(t)(tr') = 1$ and $\sigma(\alpha)(\bot) = \sigma(t)(\bot) = 0$ whenever $\mathit{last}(\alpha) = t$, and $\mu(z) = \sum_{i=0}^{\infty} (\sigma(t)(tr') \cdot \rho(t))^i \cdot (\sigma(t)(tr') \cdot \rho(z)) \cdot \sigma(z)(\bot) = \sum_{i=0}^{\infty} (1 \cdot \rho(t))^i \cdot (1 \cdot \rho(z)) \cdot 1 = \rho(z) \cdot \sum_{i=0}^{\infty} (\rho(t))^i = \rho(z) \cdot \frac{1}{1-\rho(t)} = (\rho \backslash t)(z)$, hence $\rho \backslash s \; \mathcal{L}(\mathcal{I}) \; \mu = \rho \backslash t$, as required. □

Proposition 7 (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 $\rightsquigarrow \in \{\overset{\asymp}{\rightsquigarrow}, \overset{C}{\rightsquigarrow}, \overset{T}{\rightsquigarrow}, \overset{R}{\rightsquigarrow}\}$ and $\asymp \in \{\sim, \sim_{\scriptscriptstyle LTS}, \approx, \approx_{\scriptscriptstyle LTS}\}$.*

Proof outline. The result for $\overset{\asymp}{\rightsquigarrow}$ follows by the corresponding polynomial decision procedures [7, 23, 30, 34, 41] and reachability analysis; $\overset{C}{\rightsquigarrow}$ requires for each state and each enabled action to solve $\mathcal{O}(|\mathcal{T}|)$ linear programming problems (cf. [7, Section 6]) in order to find the set of generators of reachable distributions; $\overset{R}{\rightsquigarrow}$ can be obtained by computing for each transition $s \xrightarrow{\tau} \mu$ the distribution $\mu \backslash s$, which 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 [41] 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 their 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 $NF_{\asymp}\colon PA \to PA$ a *normal form*, if it satisfies for every *PA* $\mathcal{A}$

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

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 $NF_{\sim_{LTS}} = \overset{\sim_{LTS}}{\rightsquigarrow}$.
- Let $NF_{\approx_{LTS}}$ be an arbitrary function in $\overset{\approx_{LTS}}{\rightsquigarrow} \circ \overset{T}{\rightsquigarrow}$.
- Let $NF_{\sim} = \overset{\sim}{\rightsquigarrow} \circ \overset{C}{\rightsquigarrow}$.
- Let $NF_{\approx}$ be an arbitrary function in $\overset{\approx}{\rightsquigarrow} \circ \overset{T}{\rightsquigarrow} \circ \overset{R}{\rightsquigarrow}$.

Theorem 8. *Let $\asymp \in \{\sim, \sim_{LTS}, \approx, \approx_{LTS}\}$.*

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 8. We begin with a lemma that we use often.

Lemma 9 (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 10 (State Minimality of Quotient Automata)**.** *For every $\mathcal{A} \in PA$, the automaton $\mathcal{A}'$ with $\mathcal{A} \overset{\asymp}{\rightsquigarrow} \mathcal{A}'$ is $\underset{\smile}{\preceq}^{|S|}$-minimal.*

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

Lemma 11 (Compatibility of $\underset{\smile}{\preceq}^{|S|}$ and $\underset{\smile}{\preceq}^{|\mathcal{T}|}$-minimality)**.** *For every PA $\mathcal{A}$ there exists a PA $\mathcal{A}'$ with $\mathcal{A}' \asymp \mathcal{A}$, which is $\underset{\smile}{\preceq}^{|S|}$ and $\underset{\smile}{\preceq}^{|\mathcal{T}|}$-minimal.*

Proof. By Lemma 3, there exists a *PA* $\mathcal{A}$ that is $\underset{\smile}{\preceq}^{|\mathcal{T}|}$-minimal. Consider $\mathcal{A}' = [\mathcal{A}]_{\asymp}$. From Definition 7 it is clear that for every transition of $[\mathcal{A}]_{\asymp}$ there exists a transition in $\mathcal{A}$. Thus, $[\mathcal{A}]_{\asymp} \underset{\smile}{\preceq}^{|\mathcal{T}|} \mathcal{A}$, and hence, $[\mathcal{A}]_{\asymp}$ must also be $\underset{\smile}{\preceq}^{|\mathcal{T}|}$-minimal. Furthermore, by Lemma 10, $[\mathcal{A}]_{\asymp}$ must also be $\underset{\smile}{\preceq}^{|S|}$-minimal, and finally by Lemma 6 we have that $\mathcal{A} \asymp [\mathcal{A}]_{\asymp} = \mathcal{A}'$. ☐

6.1. Strong Bisimilarities

Lemma 12 (Canonicity of Normal Form)**.** *Let $\asymp \in \{\sim_{LTS}, \sim\}$, $\mathcal{A} \in PA$, and $\mathcal{A}' = NF_{\asymp}(\mathcal{A})$. For every $\underset{\smile}{\preceq}^{|S|}$ and $\underset{\smile}{\preceq}^{|\mathcal{T}|}$-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 $\asymp = \sim_{LTS}$ and proceed with the more complicated case of $\asymp = \sim$. Recall that $NF_{\sim} = \overset{\sim}{\rightsquigarrow} \circ \overset{C}{\rightsquigarrow}$. As applying $\overset{\sim}{\rightsquigarrow}$ to $\mathcal{A}$ leads to a $\precsim^{|S|}$-minimal automaton according to Lemma 10, and $\overset{C}{\rightsquigarrow}$ obviously does not alter the number of states, $\mathcal{A}' = NF_{\sim}(\mathcal{A})$ is $\precsim^{|S|}$-minimal, and thus $|S_m| = |S'|$, as $\mathcal{A}_m$ is $\precsim^{|S|}$-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 $b(s_2) = t$. This implies that $s_1 \sim t$ and $s_2 \sim t$. By transitivity, this implies $s_1 \sim s_2$, contradicting Lemma 10. 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 $\precsim^{|S|}$-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 $\mathcal{A}' \sim \mathcal{A}_m$.

The condition that b maps $\bar{s}'$ to $\bar{s}_m$ is immediate by definition of b; the last condition to be shown to have b being an isomorphism is relative to the transitions, i.e.,

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

The set of combined transitions any state s of $\mathcal{A}'$ can take must equal (modulo b) the set of combined transitions that $b(s)$ can take 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 $\sim_{LTS}$ and $\sim$, Theorem 8 now follows almost immediately by Lemma 11, Lemma 12, and Lemma 6. For $\sim_{LTS}$, we in addition need the observation that $\mathcal{A}$ is $\precsim^{\|\mathcal{T}\|}$-minimal if and only if it is $\precsim^{|\mathcal{T}|}$-minimal, as we remarked before Lemma 11. For $\sim$, the same observation holds, but follows from the uniqueness of the minimal set of generators (Lemma 4).

6.2. Weak Bisimilarities

The following two lemmas are the weak counterparts to Lemma 12.

Lemma 13. *Let $\mathcal{A}$ be a PA and $\mathcal{A}' = NF_{\approx_{LTS}}(\mathcal{A})$. Let $\mathcal{A}_m$ be a $\precapprox^{|S|}_{LTS}$ and $\precapprox^{|\mathcal{T}|}_{LTS}$-minimal PA satisfying $\mathcal{A}_m \approx_{LTS} \mathcal{A}$. 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. The *LTS* part of Theorem 8 can then be proven in complete analogy to the proof for $\sim$.

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 $\precsim^{|S|}$ and $\precsim^{|\mathcal{T}|}$-minimal automata, but they will not agree up to $=_{iso}$, in full generality. Different to Lemmas 13 and 12, the following lemma is slightly more general.

Lemma 14. *Let $\mathcal{A}$ be a $\precapprox^{|S|}$-minimal PA, $\mathcal{A} \overset{T}{\rightsquigarrow} \circ \overset{R}{\rightsquigarrow} \mathcal{A}'$, and $\mathcal{A}'_m$ be a $\precapprox^{|S|}$ and $\precapprox^{|\mathcal{T}|}$-minimal PA satisfying $\mathcal{A}'_m \approx \mathcal{A}$. Finally, 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. By following the same argumentation as in the proof of Lemma 12, we can show that $b = \approx \cap (S_m \times S')$ is a bijection. To complete the proof, we need to establish that b is a suitable mapping so that $\mathcal{A}_m =_{iso} \mathcal{A}'$ follows.

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 are required to be $\precapprox^{|S|}$-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 $s \xrightarrow{a} \mu_s \in \mathcal{T}_m$ but there does not exist $t \xrightarrow{a} \mu_t \in \mathcal{T}'$ such that $\mu_s \; \mathcal{L}(\approx) \; \mu_t$, i.e., there does not exist a transition $t \xrightarrow{a} \mu_t \in \mathcal{T}'$ such that $\mu_t = b(\mu_s)$, or

(ii) a transition $t \xrightarrow{a} \mu_t \in \mathcal{T}'$ but there does not exist $s \xrightarrow{a} \mu_s \in \mathcal{T}_m$ such that $\mu_s \; \mathcal{L}(\approx) \; \mu_t$.

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

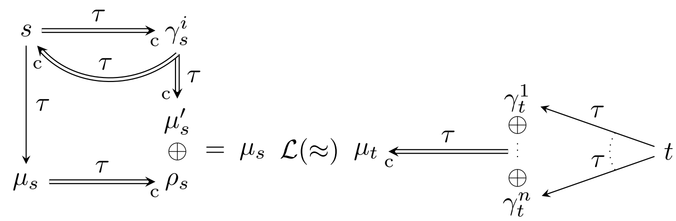

Let σ_t be the scheduler inducing $t \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \mu_t$ and $t \xrightarrow{\tau} \gamma_t^1, \ldots, t \xrightarrow{\tau} \gamma_t^n$ be all transitions such that $p_i = \sigma_t(t)(t \xrightarrow{\tau} \gamma_t^i) > 0$ and $\gamma_t^i \not\mathcal{L}(\approx) \mu_s$,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

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

So, suppose that $s \xrightarrow{\tau} \mu_s$ is used in order to reach μ_s. Since $(\bigoplus_{i=1}^n p_i \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \mu_s$, we may split this hyper-transition into two parts according to Lemma 2, depending on whether $s \xrightarrow{\tau} \mu_s$ is chosen by the scheduler with non-zero probability: $(\bigoplus_{i=1}^n p_i \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \mu_s'$ with weight $c_1 \geq 0$ that does not involve $s \xrightarrow{\tau} \mu_s$, and $(\bigoplus_{i=1}^n p_i \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s$ with weight $c_2 > 0$ that involves $s \xrightarrow{\tau} \mu_s$ such that $c_1 + c_2 = 1$ and there exists ρ_s such that ($s \xrightarrow{\tau} \mu_s$ and) $\mu_s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \rho_s$ and $\mu_s = (c_1 \mu_s' \oplus c_2 \rho_s)$. Note that we use ρ_s instead of μ_s since it may be that, in order to reach a distribution equivalent to μ_s, we have to adjust probabilities by performing more steps. Now, consider the convex combination of the two weak combined transitions $Tr_1 = s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^n p_i \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \mu_s'$ and $Tr_2 = s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^n p_i \gamma_s^i) \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s \xrightarrow{\tau} \mu_s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \rho_s$, with weights c_1 and c_2, respectively. Since $(c_1 \mu_s' \oplus c_2 \rho_s) = \mu_s$, we have that such convex combination corresponds to the weak transition $s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \mu_s$, so we can replace the transition $s \xrightarrow{\tau} \mu_s$ by the weak combined transition $Tr = c_1 \cdot Tr_1 \oplus c_2 \cdot Tr_2$ with $\mu_s = c_1 \mu_s' \oplus c_2 \rho_s$. Since $s \xrightarrow{\tau} \mu_s$ still occurs in $Tr_2 = s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s \xrightarrow{\tau} \mu_s \overset{\tau}{\Longrightarrow}_{\mathrm{c}} \rho_s$, we can recursively replace it by the same weak combined transition Tr, hence, after k replacements, we have that $\mu_s = c_1 \mu_s' \oplus c_2 c_1 \mu_s' \oplus c_2^2 c_1 \mu_s' \oplus \cdots \oplus c_2^k \rho_s = (\bigoplus_{l=0}^{k-1} c_1 c_2^l \mu_s') \oplus c_2^k \rho_s$, that is, $(\bigoplus_{l=0}^{k-1} (1 - c_2) c_2^l \mu_s') \oplus c_2^k \rho_s$. If we let k tend to infinite, since $c_2 < 1$, we derive that $\mu_s = \mu'_s$, therefore there exists the weak combined transition $s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^{n} p_i \gamma_s^i) \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \mu_s$ that does not involve $s \xrightarrow{\tau} \mu_s$, hence again $s \xrightarrow{\tau} \mu_s$ can be omitted. This contradicts the $\precapprox^{|\mathcal{T}|}$-minimality of $\mathcal{A}_m$.

