



ELSEVIER

Contents lists available at ScienceDirect

Science of Computer Programming

www.elsevier.com/locate/scico


On probabilistic term rewriting

 Martin Avanzini^{a,*}, Ugo Dal Lago^{a,b,*}, Akihisa Yamada^{c,*}
^a INRIA Sophia Antipolis, France^b University of Bologna, Italy^c National Institute of Informatics, Japan

ARTICLE INFO

Article history:

Received 14 November 2018

Received in revised form 8 October 2019

Accepted 13 October 2019

Available online 21 October 2019

Keywords:

Probabilistic abstract reduction systems

Probabilistic term rewriting

Almost sure termination

Interpretation method

ABSTRACT

We study the termination problem for probabilistic term rewrite systems. We prove that the interpretation method is sound and complete for a strengthening of positive almost sure termination, when abstract reduction systems and term rewrite systems are considered. Two instances of the interpretation method—polynomial and matrix interpretations—are analyzed and shown to capture interesting and nontrivial examples when automated. We capture probabilistic computation in a novel way by means of multidistribution reduction sequences, thus accounting for both the nondeterminism in the choice of the redex and the probabilism intrinsic in firing each rule.

© 2019 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Interactions between computer science and probability theory are pervasive and extremely useful to the first discipline. Probability theory indeed offers models that enable *abstraction*, but it also suggests a new model of *computation*, like in randomized algorithmics [32] or cryptography [19]. All this has stimulated the study of probabilistic computational models and programming languages: probabilistic variations on well-known models like automata [12,35], Turing machines [37,17], and the λ -calculus [36,23] are known from the early days of theoretical computer science.

The simplest way probabilistic choice can be made available in programming languages consists in endowing the language of programs with an operator modeling sampling from (one or many) distributions. Fair, binary, probabilistic choice is for example perfectly sufficient to get universality if the underlying programming language is itself universal (e.g., see [11]).

Term rewriting [38] is a well-studied model of computation where no probabilistic behavior is involved. It provides a faithful abstraction of pure functional programming which is, up to a certain extent, also adequate for modeling higher-order parameter passing [27]. What is peculiar in term rewriting is that, in principle, rule selection turns reduction into a potentially nondeterministic process. The following question is then a natural one: is there a way to generalize term rewriting to a fully-fledged *probabilistic* model of computation? Actually, not much is known about probabilistic term rewriting: we are only aware of the definitions due to Agha et al. [1] and Bournez and Garnier [6]. We base our work on the latter, where probabilistic rewriting is captured as a Markov decision process; rule selection remains nondeterministic, but each rule can have one of many possible outcomes, each with its own probability. Rewriting thus becomes a process in which both nondeterministic and probabilistic aspects are present and intermingled. When firing a rule, the reduction process implicitly samples from a distribution, much in the same way as when performing binary probabilistic choice in one of the models mentioned above.

* Corresponding authors.

E-mail addresses: martin.avanzini@inria.fr (M. Avanzini), ugo.dallago@unibo.it (U. Dal Lago), akihisayamada@nii.ac.jp (A. Yamada).

In this paper, we first define a new, simple framework for discrete probabilistic reduction systems, which properly generalizes standard abstract reduction systems [38] (Section 3). In particular, what plays the role of a reduction sequence, usually a (possibly infinite) sequence $a_1 \rightarrow a_2 \rightarrow \dots$ of states, is a sequence $\mu_1 \rightarrow^{\mathcal{M}} \mu_2 \rightarrow^{\mathcal{M}} \dots$ of (multi)distributions over the set of states. A multidistribution is not merely a distribution, and this is crucial to appropriately account for both the probabilistic behavior of each rule and the nondeterminism in rule selection. Such a correspondence does not exist in Bournez and Garnier's framework, as nondeterminism has to be resolved by a strategy, in order to define reduction sequences. Consequently, our framework is not only conceptually simpler, it is closer to ordinary rewriting. Indeed, our reduction relation is simply a reduction system over multidistributions, and can be studied without particular knowledge of probability theory. Nevertheless, the two frameworks turn out to be equiexpressive (Section 3.4).