As a final note, consider the weight c_2 and suppose that $c_2 = 1$. Since $s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^{n} p_i \gamma_s^i) \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s$ with $(\bigoplus_{i=1}^{n} p_i \gamma_s^i) \not\mathcal{L}(\approx)\ \delta_s$, it follows that each state in the support of $\bigoplus_{i=1}^{n} p_i \gamma_s^i$ is actually weak bisimilar to s as the states visited along the loop $s \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} (\bigoplus_{i=1}^{n} p_i \gamma_s^i) \stackrel{\tau}{\Longrightarrow}_{\mathrm{c}} \delta_s$ form a strongly connected component. This contradicts the $\precapprox^{|S|}$-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 $\stackrel{T}{\rightsquigarrow}$. More precisely, since it is not possible to find a transition $s \xrightarrow{a} \mu_s \in \mathcal{T}_m$ such that there does not exist a transition $t \xrightarrow{a} \mu_t \in \mathcal{T}'$ such that $\mu_t = b(\mu_s)$, it follows that $|\mathcal{T}'| \geq |\mathcal{T}_m|$. In order to complete the proof that $\mathcal{A}' =_{iso} \mathcal{A}_m$, we have to show that $|\mathcal{T}'| = |\mathcal{T}_m|$. Suppose, for the sake of contradiction, that $|\mathcal{T}'| > |\mathcal{T}_m|$, that is, there exists a transition $t \xrightarrow{a} \mu_t \in \mathcal{T}'$ such that there does not exist a transition $s \xrightarrow{a} \mu_s \in \mathcal{T}_m$ with $b(s) = t$ such that $\mu_t = b(\mu_s)$. By following the same approach as before, we derive that $t \xrightarrow{a} \mu_t$ can be replaced by the weak combined transition $t \stackrel{a}{\Longrightarrow}_{\mathrm{c}} \mu_t$ that does not involve $t \xrightarrow{a} \mu_t$, but this contradicts the fact that $\mathcal{A}'$ is the outcome of $\stackrel{T}{\rightsquigarrow}$. Hence, since $|\mathcal{T}'| = |\mathcal{T}_m|$ and for each transition $s \xrightarrow{a} \mu_s \in \mathcal{T}_m$ there exists $t \xrightarrow{a} \mu_t \in \mathcal{T}'$ with $t = b(s)$ such that $\mu_t = b(\mu_s)$ and vice-versa, b is an isomorphism between $\mathcal{A}'$ and $\mathcal{A}_m$, thus $\mathcal{A}' =_{iso} \mathcal{A}_m$, as required. □

Corollary 15. *Let $\mathcal{A}$ be a $\precapprox^{|S|}$-minimal PA.*
$\mathcal{A}$ is $\subseteq_{\mathcal{T}}$-minimal if and only if it is $\precapprox^{|\mathcal{T}|}$-minimal.

Proof. Let $\mathcal{A}$ be $\precapprox^{|S|}$-minimal. For the first direction of the *if and only if*, note first that by Lemma 11, a *PA* $\mathcal{A}'_m$ must exist, which is minimal with respect to $\precapprox^{|\mathcal{T}|}$ and $\precapprox^{|S|}$. Let $\mathcal{A}'_m \stackrel{R}{\rightsquigarrow} \mathcal{A}_m$. Clearly, $\mathcal{A}_m$ must be $\precapprox^{|S|}$ and $\precapprox^{|\mathcal{T}|}$-minimal, too. As by assumption, $\mathcal{A}$ is $\subseteq_{\mathcal{T}}$-minimal, $\mathcal{A} \stackrel{T}{\rightsquigarrow} \mathcal{A}$. Let $\mathcal{A}'$ satisfy $\mathcal{A} \stackrel{R}{\rightsquigarrow} \mathcal{A}'$. We combine the two reductions and see that $\mathcal{A} \stackrel{T}{\rightsquigarrow} \circ \stackrel{R}{\rightsquigarrow} \mathcal{A}'$. This allows us to apply Lemma 14 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 $\precapprox^{|\mathcal{T}|}$-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 $\mathcal{A} \stackrel{R}{\rightsquigarrow} \mathcal{A}'$, this is only possible if there are two transitions (s, τ, μ) and (s, τ, γ) in $\mathcal{A}$ such that $\mu \backslash s = \gamma \backslash s$. This means that 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 $\precapprox^{|\mathcal{T}|}$-minimal. As removing transitions from $\mathcal{A}$ would lead to an automaton that is smaller with respect to $\precapprox^{|\mathcal{T}|}$, it must be the case that any such automaton $\mathcal{A}'$ does not satisfy $\mathcal{A}' \approx \mathcal{A}$, otherwise this would contradict the assumption that $\mathcal{A}$ is $\precapprox^{|\mathcal{T}|}$-minimal. It immediately follows that $\mathcal{A}$ is also $\subseteq_{\mathcal{T}}$-minimal. $\square$

Lemma 16. *If $\mathcal{A}$ is $\precapprox^{\|\mathcal{T}\|}$-minimal, then there also exists $\mathcal{A}'$, such that $\mathcal{A} \approx \mathcal{A}'$ and $\mathcal{A}'$ is $\precapprox^{|S|}$, $\precapprox^{|\mathcal{T}|}$, and $\precapprox^{\|\mathcal{T}\|}$-minimal.*

Proof. We first show that for every $\precapprox^{\|\mathcal{T}\|}$-minimal automaton $\mathcal{A}$ there is one that is also $\precapprox^{|S|}$-minimal. As candidate, we take the unique automaton $\mathcal{A}'$ such that $\mathcal{A} \overset{\approx}{\rightsquigarrow} \mathcal{A}'$. From Definitions 7 and 8 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 $\precapprox^{\|\mathcal{T}\|}$ is preserved.

Now we show that any $\mathcal{A}''$, which satisfies $\mathcal{A}' \overset{T}{\rightsquigarrow} \mathcal{A}''$ is in addition $\precapprox^{|\mathcal{T}|}$-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 $\precapprox^{\|\mathcal{T}\|}$-minimal, also $\mathcal{A}''$ must be $\precapprox^{\|\mathcal{T}\|}$-minimal. By Definition 12, $\mathcal{A}''$ is minimal with respect to $\subseteq_{\mathcal{T}}$, and thus, by Corollary 15, also with respect to $\precapprox^{|\mathcal{T}|}$. $\square$

Corollary 17. *For every PA $\mathcal{A}$ there exists a PA $\mathcal{A}'$ with $\mathcal{A}' \approx \mathcal{A}$, which is $\precapprox^{|S|}$, $\precapprox^{|\mathcal{T}|}$, and $\precapprox^{\|\mathcal{T}\|}$-minimal.*

Proof. Follows immediately from Lemmas 3 and 16. $\square$

Lemma 18 (Canonicity of Normal Form)**.** *Let $\mathcal{A}_{NF_\approx} = NF_\approx(\mathcal{A})$. Let $\mathcal{A}_m$ be a $\precapprox^{|S|}$, $\precapprox^{|\mathcal{T}|}$, and $\precapprox^{\|\mathcal{T}\|}$-minimal automaton satisfying $\mathcal{A}_m \approx \mathcal{A}$. Then $\mathcal{A}_{NF_\approx} =_{iso} \mathcal{A}_m$.*

Proof. By Corollary 17 we know that $\mathcal{A}_m$ exists such that $\mathcal{A}_m \approx \mathcal{A}$ and $\mathcal{A}_m$ is $\precapprox^{|S|}$, $\precapprox^{|\mathcal{T}|}$ and $\precapprox^{\|\mathcal{T}\|}$-minimal. Furthermore, as $\mathcal{A}_m$ is $\precapprox^{\|\mathcal{T}\|}$-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 10, $\mathcal{A}''$ is $\precapprox^{|S|}$-minimal. Thus, we can apply Lemma 14 to obtain the desired result. $\square$

Theorem 8 now follows for $\approx$ with Corollary 17 and Lemma 18.

7. Weak Distribution Bisimulation

In the previous sections, we have seen how to obtain the normal forms for the state-based bisimulations. These normal forms are generated by applying in sequence different reductions working on the states and the transitions of the automaton for which we want to find the normal form. The first reduction minimizes the number of states (and reduces the number of transitions as byproduct) by taking the quotient under the considered bisimulation. Then, the following reductions take care to remove transitions that are not necessary and to reduce the fanout. Extending blindly this approach to a distribution-based bisimulation turns out to be not achievable, since the first step, the generation of the quotient automaton, does not provide an automaton that is $\underset{\approx}{\prec}^{|S|}$-minimal, and thus a different construction has to be adopted to obtain a normal form. This is the topic of this section where we show how to define a normal form for the weak distribution bisimulation. The key operation is the elimination of vanishing states, i.e., states that are fully replaceable by a distribution they can reach. Such vanishing states are at the core of the weak distribution bisimulation decision algorithm presented in [37] and their dual –the pivotal states– of the decision algorithm of [14].

To simplify the presentation, we first introduce some notation inspired by [37].

Definition 16. Given a *PA* $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$, a state $s \in S$, and a fresh state $t \notin S$, we introduce the following notation.

- The set of fresh internal transitions corresponding to the internal weak combined transitions from s is $\mathcal{WT}(s, \tau) = \{\, (s, \tau, \mu) \mid s \xRightarrow{\tau}_{\mathrm{c}} \mu \,\}$;
- The renamed automaton $\mathcal{A}[t/s]$ is the *PA* $\mathcal{A}[t/s] = (S', \bar{s}', \Sigma, \mathcal{T}')$ where $S' = (S \setminus \{s\}) \cup \{t\}$, $\bar{s}' = \bar{t}$ if $\bar{s} = s$, $\bar{s}' = \bar{s}$ otherwise, and $\mathcal{T}' = \{\, (v, a, \mu[t/s]) \mid (v, a, \mu) \in \mathcal{T} \setminus \mathcal{T}(s, \cdot\,) \,\} \cup \{\, (t, a, \mu[t/s]) \mid (s, a, \mu) \in \mathcal{T}(s, \cdot\,) \,\}$. Essentially, by renaming s as t we generate a copy of $\mathcal{A}$ where each occurrence of s has been replaced by t;
- For a finite $T \subseteq \mathcal{WT}(s, \tau)$, the T-replaced automaton $\mathcal{A}_T$ is the *PA* $\mathcal{A}_T = (S, \bar{s}, \Sigma, (\mathcal{T} \setminus \mathcal{T}(s, \cdot\,)) \cup T)$. If $T = \{(s, \tau, \mu)\}$ we may write $\mathcal{A}_{(s,\mu)}$ instead of $\mathcal{A}_{\{(s,\tau,\mu)\}}$. Essentially, we replace all transitions with source s with the new internal transitions from T.

We are now ready to define the vanishing states:

Definition 17. Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ be a *PA* and $s \in S$ be instable. We say that the state s is

trivially vanishing if $\mathcal{T}(s, \cdot\,) = \{(s, \tau, \mu)\}$ for some $\mu \in \mathrm{Disc}(S)$;

vanishing if there exists $(s, \tau, \mu) \in \mathcal{WT}(s, \tau)$ such that $s \approx_\delta t$ when comparing $\mathcal{A}[t/s]$ and $\mathcal{A}_{(s,\mu)}$ for $t \notin S$. In this case $\mathcal{A}_{(s,\mu)}$ – or (s, μ), for short – is called a *vanishing representation* of s;

naïvely vanishing if it is vanishing and for all vanishing representations $\mathcal{A}_{(s,\mu)}$ and all $t \in \mathrm{Supp}(\mu)$, we have that $s \approx_\delta t$; and

non-naïvely vanishing (denoted by $s_{\boxtimes}$) if it is vanishing but not naïvely vanishing, that is, there is a vanishing representation $\mathcal{A}_{(s,\mu)}$ such that there exists $t \in \mathrm{Supp}(\mu)$ such that $s \not\approx_\delta t$.

We say that s is *non-naïvely tangible*, denoted by $s_{\square}$, if it is not non-naïvely vanishing.

We extend the notation to set of states in the expected way, that is, we denote the set of all non-naïvely vanishing states by $S_{\boxtimes}$ and the set of all non-naïvely tangible states by $S_{\square}$.

Example 3. As an example of vanishing state, consider the *PA* in Figure 4a and the state Y. Y is trivially vanishing, since it enables exactly one transition $Y \xrightarrow{\tau} \mu$, where $\mu = \{(A, 2/5), (B, 1/5), (C, 1/5), (D, 1/5)\}$. Y is also non-naïvely vanishing: it is vanishing but not naïvely vanishing because for $A \in \mathrm{Supp}(\mu)$ it clearly holds that $Y \not\approx_{\delta} A$. We refer the interested reader to [37, Example 2] for more examples of vanishing states. ■

The elimination of vanishing states has been defined in [37]. We present these ideas now in the context of reductions. For this paper we do not consider general elimination, but we restrict ourselves to the elimination of only states in $S_{\boxtimes}$: it is known (cf. [37, Theorem 2]) that the elimination of such states reduces the weak distribution bisimulation to weak probabilistic bisimulation. Moreover, it is shown (cf. [37, Lemma 6]) that all states in $S_{\boxtimes}$ can be eliminated, since if $s, t \in S_{\boxtimes}$, after the removal of t it still holds that $s \in S_{\boxtimes}$ (cf. [37, Corollary 1]). Note that naïvely vanishing states do not necessarily have to be eliminated in advance: after quotient reduction with respect to $\approx$, all naïvely vanishing states are collapsed with the bisimilar states in the support of their respective vanishing representations.

Definition 18 (Elimination of the non-naïvely vanishing states $S_{\boxtimes}$, cf. [37])**.** Given a *PA* $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$, a state $s \in S_{\boxtimes}$, and a *PA* $\mathcal{A}' = (S, \bar{s}, \Sigma, \mathcal{T}')$ such that $\mathcal{A} \overset{R}{\rightsquigarrow} \mathcal{A}'$, let $\mathcal{A}'_{(s,\mu)}$ be the vanishing representation of s in $\mathcal{A}'$. Let

$$\mathcal{T}'' = \{\, (t, a, (\rho(s) \cdot \mu) \oplus (\rho - s)) \mid (t, a, \rho) \in \mathcal{T}', t \neq s \,\}.$$

The elimination of s from $\mathcal{A}'$ is defined as:

$$\mathcal{A}'^{\widehat{s}} = \begin{cases} (S \setminus \{s\}, \bar{s}, \Sigma, \mathcal{T}'') & \text{if } s \neq \bar{s}, \\ (S \setminus \{s\} \cup \{\bar{s}_f\}, \bar{s}_f, \Sigma, \mathcal{T}'' \cup \{\bar{s}_f \xrightarrow{\tau} \mu\}) & \text{otherwise,} \end{cases}$$

where $\bar{s}_f \notin S$ is a fresh state.

For $S'_{\boxtimes} = \{s_1, \ldots, s_n\}$, consider $\mathcal{A}'' = ((\mathcal{A}'^{\widehat{s_1}})'^{\widehat{s_2}} \ldots)'^{\widehat{s_n}}$; we write $\mathcal{A} \overset{V}{\rightsquigarrow} \mathcal{A}''$ for the elimination of *all non-naïvely vanishing* states.

Remark 1. In the above elimination, we have to take particular care of the case $\bar{s} \in S_{\boxtimes}$. In this case, we have to ensure to provide a new initial state for the resulting automaton: this is obtained by adding $\bar{s}_f \notin S$ as new initial state and the transition $\bar{s}_f \xrightarrow{\tau} \mu$. It is immediate to see that $\bar{s}_f$ is trivially vanishing and that it is transient since it has no incoming transitions.

As pointed out in [39, Remark 1], *PAs* can also be defined by considering an initial probability distribution instead of a single initial state. In such a scenario, for $\mathcal{A}' = (S, \bar{\iota}, \Sigma, \mathcal{T}')$ and $s \in S_{\boxtimes}$, we can define $\mathcal{A}'^{\widehat{s}}$ as the *PA* $\mathcal{A}'^{\widehat{s}} = (S \setminus \{s\}, \bar{\iota}', \Sigma, \mathcal{T}'')$ where $\bar{\iota}' = (\bar{\iota}(s)\mu) \oplus (\bar{\iota} - s)$.

It is worth mentioning that the resulting automaton $\mathcal{A}''$ does not depend on the actual order of removal of states in $S_\boxtimes$. More precisely, for the symmetric group of n elements G_n and each pair of permutations $\pi_1, \pi_2 \in G_n$, we have that $\mathcal{A}''_1 = \mathcal{A}''_2$ where $\mathcal{A}''_i = (\mathcal{A}'^{\widehat{s_{\pi_i(1)}}} \dots)^{\widehat{s_{\pi_i(n)}}}$, $i \in \{1, 2\}$. This follows from the facts that all states in $S_\boxtimes$ can be eliminated (cf. [37, Lemma 6]) and that the elimination commutes (cf. [36, Lemma 5]).

The following theorem relates state-based and distribution-based bisimulation, that is, after elimination of the states in $S_\boxtimes$, distribution-based bisimulation boils down to state-based bisimulation.

Theorem 19 (cf. [37, Theorem 2])**.** *Let $\mathcal{A}_1$, $\mathcal{A}_2$ be two PAs and $\mathcal{A}'_1$, $\mathcal{A}'_2$ be such that $\mathcal{A}_1 \overset{V}{\rightsquigarrow} \mathcal{A}'_1$ and $\mathcal{A}_2 \overset{V}{\rightsquigarrow} \mathcal{A}'_2$. Then,*

$$\mathcal{A}_1 \Approx \mathcal{A}_2 \iff \mathcal{A}'_1 \approx \mathcal{A}'_2.$$

Intuitively, Theorem 19 implies that after elimination of non-naïvely vanishing states from $\mathcal{A}_1$ and $\mathcal{A}_2$, in the resulting *PAs* $\mathcal{A}'_1$ and $\mathcal{A}'_2$ are such that whenever $s' \Approx_\delta t'$, then $s' \approx t'$, i.e., $[u']_{\Approx_\delta} = [u']_\approx$ for each state $u' \in S'_1 \cup S'_2$. By Theorem 19 we immediately get the following result:

Corollary 20. *The following mappings from PA to PA are normal forms with respect to $\Approx$: $\overset{V}{\rightsquigarrow} \circ \overset{\approx}{\rightsquigarrow} \circ \overset{T}{\rightsquigarrow} \circ \overset{R}{\rightsquigarrow}$ and $\overset{\Approx_\delta}{\rightsquigarrow} \circ \overset{V}{\rightsquigarrow} \circ \overset{T}{\rightsquigarrow} \circ \overset{R}{\rightsquigarrow}$. Moreover, for each PA $\mathcal{A}$, for $\mathcal{A}'$ and $\mathcal{A}''$ such that $\mathcal{A} \overset{V}{\rightsquigarrow} \circ \overset{\approx}{\rightsquigarrow} \circ \overset{T}{\rightsquigarrow} \circ \overset{R}{\rightsquigarrow} \mathcal{A}'$ and $\mathcal{A} \overset{\Approx_\delta}{\rightsquigarrow} \circ \overset{V}{\rightsquigarrow} \circ \overset{T}{\rightsquigarrow} \circ \overset{R}{\rightsquigarrow} \mathcal{A}''$, it holds that $\mathcal{A}' =_{iso} \mathcal{A}''$. We denote both normal forms by $NF_{\Approx}$.*

Proof. The following diagram justifies the first normal form (left equivalences follow by Theorem 19, right equivalences by Theorem 8):

$$\begin{tikzcd}
\mathcal{A}_1 \arrow[r, "\overset{V}{\rightsquigarrow}"] \arrow[d, dash, dashed, "\Approx"] & \mathcal{A}'_1 \arrow[r, "\overset{\approx}{\rightsquigarrow}"] \arrow[d, dash, dashed, "\approx"] & \mathcal{A}''_1 \arrow[r, "\overset{T}{\rightsquigarrow}"] & \mathcal{A}'''_1 \arrow[r, "\overset{R}{\rightsquigarrow}"] & \mathcal{A}''''_1 \arrow[d, dash, dashed, "=_{iso}"] \\
\mathcal{A}_2 \arrow[r, "\overset{V}{\rightsquigarrow}"] & \mathcal{A}'_2 \arrow[r, "\overset{\approx}{\rightsquigarrow}"] & \mathcal{A}''_2 \arrow[r, "\overset{T}{\rightsquigarrow}"] & \mathcal{A}'''_2 \arrow[r, "\overset{R}{\rightsquigarrow}"] & \mathcal{A}''''_2
\end{tikzcd}$$

The second normal form follows from the first one by the following diagram:

$$\begin{tikzcd}
\mathcal{A} \arrow[r, "\overset{\Approx_\delta}{\rightsquigarrow}"] \arrow[d, "\overset{V}{\rightsquigarrow}"] & \mathcal{A}' \arrow[d, "\overset{V}{\rightsquigarrow}"] \\
\mathcal{A}'' \arrow[r, dash, dashed, "\approx"] \arrow[d, "\overset{\approx}{\rightsquigarrow}"] & \mathcal{A}''' \arrow[dl, dash, dashed, "\approx"'] \\
\mathcal{A}''''
\end{tikzcd}$$