On top of this framework, we then introduce abstract embeddings and more concrete probabilistic ranking functions, sound and complete methods for proving strong almost sure termination, a strengthening of positive almost sure termination [6]. We moreover show that ranking functions provide bounds on expected runtimes (Section 3.3).

This paper's main contribution, then, is the definition of a simple framework for probabilistic term rewrite systems as an example of this abstract framework (Section 4). Our main aim is studying whether any of the well-known techniques for termination of term rewrite systems can be generalized to the probabilistic setting, and whether they can be automated. We give positive answers to these two questions, by describing how polynomial and matrix interpretations can indeed be turned into instances of probabilistic ranking functions, thus generalizing them to the more general context of probabilistic term rewriting. We moreover implement these new techniques into the termination tool NaTT [39] (Section 5).

This paper is revised version of the conference paper [3]. Apart from giving all missing proofs and more detailed examples and explanations, we generalized the definition of probabilistic reduction systems so that reducts may be sampled from infinite distributions. Another notable change lies in the reformulation of probabilistic ranking functions in terms of embeddings, leading also to a more elegant formulation of barycentric algebras.

2. Related work

Termination is a crucial property of programs, and has been widely studied in term rewriting. Tools checking and certifying termination of term rewrite systems are nowadays capable of implementing tens of different techniques, and can prove termination of a wide class of term rewrite systems, although the underlying verification problem is well known to be undecidable [38].

Termination remains an interesting and desirable property in a probabilistic setting, e.g., in probabilistic programming [20] where inference algorithms often rely on the underlying program to terminate. But what does termination mean when systems become probabilistic? If one wants to stick to a qualitative definition, almost-sure termination is a well-known answer: a probabilistic computation is said to almost surely terminate iff non-termination occurs with null probability. One could even require positive almost-sure termination, which asks the expected time to termination to be finite. Recursion-theoretically, checking (positive) almost-sure termination is harder than checking termination of non-probabilistic programs, where termination is at least recursively enumerable, although undecidable: in a universal probabilistic imperative programming language, the termination questions for almost-sure and positive almost-sure termination on a single input are already Π_2^0 and Σ_2^0 complete, respectively [24].

Many sound verification methodologies for probabilistic termination have recently been introduced (see, e.g., [6,7,18,14,10]). In particular, the use of ranking martingales has turned out to be quite successful when the analyzed program is imperative, and thus does not have an intricate recursive structure. When the latter holds, techniques akin to sized types have been shown to be applicable [26]. More recently, Ngo et al. [33] implemented the `ert`-calculus of Kaminski et al. [25] for reasoning about the expected runtime of imperative integer programs, showing promising results.¹

Finally, as already mentioned, the current work can be seen as stemming from the work by Bournez et al. [8,6,7]. The added value compared to their work are first of all the notion of multidistribution as a way to give an instantaneous description of the state of the underlying system which exhibits both nondeterministic and probabilistic features. Our completeness result can be seen as a correction to their claim [6, Theorem 3], which was already refuted [14]. In fact, incompleteness claims in [14] also contradict our result, but their counterexample is invalid as part of the reduction steps are not counted.² Moreover, an interpretation method inspired by ranking functions is made more general here, this way accommodating not only interpretations over the real numbers, but also interpretations over vectors, in the spirit of matrix interpretations. Finally, we provide an automation of polynomial and matrix interpretation inference, whereas nothing about implementation was presented in Bournez and Garnier's work.

Some of our results have been independently obtained by Fu and Chatterjee [15], most notably, Theorem 2. Their notion of bounded termination is what we call strong almost-sure termination in this work. The main peculiarity of our work is that we focus on probabilistic term rewriting, while their focus is on procedural programs with probabilistic sampling. Besides, the authors also investigate lower bounds on expected runtimes and tail-probabilities, topics not covered in our work.