By definition of the quotient mapping it is clear that $\mathcal{A} \Approx \mathcal{A}'$. Theorem 19 then gives $\mathcal{A}'' \approx \mathcal{A}'''$. Now by definition of the quotient mapping it is also clear that $\mathcal{A}'' \approx \mathcal{A}''''$, so by transitivity we get also $\mathcal{A}''' \approx \mathcal{A}''''$. The claim is now that it even holds that $\mathcal{A}''' =_{iso} \mathcal{A}''''$. By Theorem 19 it follows that an equivalence class of states in $S_{\square}$ (and of an initial state in $S_{\boxtimes}$) wrt. $\approx_\delta$ after elimination of states in $S_{\boxtimes}$ corresponds to an equivalence class induced by $\approx$. It remains to show that also the transitions coincide. Note that by [37, Lemma 9], elimination of states in $S_{\boxtimes}$ corresponds to substituting them by their unique-up-to-bisimilarity canonical vanishing representations[1] (that is, vanishing representations (s, μ) where $\mathrm{Supp}(\mu) \subseteq S_{\square}$). Now quotienting makes these representations unique. Note that still $\mathcal{A}'''$ and $\mathcal{A}''''$ are not identical because by the current definition of elimination the fresh symbol $\bar{s}_f$ for an eliminated initial state $\bar{s}_{\boxtimes}$ is not defined more specifically. However, the fresh symbols $\bar{s}_f$ are trivially isomorphic. □

Lemma 21 (State Minimality of Quotient Automata after Elimination)**.** *For every $\mathcal{A} \in PA$, the PA $\mathcal{A}'$ such that $\mathcal{A} \overset{\approx_\delta}{\rightsquigarrow} \circ \overset{V}{\rightsquigarrow} \mathcal{A}'$ is $\underset{\approx}{\prec}^{|S|}$-minimal.*

Proof. By Theorem 19 it is clear that elimination leads to the weak bisimulation case. In [37] it has been shown that all states in $S_{\boxtimes}$ can be eliminated and the proof of Corollary 20 shows that after quotienting with respect to $\approx_\delta$ and elimination the result is already a quotient with respect to $\approx$. Lemma 10 shows that this quotient is already state-minimal. □

Corollary 22. *$NF_{\approx}$ is $\underset{\approx}{\prec}^{|S|}$-minimal and $\underset{\approx}{\prec}^{|\mathcal{T}|}$-minimal.*

Proof. It remains to show that the number of transitions is minimal. This follows directly from Theorem 8 and the fact that – even when $\bar{s} \in S_{\boxtimes}$ – elimination never results in more transitions than before elimination. □

Example 4. Figure 4 shows that the fanout may increase by elimination. The *PA* in Figure 4a has a cumulative fanout of 14. State Y is trivially (non-naïvely) vanishing. The *PA* obtained after elimination of Y (cf. Figure 4b) has a cumulative fanout of 15. ■

Therefore we have shown:

Lemma 23. *$NF_{\approx}$ is in general* not *$\underset{\approx}{\prec}^{\|\mathcal{T}\|}$-minimal.*

This lemma seems to be in contrast with Theorem 8 for the state-based weak bisimulation, where state-, transition-, and fanout-minimality are reached simultaneously. The key difference is that by removing non-naïvely vanishing states, it is the case that some state of the original automaton has no more a representative in the normal form (so we are able to reduce further the number of states), and this is reflected by an increase of the fanout that is caused by the replacement of the removed state with a distribution. Still we can state:

[1] As usual, initial states in $S_{\boxtimes}$ are made transient by elimination.

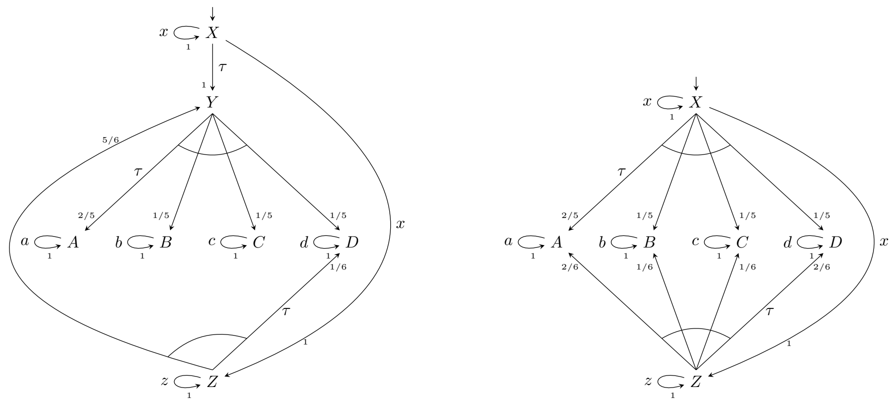

(a) Before elimination (b) After elimination

Figure 4: Fanout may increase after elimination

Lemma 24 (Canonicity of Normal Form)**.** *Given a PA $\mathcal{A}$, with $\bar{s}$ being non-naïvely tangible, let $\mathcal{A}_{NF_{\approx}} = NF_{\approx}(\mathcal{A})$. Let $\mathcal{A}_m \approx \mathcal{A}$ be a $|S|$ minimal PA. Let further $\mathcal{A}_m$ be $\|\mathcal{T}\|$ minimal among all such $|S|$ minimal PAs. In other words $\mathcal{A}_m$ is $(|S|, \|\mathcal{T}\|)$ minimal with respect to the usual lexicographic order $(\cdot, \cdot)$. Then $\mathcal{A}_{NF_{\approx}} =_{iso} \mathcal{A}_m$.*

Proof. When the initial state $\bar{s}$ is a non-naïvely tangible state, minimality ensures that we have a quotient and all states in $S_{\boxtimes} \setminus \{\bar{s}_m\}$ have been eliminated (cf. Lemma 21). Among the state minimal automata, fanout minimality ensures that after elimination convex-transitive reduction and rescaling have been performed (cf. Lemma 18). So with Theorem 19, Corollary 20 and Corollary 18 the claim follows. □

Remark 2 (Stronger canonicity result)**.** According to the definition, each *PA* has a unique initial state; as noted in [39], nothing prevents us to define *PAs* so that they can have multiple initial states or even an *initial distribution*, similarly to Markov chains and *MDPs* (cf., e.g, [2]). If one drops the requirement that an explicit start state shall exist, even stronger canonicity results are possible. Assume that initial distributions are allowed and whenever there is a non-naïvely vanishing state, it is described by its canonical vanishing representation, i.e., the vanishing representation containing only non-naïvely tangible states (which is unique up to isomorphism, cf. [37, Lemma 9]). In this context, non-naïvely vanishing initial states may be omitted from the state space. Therefore the statement of the lemma can be simplied as: *Given a PA $\mathcal{A}$, let $\mathcal{A}_{NF_{\approx}} = NF_{\approx}(\mathcal{A})$. Let $\mathcal{A}_m \approx \mathcal{A}$ be a $|S|$-minimal PA. Let further $\mathcal{A}_m$ be $\|\mathcal{T}\|$-minimal among all such $|S|$-minimal PAs. Then $\mathcal{A}_{NF_{\approx}} =_{iso} \mathcal{A}_m$.*

Proposition 25. *For every PA $\mathcal{A}$, a PA $\mathcal{A}'$ can be found in exponential time, such that $\mathcal{A} \rightsquigarrow \mathcal{A}'$ for $\rightsquigarrow \in \{\overset{\approx_\delta}{\rightsquigarrow}, \overset{V}{\rightsquigarrow}\}$.*

Proof. The result follows by the exponential complexity of computing the set $\mathcal{WT}(s, \tau)$ for each $s \in S$, see [37]. □

Theorem 26 (Complexity of Normal Forms)**.** *For every PA $\mathcal{A}$,*

- *$NF_{\sim_{LTS}}(\mathcal{A})$, $NF_{\approx_{LTS}}(\mathcal{A})$, $NF_{\sim}(\mathcal{A})$, and $NF_{\approx}(\mathcal{A})$ can be computed in polynomial time;*
- *$NF_{\approxeq}(\mathcal{A})$ can be computed in exponential time.*

Proof. The result follows immediately by Propositions 7 and 25 by simply adding the complexity of each reduction involved in the corresponding definition of $NF_{\asymp}$. □

8. Fanout minimal automata

In this section we introduce a sketch of an algorithm which, starting from the normal form of an automaton, calculates its fanout minimal automata. We will see that in general there are many of those and there is no canonical fanout minimal automaton.

By Lemma 23, we know that automata that are minimal with respect to the number of states and transitions may not be minimal with respect to the fanout. This implies that there exist automata that are minimal with respect to the fanout but they have more states and transitions. Note that having the same states but more transitions would just increase the fanout and anyway the extra transitions would be redundant. A way to reduce the fanout is to add some new state and transition that are then shared among other transitions; essentially, we act as the inverse of $\overset{V}{\rightsquigarrow}$. The following basic lemma shows that non-naïvely vanishing states play a key role when fanout minimality will be reached:

Lemma 27. *Given a PA $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ with fanout n, let n' be the fanout of the PA $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})'$ obtained from $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ by adding a vanishing state v' and merging transitions. If v' is a naïvely vanishing state or a non-trivially vanishing state, then $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})'$ is not fanout-minimal.*

Proof. Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ be a *PA* with fanout n. We show first that adding naïvely vanishing states is not beneficial for reaching minimal fanout. Assume that we add a naïvely vanishing state v' to get –after merging transitions– an automaton $\mathcal{A}'$ with fanout $n' < n$. By definition of a naïvely vanishing state it must hold that $v' \approxeq v$ for some $v \in S$. Therefore all transitions which share v' could equivalently share v without losing bisimilarity. This implies that there is surely a bisimilar automaton $\mathcal{A}''$ with state set S with fanout $n'' < n'$ since v' and its transitions can be replaced by v and its transitions.

Next we show that adding a non-trivially vanishing state v' to the automaton $\mathcal{A}$ is not beneficial for reaching minimal fanout. By definition, non-trivially vanishing states have other emanating transitions in addition to their vanishing representation (v', τ, μ). Assume that adding the non-trivially vanishing state v' leads –after merging transitions– to an automaton $\mathcal{A}'$ with fanout $n' < n$. By definition of non-trivially vanishing states it is clear that without losing bisimilarity v' can be replaced by a state v'' having only the emanating transition (v'', τ, μ). This leads to an automaton $\mathcal{A}''$ with fanout $n'' < n'$, because the emanating transitions of v' which are not equal to the vanishing representation have been omitted. □

So the proposal is to successively add non-naïvely vanishing states which are trivially vanishing.