¹ While it is possible to translate such programs into probabilistic term rewrite systems [4], our method will not capture typical termination arguments on probabilistic integer programs, since properties of integers will be unavailable after the encoding.

² We thank Luis María Ferrer Fioriti for this analysis.

3. Probabilistic abstract reduction systems

An *abstract reduction system (ARS)* on a set A is a binary relation $\rightarrow \subseteq A \times A$. Having $a \rightarrow b$ means that a reduces to b in one step, or b is a one-step reduct of a . Bournez and Garnier [6] extended the ARS formalism to probabilistic computations by allowing reducts to be sampled from a (*probability*) *distribution*. Throughout the paper, we denote by $\mathbb{R}_{\geq 0}$ the set of non-negative reals.

Definition 1 (Distribution). A *distribution* on a (countable) set A is a mapping $d : A \rightarrow \mathbb{R}_{\geq 0}$ assigning to each $a \in A$ a probability $d(a)$, such that $\sum_{a \in A} d(a) = 1$. We write $\mathcal{D}(A)$ for the set of distributions over A . The *support* of a distribution d is the set $\text{Supp}(d) := \{a \in A \mid d(a) > 0\}$. We may denote a distribution d by $\{d(a) : a \mid a \in \text{Supp}(d)\}$, or $\{d(a_1) : a_1, \dots, d(a_n) : a_n\}$ when $\text{Supp}(d)$ is the finite set $\{a_1, \dots, a_n\}$.

A *probabilistic ARS (PARS)* in terms of Bournez and Garnier [6] is defined as a relation $\rightarrow \subseteq A \times \mathcal{D}(A)$, and having $a \rightarrow d$ means that d is the distribution of the one-step reducts of a , or a reduces to b with probability $p = d(b)$, in notation $p : a \rightarrow b$. We extend this notation to sequences, also called *runs* below. For instance, if further $q : b \rightarrow c$, we write $p \cdot q : a \rightarrow b \rightarrow c$. Notice that since \rightarrow is a relation, the PARS may indeed specify more than one distribution of one-step reductions for $a \in A$, i.e., $a \rightarrow d_1$ and $a \rightarrow d_2$ with $d_1 \neq d_2$. The distribution of one-step reducts of a is nondeterministically chosen from d_1 and d_2 in this case, and \rightarrow is called *nondeterministic*.

Example 1 (Random walk). A random walk over \mathbb{N} with *bias* probability p is modeled by the probabilistic ARS $\xrightarrow{\mathcal{W}_p}$ such that

$$n + 1 \xrightarrow{\mathcal{W}_p} \{p : n; 1 - p : n + 2\} \quad \text{for all } n \in \mathbb{N}.$$

For a PARS \rightarrow , we aim at defining a reduction relation $\xrightarrow{\mathcal{D}}$, as an ARS on distributions. Taking Example 1, we would like to have

$$\{1 : 1\} \xrightarrow{\mathcal{W}_{\frac{1}{2}}} \left\{ \frac{1}{2} : 0; \frac{1}{2} : 2 \right\},$$

meaning that the distribution of the one-step reducts of 1, or more precisely the singleton distribution assigning probability one to 1, is $\{\frac{1}{2} : 0; \frac{1}{2} : 2\}$. Continuing the reduction, what should the distribution of two-step reducts of 1 be? Actually, it cannot be a distribution (on A): since 0 cannot be reduced, by probability $\frac{1}{2}$ we have no two-step reduct of 1. Hence we consider *subdistributions*, i.e., generalizations of distributions where probabilities may sum up to less than one, allowing

$$\{1 : 1\} \xrightarrow{\mathcal{W}_{\frac{1}{2}}} \left\{ \frac{1}{2} : 0; \frac{1}{2} : 2 \right\} \xrightarrow{\mathcal{W}_{\frac{1}{2}}} \left\{ \frac{1}{4} : 1; \frac{1}{4} : 3 \right\}.$$

Further continuing the reduction, one would expect $\{\frac{1}{8} : 0; \frac{1}{4} : 2; \frac{1}{8} : 4\}$ as the next step, but note that a half of the probability $\frac{1}{4}$ of 2 is due to the run $\frac{1}{8} : 1 \rightarrow 2 \rightarrow 1 \rightarrow 2$, and the other half is due to the run $\frac{1}{8} : 1 \rightarrow 2 \rightarrow 3 \rightarrow 2$. It turns out that, in the presence of nondeterminism, we should distinguish the two possibilities.