8.1. Equations for adding a single state

Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ be a *PA*, $n = |S|$, and $m = |\mathcal{T}|$. Our aim is to reduce the fanout of $\mathcal{A}$ by modifying it while being weak distribution bisimilar; we obtain this by adding a new state and a new transition that are then shared among the original transitions so that the overall fanout is reduced. As an example, consider an automaton having three transitions leading to the distributions $\mu_1 = \{(s_1, \frac{1}{4}), (s_4, \frac{1}{4}), (s_5, \frac{1}{4}), (s_6, \frac{1}{4})\}$, $\mu_2 = \{(s_2, \frac{1}{4}), (s_4, \frac{1}{4}), (s_5, \frac{1}{4}), (s_6, \frac{1}{4})\}$, and $\mu_3 = \{(s_3, \frac{1}{4}), (s_4, \frac{1}{4}), (s_5, \frac{1}{4}), (s_6, \frac{1}{4})\}$. We can see that such distributions share the subdistribution $\{(s_4, \frac{1}{4}), (s_5, \frac{1}{4}), (s_6, \frac{1}{4})\}$ so a way to reduce the overall fanout is the following: if we add a fresh state s_7 and the transition (s_7, τ, ν_4) where $\nu_4 = \{(s_4, \frac{1}{3}), (s_5, \frac{1}{3}), (s_6, \frac{1}{3})\})$ and we replace the above distributions by $\nu_1 = \{(s_1, \frac{1}{4}), (s_7, \frac{3}{4})\}$, $\nu_2 = \{(s_2, \frac{1}{4}), (s_7, \frac{3}{4})\}$, and $\nu_3 = \{(s_3, \frac{1}{4}), (s_7, \frac{3}{4})\}$, the overall fanout reduces from 12 to 9 and the two automata are weak distribution bisimilar. We now formalise such a construction.

Definition 19. Let $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$ be a *PA*. Let $n = |S|$, $m = |\mathcal{T}|$ and assume two sets of fixed indices $I_S = \{1, \ldots, n\}$ for states and $I_{\mathcal{T}} = \{1, \ldots, m\}$ for transitions, respectively. In the following, let $\mu_i = \mathit{trg}(tr_i)$ for each $i \in I_{\mathcal{T}}$ and denote by $\mu_{i,j}$ the value $\mu_i(s_j)$ for each $i \in I_{\mathcal{T}}$ and $j \in I_S$. This gives rise to a mapping $\mathcal{M}\colon \mathit{PA} \to \mathbb{R}^{m \times n}$, where $\mathcal{A} \mapsto (\mu_{i,j})_{i \in I_{\mathcal{T}}, j \in I_S}$.

Note that all labels of states and transitions as well as the information about the initial state are lost by this mapping. However this mapping is sufficient as it is always clear that we add a τ transition and leave the labels of the existing transitions as they are. For an automaton as given in Definition 19 we construct a new automaton $\mathcal{A}' = (S', \bar{s}, \Sigma, \mathcal{T}')$ by adding a fresh state s_{n+1} and a fresh transition $tr_{m+1} = (s_{n+1}, \tau, \nu_{m+1})$, that is, $S' = S \cup \{s_{n+1}\}$ and $\mathcal{T}' = \mathcal{T} \cup \{tr_{m+1}\}$. It will suffice to consider $\nu_{m+1} \in \mathrm{Disc}(S)$ instead of $\nu_{m+1} \in \mathrm{Disc}(S')$ since a loop to the new state s_{n+1} will surely not lead to fanout-minimality. In fact, a trivial fact we can observe is that the newly introduced state cannot have a non-trivial probability $0 < \nu_{m+1,n+1} < 1$. This is justified by the fact that rescaling at s_{n+1} would lead to a bisimilar automaton with fanout reduced by one.

Fact 28. *For fanout-minimal automata it must hold that $\nu_{m+1,n+1} = 0$.*

Denote by $I_{S'} = \{1, \dots, n, n+1\}$ and $I_{\mathcal{T}'} = \{1, \dots, m, m+1\}$ the index sets for S' and $\mathcal{T}'$, respectively.

According to Definition 19, the mapping $\mathcal{M}(\mathcal{A}')$ would result in the following matrix:

$$\left(\begin{array}{ccc|c} \mu_{1,1} & \cdots & \mu_{1,n} & 0 \\ \vdots & \ddots & \vdots & \vdots \\ \mu_{m,1} & \cdots & \mu_{m,n} & 0 \\ \hline \nu_{m+1,1} & \cdots & \nu_{m+1,n} & 0 \end{array}\right)$$

From the matrix it can be seen that the fresh state is so far unreachable. The aim is to construct a new matrix $(\nu_{i,j})_{i \in I_{\mathcal{T}'}, j \in I_{S'}}$ (that is, a new automaton) with at least one non-zero entry in the $(n+1)$-th column of the first m rows.

In order to get identical successor distributions after the elimination of the state s_{n+1}, the following system of equations enriched by non-negativity constraints has to hold:

$$\mathbf{NSS} = \begin{cases} \nu_{i,j} + \nu_{m+1,j} \cdot \nu_{i,n+1} = \mu_{i,j} & \text{for each } i \in I_{\mathcal{T}},\ j \in I_S, \\ \sum_{j \in I_{S'}} \nu_{i,j} \qquad\qquad = 1 & \text{for each } i \in I_{\mathcal{T}'}, \\ \nu_{i,j} \qquad\qquad\qquad\quad\ \geq 0 & \text{for each } i \in I_{\mathcal{T}'},\ j \in I_{S'}. \end{cases}$$

A solution of such a system represents how $\mathcal{A}$ can be modified so that the resulting automaton $\mathcal{A}'$ satisfies $\mathcal{A} \approx \mathcal{A}'$. Essentially, we replace the probability of reaching s_j according to the original distribution μ (i.e., $\mu_{i,j}$) with the residual probability given by ν and the probability of the freshly introduced ν_{m+1}, weighted by the probability of reaching s_{n+1} according to ν. More precisely, for a given transition $tr_i = (s, a, \mu_i)$, let ν_i be the distribution induced by the solution **NSS**, i.e., ν_i is such that $(\nu_i - s_{n+1}) + \nu_i(s_{n+1}) \cdot \nu_{m+1} = \mu_i$; then the transitions of $\mathcal{A}'$ are $\{\, (s, a, \nu_i) \mid (s, a, \mu_i) \in \mathcal{T} \,\} \cup \{(s_{n+1}, \tau, \nu_{m+1})\}$.

For a fixed $\nu_{m+1} \in \mathrm{Disc}(S)$, the above problem can be seen as a Linear Programming Problem where the objective function is constantly 0; on the other hand, if we consider also $\nu_{m+1,j}$ as variables, the above problem can be seen as a Quadratically Constrained Quadratic Problem. Since this system encodes all transitions, including the newly added transition tr_{m+1}, it is immediate to see that the fanout of $\mathcal{A}'$ is the number of the non-zero entries $\nu_{i,j}$ of the solution of the programming problem **NSS**, for $i \in I_{\mathcal{T}'}$ and $j \in I_{S'}$. In general we can observe:

Fact 29. *For a PA $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$, its fanout is equal to the number of non-zero entries of the matrix $\mathcal{M}(\mathcal{A})$.*

A trivial fact about the above system, affecting the performance of the algorithm, is that whenever $\mu_{i,j}$ is zero, that also $\nu_{i,j}$ is zero as well. This is because all the summands in **NSS** are nonnegative.

Fact 30. *When $\mu_{i,j} = 0$, then also $\nu_{i,j} = 0$ for all $i \in I_{\mathcal{T}}$, $j \in I_S$.*

By Fact 30, we can reduce the a priori number of variables in the problem **NSS** to those which are actually needed. Of course those with index $j = n+1$ and the variables describing ν_{m+1} remain untouched.

Remark 3 (non-naïvely vanishing initial state)**.** As recalled in Remark 2, one can consider *PAs* with an initial distribution instead of an initial state. In this setting, we can further reduce the fanout of the normal form whenever it contains a (transient) non-naïvely vanishing initial state $\bar{s}$. Assume that $(\bar{s}, \tau, \mu)$ is the emanating transition of $\bar{s}$; we can directly use μ as initial distribution and solve **NSS** where we set $\nu_{m+1,i} = \mu_i$ for $i \in S$; this means that only the remaining $\nu_{i,j}$ variables of **NSS** have to be determined.

8.2. Extension to multiple states

The above approach can be easily extended to multiple additional states: a first possibility is to add k new states and distributions, so that the mapping applied to the resulting automaton $\mathcal{A}'$ would result in the following matrix:

$$\left(\begin{array}{ccc|c} \mu_{1,1} & \cdots & \mu_{1,n} & 0 \\ \vdots & \ddots & \vdots & \vdots \\ \mu_{m,1} & \cdots & \mu_{m,n} & 0 \\ \hline \nu_{m+1,1} & \cdots & \nu_{m+1,n} & 0 \\ \vdots & \ddots & \vdots & \vdots \\ \nu_{m+k,1} & \cdots & \nu_{m+k,n} & 0 \end{array}\right)$$

The resulting $\mathbf{NSS}_k$ would then be:

$$\mathbf{NSS}_k = \begin{cases} \nu_{i,j} + \sum_{l=1}^{k} \nu_{m+l,j} \cdot \nu_{i,n+l} = \mu_{i,j} & \text{for each } i \in I_{\mathcal{T}},\ j \in I_S, \\ \sum_{j \in I_{S_k}} \nu_{i,j} \qquad\qquad\qquad\qquad = 1 & \text{for each } i \in I_{\mathcal{T}_k}, \\ \nu_{i,j} \qquad\qquad\qquad\qquad\qquad\quad \geq 0 & \text{for each } i \in I_{\mathcal{T}_k},\ j \in I_{S_k}, \end{cases}$$

where $S_k = S \cup \{\, s_{n+l} \mid 1 \leq l \leq k \,\}$ and $\mathcal{T}_k = \mathcal{T} \cup \{\, (s_{n+l}, \tau, \nu_{m+l}) \mid 1 \leq l \leq k \,\}$ for fresh states $s_{n+1}, \ldots, s_{n+k}$.

Another possibility is to repeat the proposed approach for **NSS** on the obtained automaton $\mathcal{A}'$, i.e., to construct an automaton $\mathcal{A}''$ such that the mapping $\mathcal{M}(\mathcal{A}'')$ gives:

$$\left(\begin{array}{ccc|c} \mu_{1,1} & \cdots & \mu_{1,n} & 0 \\ \vdots & \ddots & \vdots & \vdots \\ \mu_{m,1} & \cdots & \mu_{m,n} & 0 \\ \nu_{m+1,1} & \cdots & \nu_{m+1,n} & 0 \\ \hline \nu_{(m+1)+1,1} & \cdots & \nu_{(m+1)+1,n} & 0 \end{array}\right)$$

When repeated k times, this approach is a generalization of the first one, since it permits to obtain the $\mathbf{NSS}_k$ system. In addition, it allows us to have an automaton where some distribution ν_{m+l} has the state $s_{n+l'}$ in its support, for $1 \leq l, l' \leq k$, $l' \neq l$. The order with which we insert new states does not affect the final resulting automaton, as the following remark states.

Remark 4 (Commutativity of insertions)**.** Given a *PA* $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$, consider $\{\, s_i \in S_{\boxtimes} \mid i \in \{1, \ldots, n\} \,\}$. We already mentioned in Section 7 that for the symmetric group of n elements G_n and each pair of permutations $\pi_1, \pi_2 \in G_n$, we have that $\mathcal{A}''_1 = \mathcal{A}''_2$ where $\mathcal{A}''_l = (\mathcal{A}'^{\widehat{s_{\pi_l(1)}}} \ldots)'^{\widehat{s_{\pi_l(n)}}}$, $l \in \{1, 2\}$. In other words, eliminations of non-naïvely vanishing states commute. Therefore also when inserting non-naïvely vanishing states one can also expect some kind of commutativity. First note that by elimination of $s_i \in S_\boxtimes$, the actual vanishing representation of every other state s_j, $i \neq j$, might be changed. This means that the *same* $s_j \in S_\boxtimes$ can have different vanishing representations before and after the elimination of s_i. It is known (cf. [37, Lemma 9]) that there exist canonical vanishing representations $(s, \hat{\mu})$ of (s, μ) which are unique on quotient automata. They only consist of non-naïvely tangible states in $\mathrm{Supp}(\hat{\mu})$. By identifying (s, μ) with its canonical vanishing representation $(s, \hat{\mu})$ one can speak of commutativity. This means inserting (s_1, μ_1) after (s_2, μ_2) leads to the same results as inserting (s_2, γ_2) after (s_1, γ_1) if the results are fanout-minimal automata and $\hat{\mu}_l = \hat{\gamma}_l$, $l \in \{1, 2\}$.