Example 2. Consider the PARS $\xrightarrow{\mathcal{N}}$ such that

$$\begin{array}{lll} a \xrightarrow{\mathcal{N}} \left\{ \frac{1}{2} : b_1; \frac{1}{2} : b_2 \right\} & b_1 \xrightarrow{\mathcal{N}} \{1 : c\} & c \xrightarrow{\mathcal{N}} \{1 : d_1\} \\ & b_2 \xrightarrow{\mathcal{N}} \{1 : c\} & c \xrightarrow{\mathcal{N}} \{1 : d_2\}. \end{array}$$

Reducing a twice always yields c , so the two-step reduct of a seen as a distribution is $\{1 : c\}$. More precisely, there are two runs from a to c , namely, $\frac{1}{2} : a \rightarrow b_1 \rightarrow c$ and $\frac{1}{2} : a \rightarrow b_2 \rightarrow c$. Each of them can be nondeterministically continued to d_1 and d_2 , so the distribution of three-step reducts of a is the nondeterministic choice among $\{1 : d_1\}$, $\{\frac{1}{2} : d_1, \frac{1}{2} : d_2\}$, $\{1 : d_2\}$. On the other hand, whereas it is obvious that $\{1 : c\}$ should reduce to $\{1 : d_1\}$ or $\{1 : d_2\}$, obtaining the third choice $\{\frac{1}{2} : d_1, \frac{1}{2} : d_2\}$ would require the reduction relation $\xrightarrow{\mathcal{D}}$ to be defined in a non-local manner.

To overcome this problem, we refine subdistributions to *multidistributions*, where a single element can be associated with *more than one* probability. With respect to the nondeterministic PARS $\xrightarrow{\mathcal{N}}$, this allows us to model that a reduces in two steps to c on two distinct runs, each with probability $\frac{1}{2}$, and continuing this reduction may yield any of the above three-step reducts.

3.1. Probabilistic ARSs and multidistribution reductions

We model multidistribution by a special form of (possibly infinite) multisets.

Definition 2 (Multiset). A multiset over a set A is a mapping $M : A \rightarrow \mathbb{N}$. The union $\biguplus_{i \in I} M_i$ of countably many multisets M_i is defined by

$$\left(\biguplus_{i \in I} M_i \right) (a) := \sum_{i \in I} M_i(a),$$

which forms a multiset if and only if $\sum_{i \in I} M_i(a)$ is finite for every $a \in A$. The sum of a multiset M with respect to $f : A \rightarrow \mathbb{R}$ is defined by

$$\sum_{a \in M} f(a) := \sum_{a \in A} M(a) \cdot f(a).$$

The *submultiset* relation is defined by $M \subseteq N \iff \forall a \in A. M(a) \leq N(a)$.

We use set-like notations for multisets: \emptyset denotes the *empty* multiset $\emptyset(a) := 0$, $\{\{a_i \mid i \in I\}\}$ is the multiset M with $M(a) = |\{i \in I \mid a_i = a\}|$, and $\{\{a_1, \dots, a_n\}\}$ is its special case where $I = \{1, \dots, n\}$ is finite. The following lemma is an easy consequence of the definition.

Lemma 1. For a family $(\{\{a_{i,j} \mid j \in J_i\}\})_{i \in I}$ of multisets we have

$$\biguplus_{i \in I} \{\{a_{i,j} \mid j \in J_i\}\} = \{\{a_{i,j} \mid i \in I, j \in J_i\}\}.$$

Definition 3 (Multidistributions). A (sub)multidistribution on a set A is a multiset μ of pairs of $a \in A$ and $0 < p \leq 1$, written $p : a$, satisfying

$$|\mu| := \sum_{p:a \in \mu} p \leq 1.$$

We call μ a *proper multidistribution* if $|\mu| = 1$. We denote the set of multidistributions on A by $\mathcal{M}_{\leq 1}(A)$, and proper ones by $\mathcal{M}(A)$.