8.3. Practical verification

To verify the effects of adding fresh states, we have developed a proof-of-concept program for finding all possible automata corresponding to the solutions of the **NSS** system. To this end, we encoded the **NSS** system as an SMT-LIB instance [4] and then used Z3 [10] to find the minimum fanout by iteratively reducing the possible value of the fanout, until the solver returns that the resulting instance is not satisfiable. We have implemented in Python a program that, given the SMT-LIB representation of **NSS** and the desired fanout, enumerates all different automata with at most the given fanout that are solutions of **NSS**. This is obtained by iteratively inserting into the SMT-LIB instance new constraints as follows: we first ask the Z3 library to return a model of the current instance; if such model does not exist (because the instance is unsatisfiable), the program terminates. Otherwise, we add a constraint requiring that not all variables are equal to their corresponding value in the returned model: for instance, suppose that for the current instance I_k we get a model M_k assigning to each variable $\nu_{i,j}$ a value $v^k_{i,j}$. We then define the new constraint $C_{k+1} = \neg \bigwedge_{i,j} (\nu_{i,j} = v^k_{i,j})$ and we add it to I_k, obtaining the new instance $I_{k+1} = I_k \wedge C_{k+1}$; if I_{k+1} is satisfiable, then a model M_{k+1} for I_{k+1} must ensure that for at least one variable $\nu_{i,j}$, it holds $M_{k+1}(\nu_{i,j}) = v^{k+1}_{i,j} \neq v^k_{i,j} = M_k(\nu_{i,j})$, i.e., M_{k+1} is different from M_k.

As an example, consider the *PA* shown in Figure 5, whose fanout is 19. By encoding it as an instance of **NSS** and running the program, we obtained 12 different automata with fanout 18 in less than 4 seconds on a desktop PC equipped with an Intel Core i7-4790 processor at 3.6GHz with 16GB of RAM; some of the obtained automata are shown in Figure 6. We have also tried to reduce the fanout to 17; for this experiment, after 16 seconds we have obtained that **NSS** is unsatisfiable, i.e., there are no automata with fanout 17 that are equivalent to the automaton in Figure 5.

We have also run the experiments on the extension to multiple states: the problem $\mathbf{NSS}_2$ has no solution when the fanout is 18; on the other hand, if

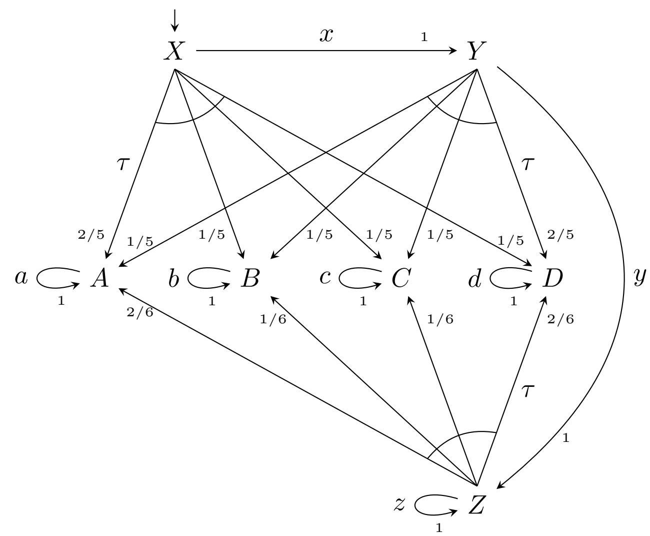

Figure 5: Example automaton

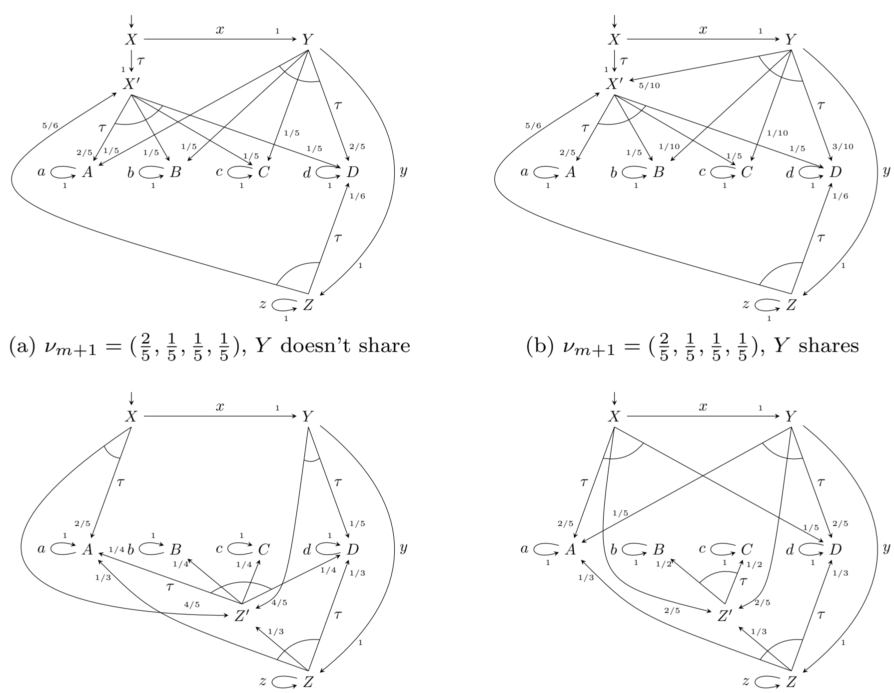

(a) $\nu_{m+1} = (\frac{2}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5})$, Y doesn't share

(b) $\nu_{m+1} = (\frac{2}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5})$, Y shares

(c) $\nu_{m+1} = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$

(d) $\nu_{m+1} = (0, \frac{1}{2}, \frac{1}{2}, 0)$

Figure 6: Some fanout minimal automata with 8 states; the distribution ν_{m+1} is represented as vector with entries A, B, C, and D, respectively

we take the more general approach where we apply twice the **NSS** approach, then we obtain 21 automata with fanout 18 but no automaton with fanout 17. We want to remark that each of the above results on multiple states has been obtained in around 20 seconds on a simplified version of the automaton, where all self-loop transitions (e.g., (Z, z, δ_Z)) have been omitted. In fact, even if it is clear that such self-loop transitions remain unchanged in the obtained automata, their presence makes Z3 much slower: after 15 minutes we stopped the execution without having obtained any automaton.

9. Application to Markov Automata

Section 2 recalled the mapping [17] $\mathcal{E}\colon MA \to PA$, defined in such a way that for a *MA* $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T}_I, \mathcal{T}_M)$, the property of a state s being stable (respectively instable) can be recovered from $\mathcal{E}(\mathcal{A})$: states with outgoing $\chi_{(\,\cdot\,)}$ transitions are stable, the others must be instable. In a similar way we say that a state s is non-naïvely vanishing (respectively non-naïvely tangible) when the corresponding state in $\mathcal{E}(\mathcal{A})$ is non-naïvely vanishing (respectively non-naïvely tangible).

While this mapping makes it possible to treat *MAs* and *PAs* uniformly, this does not mean that the original *MA* definition is without use. In fact, the latter is the one arising naturally if starting from a compositional syntax for *MAs*, such as the MAPA language supported by the MaMa tool [24]. Especially parallel composition $\|_{MA}$ is straightforward to define on *MAs* [17] (combining $\|_{PA}$ on *PAs* [40] with interleaving of Markov transitions). However, the two operators do not commute with $\mathcal{E}$: given two *MAs* $\mathcal{A}_1$ and $\mathcal{A}_2$, $\mathcal{E}(\mathcal{A}_1 \|_{MA} \mathcal{A}_2)$ can differ from $\mathcal{E}(\mathcal{A}_1) \|_{PA} \mathcal{E}(\mathcal{A}_2)$, due to the treatment of Markov transitions in the former.

As a result, a compositional minimisation approach for *MAs*, where components are replaced by their normal form (so as to alleviate state-space explosion) prior to parallel composition is best done by recovering the *MA* representation from the normal form *PA* of the embedding. This raises the question if this recovery is possible without destroying minimality on the *MA* level. In the sequel, we answer affirmatively this question.

In order to define normal forms for *MAs* we shall introduce a reverse mapping, thus lifting *PAs* back to *MAs*, provided the former are obtained from $\mathcal{E}$. However using an arbitrary element in the preimage of $\mathcal{E}(\,\cdot\,)$ might result in transitions which are superfluous due to the maximal progress assumption. This is not opportune, since we target minimality. Apart from this, the definition of the lifting provided below is quite straightforward. The only special case concerns the definition of $\mathcal{T}_I$, where we have to "pad" τ-loops wherever needed in order to preserve stability and instability of states.

Definition 20 (Backward Lifting)**.** Given the set image $\mathcal{E}(MA)$ and its corresponding normal form images $NF_{\asymp}(\mathcal{E}(MA))$ with $\asymp \in \{\sim, \approx, \approxeq\}$, let $\mathcal{E}_{\asymp}(MA)$ denote the set $\mathcal{E}(MA) \cup_{\asymp \in \{\sim,\approx,\approxeq\}} NF_{\asymp}(\mathcal{E}(MA))$. For a given $\mathcal{A} \in \mathcal{E}_{\asymp}(MA)$, $\mathcal{A} = (S, \bar{s}, \Sigma, \mathcal{T})$, we define the mapping $\mathfrak{Z}\colon \mathcal{E}_{\asymp}(MA) \to MA$ as follows: let

$\mathcal{T}|_{\chi} = \{ (s, a, \mu) \in \mathcal{T} \mid \exists \lambda \in \mathbb{R}_{\geq 0}.\, a = \chi_\lambda \}$ and $\mathcal{T}|_{\overline{\chi}} = \mathcal{T} \setminus \mathcal{T}|_{\chi}$. Then, $\mathfrak{Z}\colon \mathcal{A} \mapsto \mathcal{A}' = (S', \bar{s}', \Sigma', \mathcal{T}_I{}', \mathcal{T}_M{}')$ where

- $S' = S$,
- $\bar{s}' = \bar{s}$,
- $\Sigma' = \Sigma \setminus \{ \chi_r \in \Sigma \mid r \in \mathbb{R}_{\geq 0} \}$,
- $\mathcal{T}_I = \mathcal{T}|_{\overline{\chi}} \cup \{ (s, \tau, \delta_s) \mid s \in S' \wedge \mathcal{T}|_{\chi}(s, \cdot\,) = \emptyset \wedge \mathcal{T}|_{\overline{\chi}}(s, \tau) = \emptyset \}$, and
- $\mathcal{T}_M = \{ (s, \lambda, t) \mid \exists r \in \mathbb{R}_{>0}.\, (s, \chi_r, \mu) \in \mathcal{T} \wedge \lambda = r \cdot \mu(t) \}$.