By an abuse of notation, we identify a subdistribution $\{p_i : a_i \mid i \in I\}$ with the multidistribution $\{\{p_i : a_i \mid i \in I\}\}$. We often lift a function $f : A \rightarrow B$ to $f : \mathcal{M}_{\leq 1}(A) \rightarrow \mathcal{M}_{\leq 1}(B)$ as follows:

$$f(\{\{p_i : a_i \mid i \in I\}\}) := \{\{p_i : f(a_i) \mid i \in I\}\}.$$

Definition 4 (PARS). A *probabilistic ARS (PARS)* over a set A is a (typically infinite) set $\rightarrow \subseteq A \times \mathcal{M}(A)$. An object $a \in A$ is called *terminal*, or a *normal form* in \rightarrow , if there is no μ with $a \rightarrow \mu$. With NF_{\rightarrow} we denote the set of normal forms in \rightarrow .

We remark that in contrast to [3,6], we allow proper multidistributions rather than distributions as right-hand sides, e.g., $a \rightarrow \{\{\frac{1}{2} : a, \frac{1}{2} : a\}\}$ is permitted in our setting.

Now we lift a PARS $\rightarrow \subseteq A \times \mathcal{M}(A)$ to the reduction relation $\xrightarrow{\mathcal{M}} \subseteq \mathcal{M}_{\leq 1}(A) \times \mathcal{M}_{\leq 1}(A)$, an ARS over multidistributions. To this end, we need some basic operations on multidistributions. The *scalar multiplication* of a multidistribution is defined by

$$p \cdot \{\{q_i : a_i \mid i \in I\}\} := \{\{p \cdot q_i : a_i \mid i \in I\}\},$$

which is also a multidistribution if $0 < p \leq 1$. More generally, multidistributions are closed under *subconvex combinations*:

Lemma 2. For families $(\mu_i)_{i \in I}$ of multidistributions and $(p_i)_{i \in I}$ of positive real numbers, $\biguplus_{i \in I} p_i \cdot \mu_i$ is also a multidistribution if $\sum_{i \in I} p_i \leq 1$. More precisely, $|\biguplus_{i \in I} p_i \cdot \mu_i| = \sum_{i \in I} p_i \cdot |\mu_i| \leq 1$.

Proof. Let $\mu_i = \{q_{i,j} : a_{i,j} \mid j \in J_i\}$ for each $i \in I$. Using Lemma 1 we have

$$\begin{aligned} \left| \bigoplus_{i \in I} p_i \cdot \mu_i \right| &= |\{p_i \cdot q_{i,j} : a_{i,j} \mid i \in I, j \in J_i\}| \\ &= \sum_{i \in I} \sum_{j \in J_i} p_i \cdot q_{i,j} \\ &= \sum_{i \in I} p_i \cdot \sum_{j \in J_i} q_{i,j} = \sum_{i \in I} p_i \cdot |\mu_i|. \end{aligned}$$

Since $|\mu_i| \leq 1$, the claim then follows from the assumption $\sum_{i \in I} p_i \leq 1$. \square

Definition 5 (Probabilistic Reduction). Given a PARS $\rightarrow \subseteq A \times \mathcal{M}(A)$, we define the *probabilistic reduction relation* $\xrightarrow{\mathcal{M}} \subseteq \mathcal{M}_{\leq 1}(A) \times \mathcal{M}_{\leq 1}(A)$ as follows:

$$\frac{a \in \text{NF}_{\rightarrow}}{\{\{1 : a\}\} \xrightarrow{\mathcal{M}} \emptyset} \quad \frac{a \rightarrow \mu}{\{\{1 : a\}\} \xrightarrow{\mathcal{M}} \mu} \quad \frac{\forall i \in I. \mu_i \xrightarrow{\mathcal{M}} \nu_i}{\bigoplus_{i \in I} p_i \cdot \mu_i \xrightarrow{\mathcal{M}} \bigoplus_{i \in I} p_i \cdot \nu_i}$$