Both $\mathcal{E}$ and $\mathfrak{Z}$ respect unreachable parts of automata. For the normal forms this will not play a role because our quotient mapping is defined to consider only the reachable fragment of an automaton. The following lemma is clear by construction of $\mathcal{E}$ and $\mathfrak{Z}$:

Lemma 31. *$\mathfrak{Z}|_{\mathcal{E}(MA)}$ is injective, so $\mathfrak{Z}|_{\mathcal{E}(MA)}$ is a bijection between $\mathcal{E}(MA)$ and $\mathfrak{Z}(\mathcal{E}(MA))$.*

Weak bisimulations for *MA* is defined by $\mathcal{E}$ using weak distribution bisimulation on *PA*.

Definition 21 ([17])**.** Let $\mathcal{A}_1$, $\mathcal{A}_2$ be *MA*. Then $\mathcal{A}_1 ≋_{MA} \mathcal{A}_2$ if and only if $\mathcal{E}(\mathcal{A}_1) ≋ \mathcal{E}(\mathcal{A}_2)$.

Remark 5. In [16, 17] the symbol $\approx$ (called weak bisimilarity there) is used to denote what is here referred to as $≋_{MA}$ (and called weak distribution bisimilarity) for consistency reasons.

Normal forms for *MAs* are defined as expected (cf. Definition 14): for a given equivalence relation $\asymp$ over *MAs*, we call $NF_{\asymp}\colon MA \to MA$ a *normal form*, if it satisfies for every *MA* $\mathcal{A}$ that $NF_{\asymp}(\mathcal{A}) \asymp \mathcal{A}$, and that for every *MA* $\mathcal{A}'$ it holds that $\mathcal{A} \asymp \mathcal{A}'$ if and only if $NF_{\asymp}(\mathcal{A}) =_{iso} NF_{\asymp}(\mathcal{A}')$.

The *MA* normal form is harvested from the *PA* context via embedding and lifting.

Definition 22. The weak distribution bisimilarity normal form $NF_{≋_{MA}}\colon MA \to MA$ is defined as

$$NF_{≋_{MA}} = \mathcal{E} \circ NF_{≋} \circ \mathfrak{Z}$$

This definition can equally well be instantiated for $\sim$, $\approx$. For the canonicity statement we need to spell out the relevant metrics on *MAs*.

Definition 23. Let $\mathcal{A}_1 = (S_1, \bar{s}_1, \Sigma_1, \mathcal{T}_{I1}, \mathcal{T}_{M1})$ and $\mathcal{A}_2 = (S_2, \bar{s}_2, \Sigma_2, \mathcal{T}_{I2}, \mathcal{T}_{M2})$ be two *MAs*. We write

- $\mathcal{A}_1 \preceq^{|S|}_{≋_{MA}} \mathcal{A}_2$ if $\mathcal{A}_1 ≋_{MA} \mathcal{A}_2$ and $|S_1| \leq |S_2|$,
- $\mathcal{A}_1 \preceq^{|\mathcal{T}_I|}_{≋_{MA}} \mathcal{A}_2$ if $\mathcal{A}_1 ≋_{MA} \mathcal{A}_2$ and $|\mathcal{T}_{I1}| \leq |\mathcal{T}_{I2}|$,

- $\mathcal{A}_1 \precsim^{\|\mathcal{T}_I\|}_{\approx_{MA}} \mathcal{A}_2$ if $\mathcal{A}_1 \approx_{MA} \mathcal{A}_2$ and $\|\mathcal{T}_{I1}\| \leq \|\mathcal{T}_{I2}\|$, and
- $\mathcal{A}_1 \precsim^{|\mathcal{T}_M|}_{\approx_{MA}} \mathcal{A}_2$ if $\mathcal{A}_1 \approx_{MA} \mathcal{A}_2$ and $|\mathcal{T}_{M1}| \leq |\mathcal{T}_{M2}|$.

Canonicity with respect to weak distribution bisimilarity is induced by the *PA* result (Lemma 24). The proof differs slightly since $|\mathcal{T}_M|$-minimality needs to be ensured, taking into consideration Markovian multi-transitions in *MA*.

Lemma 32 (Canonicity of Normal Form)**.** *Let $\mathcal{A}' = NF_{\approx_{MA}}(\mathcal{A})$. Let $(\,\cdot\,,\,\cdot\,,\,\cdot\,)$ be the usual lexicographic ordering. Let $\mathcal{A}_m$ be a $(|S|, \|\mathcal{T}_I\|, |\mathcal{T}_M|)$-minimal automaton with $\bar{s}$ non-naïvely tangible satisfying $\mathcal{A}_m \approx_{MA} \mathcal{A}$. Then $\mathcal{A}' =_{iso} \mathcal{A}_m$.*

Remark 6. In a similar way as indicated in Remark 2, there is a stronger canonicity result when the requirement for an explicit initial state is dropped and initial distributions may be used instead. In this case the elimination of non-naïvely vanishing states can also eliminate a non-naïvely vanishing initial state and the canonicity result stated in Lemma 32 becomes: *Let $\mathcal{A}' = NF_{\approx_{MA}}(\mathcal{A})$. Let $(\,\cdot\,,\,\cdot\,,\,\cdot\,)$ be the usual lexicographic ordering. Let $\mathcal{A}_m$ be a $(|S|, \|\mathcal{T}_I\|, |\mathcal{T}_M|)$-minimal automaton satisfying $\mathcal{A}_m \approx_{MA} \mathcal{A}$. Then $\mathcal{A}' =_{iso} \mathcal{A}_m$.*

As a result, the *MA* counterpart of Theorem 8 can be established.

Theorem 33. *Let $(\,\cdot\,,\,\cdot\,,\,\cdot\,)$ be the usual lexicographic ordering.*

1. ***Minimality:*** *$NF_{\approx_{MA}}(\mathcal{A})$ is $(|S|, |\mathcal{T}_I|, |\mathcal{T}_M|)$-minimal for any $\mathcal{A} \in MA$.*
2. ***Uniqueness of minimals:*** *If $\mathcal{A}, \mathcal{A}' \in MA$ are $(|S|, \|\mathcal{T}_I\|, |\mathcal{T}_M|)$-minimal (and $(|S'|, \|\mathcal{T}_I{}'\|, |\mathcal{T}_M{}'|)$-minimal, respectively) automata with non-naïvely tangible initial states and $\mathcal{A}' \approx_{MA} \mathcal{A}$, then also $\mathcal{A}' =_{iso} \mathcal{A}$,*
3. ***Normal forms:*** *$NF_{\approx_{MA}}$ is uniquely defined up to $=_{iso}$, and is a normal form.*

Remark 7. For claim 1 in contrast to Theorem 8 fanout-minimality will not be achievable when (on a quotient) the only non naïvely vanishing state is the start state (which is assumed to be recurrent). Elimination of such a state would increase the fanout as discussed in Section 8. Claim 2 follows immediately from Lemma 32: both automata are isomorphic to their normal forms, and by the definition of a normal form (cf. Definition 14) the result follows. The precondition of non-naïvely tangible initial states is crucial here. For non-naively vanishing initial states counterexamples to this statement can be constructed, as Fig. 7 demonstrates. Both automata have a non-naïvely vanishing initial state while all other states are non-naïvely tangible. It is easy to see that elimination of Y' (that is: making it transitive) would increase the fanout. Both automata have 8 states, fanout 19 and no Markovian transitions. With the methods of Section 8 one verifies that this is the minimal fanout when using eight states. The two automata are clearly *not* isomorphic, since X can (using a τ transition) reach Y' with probability 0.5 in the automaton 2 but this is clearly not possible in automaton 1.

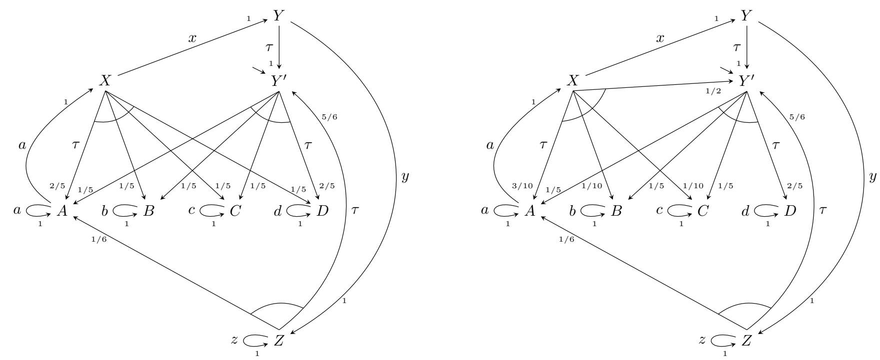

(a) Automaton 1 (b) Automaton 2

Figure 7: Bisimilar but non isomorphic minimal automata with non-naïvely vanishing initial states

Example 5. The example in Figure 8 shows the *MA* weak distribution bisimulation normal form by means of the examples from Figure 1 in [17]. Let us assume that all state symbols with the same color and shape are weakly distribution bisimilar. The mapping $\mathcal{E}$ then converts the automata to *PAs* by transforming the Markovian transitions into χ transitions and by adding χ_0-loops to every stable state without any outgoing Markovian transition.

- The left automaton the normal form mapping reduces to quotienting, that is aggregating the blue states. Note that the left of the blue states is naïvely and trivially vanishing. Every other aspect is already in reduced form. Now the normal form is lifted back to MA by reconstructing the Markovian transition out of the transition labelled with $\chi_{3\lambda}$.
- The middle automaton is already in normal form.
- The right automaton the normal form mapping reduces to elimination of the trivially vanishing non-naïvely vanishing yellow hexagonal state. Every other aspect is already in reduced form.

■

10. Conclusion

In this paper we have successfully answered the question how to compute the minimal, canonical representation of probabilistic automata under strong and weak state-based bisimilarity, together with polynomial time minimization algorithms. We considered as well the state- and transition-minimal, canonical representation of probabilistic and Markov automata under the weak distribution-based bisimilarity; in such a scenario, the minimization algorithms inherit the

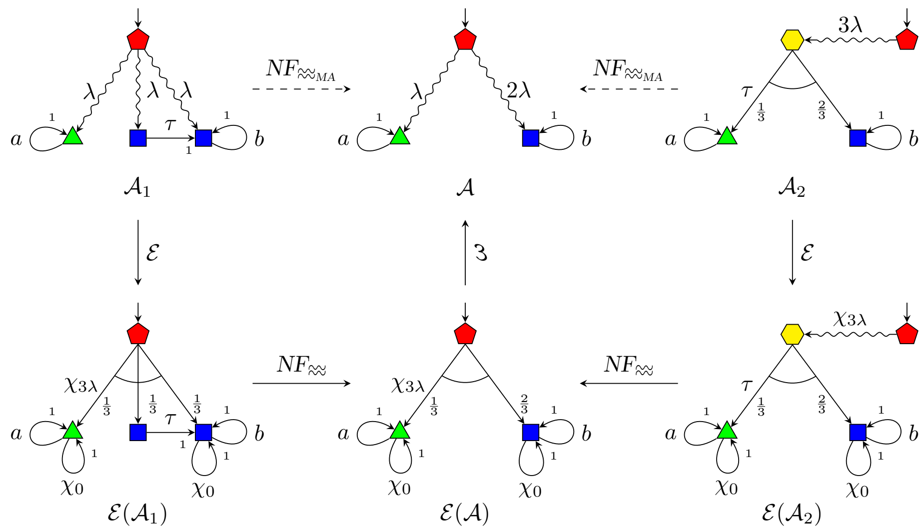

Figure 8: Normal form overview for Markov Automata

exponential complexity of computing the bisimulation relation [14, 37]. We have also proposed an algorithm to compute all fanout minimal representations of a given automaton under the weak distribution-based bisimilarity, but there is no canonical representation for fanout-minimality.

Canonical forms have also appeared in axiomatic treatments of probabilistic calculi [12], but are obtained by adding transitions via saturation, so without aiming for minimality. Figure 9 summarizes what steps are needed to perform the minimization in labelled transition systems (left), probabilistic automata (center) and Markov automata (right). For the *LTS* and *PA* cases, the triplets indicate minimality (✓) or non-minimality (×) with respect to $|S|$, then $|\mathcal{T}|$, then $\|\mathcal{T}\|$; for the *MA* case, the quadruplets indicate minimality (✓) or non-minimality (×) with respect to $|S|$, then $|\mathcal{T}_I|$, then $\|\mathcal{T}_I\|$, and finally $|\mathcal{T}_M|$. For example, a triplet ✓✓× for *PAs* indicates that state and transition numbers are minimal, but transition fanout size can be non-minimal.

The algorithms we developed can be exploited in an effective compositional minimization strategy for *PAs* (or *MDPs*), 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.

Lastly, we give a short overview of some recent results on distribution based bisimulation relations for *PAs*. For Rabin's notion of probabilistic automata [35], which is a special class of *deterministic PAs* where $|\mathcal{T}(s,a)| = 1$ for all states s and actions a, a distribution based bisimulation was proposed in [13]. For such a sub-class of *PAs*, this relation can be decided in polynomial time [13, 42]. This definition is further extended in [22, 26, 28] to handle general *PAs*. The relation between these bisimulations can be found in [21]. Further, in [43], an

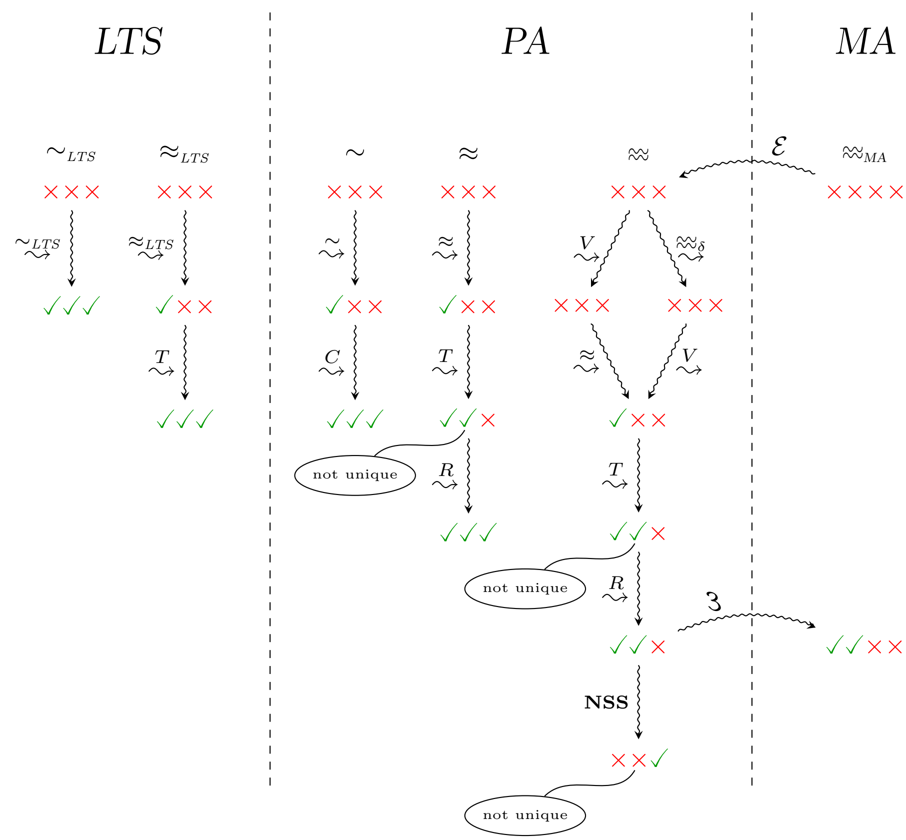

Figure 9: Algorithmic steps in minimal quotient computation.

even weaker notion of distribution based weak bisimilarity has been proposed.

Acknowledgements. This work is supported by the National Natural Science Foundation of China (Grants No. 61532019, 61761136011, 61472473, 61650410658), by the ERC Advanced Investigators Grant 695614 (POWVER), and by the CDZ project CAP (GZ 1023).

References

[1] C. Baier and H. Hermanns. Weak bisimulation for fully probabilistic processes. In *CAV*, volume 1254 of *LNCS*, pages 119–130, 1997.

[2] C. Baier and J.-P. Katoen. *Principles of Model Checking.* The MIT Press, 2008.

[3] B. Barbot, T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Efficient CTMC model checking of linear real-time objectives. In *TACAS*, volume 6605 of *LNCS*, pages 128–142, 2011.

[4] C. Barrett, A. Stump, and C. Tinelli. The SMT-LIB standard: Version 2.0. In *SMT*, 2010.

[5] J. A. Bergstra and J. W. Klop. Algebra of communicating processes with abstraction. *Theor. Comput. Sci.*, 37:77–121, 1985.

[6] J. A. Bergstra and J. W. Klop. A complete inference system for regular processes with silent moves. *Studies in Logic and the Foundations of Mathematics*, 124:21–81, 1987.

[7] S. Cattani and R. Segala. Decision algorithms for probabilistic bisimulation. In *CONCUR*, volume 2421 of *LNCS*, pages 371–385, 2002.

[8] 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*, pages 435–450, 1996.

[9] P. Crouzen and F. Lang. Smart reduction. In *FASE*, volume 6603 of *LNCS*, pages 111–126, 2011.

[10] L. M. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, volume 4963 of *LNCS*, pages 337–340, 2008.

[11] Y. Deng and M. Hennessy. On the semantics of Markov automata. In *ICALP*, volume 6756 of *LNCS*, pages 307–318, 2011.

[12] Y. Deng and M. Hennessy. On the semantics of Markov automata. *I&C*, 222:139–168, 2012.

[13] L. Doyen, T. A. Henzinger, and J.-F. Raskin. Equivalence of Labeled Markov Chains. *Int. J. Found. Comput. Sci.*, 19(3):549–563, 2008.

[14] C. Eisentraut, H. Hermanns, J. Krämer, A. Turrini, and L. Zhang. Deciding bisimilarities on distributions. In *QEST*, volume 8054 of *LNCS*, pages 72–88, 2013.

[15] C. Eisentraut, H. Hermanns, J. Schuster, A. Turrini, and L. Zhang. The quest for minimal quotients for probabilistic automata. In *TACAS*, volume 7795 of *LNCS*, pages 16–31, 2013.

[16] C. Eisentraut, H. Hermanns, and L. Zhang. Concurrency and composition in a stochastic world. In *CONCUR*, volume 6269 of *LNCS*, pages 21–39, 2010.

[17] C. Eisentraut, H. Hermanns, and L. Zhang. On probabilistic automata in continuous time. In *LICS*, pages 342–351, 2010.

[18] C. Eisentraut, H. Hermanns, and L. Zhang. On probabilistic automata in continuous time. Reports of SFB/TR 14 AVACS 62, SFB/TR 14 AVACS, 2010.

[19] J. Eloranta. Minimizing the number of transitions with respect to observation equivalence. *BIT*, 31(4):576–590, 1991.

[20] J. Eloranta, M. Tienari, and A. Valmari. Essential transitions to bisimulation equivalences. *Theor. Comput. Sci.*, 179(1-2):397–419, 1997.

[21] Y. Feng, L. Song, and L. Zhang. Distribution-based bisimulation and bisimulation metric in probabilistic automata. *CoRR*, abs/1512.05027, 2015.

[22] Y. Feng and L. Zhang. When equivalence and bisimulation join forces in probabilistic automata. In *FM*, volume 8442 of *Lecture Notes in Computer Science*, pages 247–262. Springer, 2014.

[23] J.-C. Fernandez and L. Mounier. A tool set for deciding behavioral equivalences. In *CONCUR*, volume 527 of *LNCS*, pages 23–42, 1991.

[24] D. Guck, H. Hatefi, H. Hermanns, J.-P. Katoen, and M. Timmer. Modelling, reduction and analysis of Markov automata. In *QEST*, volume 8054 of *LNCS*, pages 55–71, 2013.

[25] E. M. Hahn, Y. Li, S. Schewe, A. Turrini, and L. Zhang. IscasMc: A web-based probabilistic model checker. In *FM*, volume 8442 of *LNCS*, pages 312–317, 2014.

[26] M. Hennessy. Exploring probabilistic bisimulations, part I. *Formal Asp. Comput.*, 24(4-6):749–768, 2012.

[27] 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.

[28] H. Hermanns, J. Krcál, and J. Kretínský. Probabilistic bisimulation: Naturally on distributions. In *CONCUR 2014*, volume 8704 of *Lecture Notes in Computer Science*, pages 249–265. Springer, 2014.

[29] B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *LICS*, pages 266–277, 1991.

[30] P. C. Kanellakis and S. A. Smolka. CCS expressions, finite state processes, and three problems of equivalence. *I&C*, 86(1):43–68, 1990.

[31] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In *CAV*, volume 6806 of *LNCS*, pages 585–591, 2011.

[32] N. A. Lynch, R. Segala, and F. W. Vaandrager. Observing branching structure through probabilistic contexts. *SIAM J. on Computing*, 37(4):977–1013, 2007.

[33] R. Milner. *Communication and Concurrency.* Prentice-Hall International, Englewood Cliffs, 1989.

[34] R. Paige and R. E. Tarjan. Three partition refinement algorithms. *SIAM J. on Computing*, 16(6):973–989, 1987.

[35] M. Rabin. Probabilistic automata. *Information and Control*, 6(3):230–245, 1963.

[36] J. Schuster. *Towards faster numerical solution of Continuous Time Markov Chains stored by symbolic data structures.* PhD thesis, Universität der Bundeswehr München, 2011.

[37] J. Schuster and M. Siegle. Markov automata: Deciding weak bisimulation by means of "non-naïvely" vanishing states. *I&C*, 237:151–173, 2014.

[38] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems.* PhD thesis, MIT, 1995.

[39] R. Segala. Probability and nondeterminism in operational models of concurrency. In *CONCUR*, volume 4137 of *LNCS*, pages 64–78, 2006.

[40] R. Segala and N. A. Lynch. Probabilistic simulations for probabilistic processes. *Nordic J. Computing*, 2(2):250–273, 1995.

[41] A. Turrini and H. Hermanns. Polynomial time decision algorithms for probabilistic automata. *Information and Computation*, 244:134–171, 2015.

[42] W. Tzeng. A polynomial-time algorithm for the equivalence of probabilistic automata. *SIAM Journal on Computing*, 21(2):216–227, 1992.

[43] L. Zhang, P. Yang, L. Song, H. Hermanns, C. Eisentraut, D. N. Jansen, and J. C. Godskesen. Probabilistic bisimulation for realistic schedulers. *Acta Inf.*, 2018. Accepted for publication.