In the last rule, I is an arbitrary (possibly empty) countable set, $p_i > 0$ for every $i \in I$, and $\sum_{i \in I} p_i \leq 1$. We denote by $\text{red}_{\rightarrow}(\mu)$ the set of all possible reduction sequences from μ , i.e., $(\mu_n)_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}(\mu)$ iff $\mu_0 = \mu$ and $\mu_n \xrightarrow{\mathcal{M}} \mu_{n+1}$ for any $n \in \mathbb{N}$. We overload this notation for $a \in A$ and denote by $\text{red}_{\rightarrow}(a)$ the set of all possible reduction sequences from $\{\{1 : a\}\}$.

In essence, the reduction relation $\xrightarrow{\mathcal{M}}$ induced by a PARS \rightarrow is given by a pointwise extension of \rightarrow to multidistributions, removing normal forms along reductions. We abbreviate $\xrightarrow{\mathcal{M}}$ by \rightarrow when it is clear from the context.

Let us illustrate the definition on the two examples from above.

Example 3 (Example 1, Revisited). The informal reduction of the PARS $\xrightarrow{\mathcal{W}_{\frac{1}{2}}}$ on distributions as outlined above is given by the following reduction sequence on multidistributions:

$$\begin{aligned} \{\{1 : 1\}\} &\xrightarrow{\mathcal{W}_{\frac{1}{2}}} \{\{\frac{1}{2} : 0; \frac{1}{2} : 2\}\} \xrightarrow{\mathcal{W}_{\frac{1}{2}}} \{\{\frac{1}{4} : 1; \frac{1}{4} : 3\}\} \\ &\xrightarrow{\mathcal{W}_{\frac{1}{2}}} \{\{\frac{1}{8} : 0; \frac{1}{8} : 2; \frac{1}{8} : 2; \frac{1}{8} : 4\}\} \xrightarrow{\mathcal{W}_{\frac{1}{2}}} \dots \end{aligned}$$

Here, the first step follows from the second rule of Definition 5. The second is obtained by a combination of all three rules:

$$\frac{\frac{0 \in \text{NF}_{\rightarrow}}{\{\{1 : 0\}\} \xrightarrow{\mathcal{W}_{\frac{1}{2}}} \emptyset} \quad \frac{2 \xrightarrow{\mathcal{W}_{\frac{1}{2}}} \{\{\frac{1}{2} : 1; \frac{1}{2} : 3\}\}}{\{\{1 : 2\}\} \xrightarrow{\mathcal{W}_{\frac{1}{2}}} \{\{\frac{1}{2} : 1; \frac{1}{2} : 3\}\}}}{\{\{\frac{1}{2} : 0; \frac{1}{2} : 2\}\} \xrightarrow{\mathcal{W}_{\frac{1}{2}}} \{\{\frac{1}{4} : 1; \frac{1}{4} : 3\}\}}$$

The third step is derived similarly.

In Example 3, the third reduct $\{\{\frac{1}{8} : 0; \frac{1}{8} : 2; \frac{1}{8} : 2; \frac{1}{8} : 4\}\}$ records that 2 derives from 1 via two distinct paths. In consequence, this resolves the issues indicated in Example 2 when dealing with nondeterministic systems.

Example 4 (Example 2, Revisited). Consider again the nondeterministic PARS $\xrightarrow{\mathcal{N}}$. Besides the two reductions yielding d_1 and d_2 outlined above, we have the reduction

$$\{\{1 : a\}\} \xrightarrow{\mathcal{N}} \{\{\frac{1}{2} : b_1, \frac{1}{2} : b_2\}\} \xrightarrow{\mathcal{N}} \{\{\frac{1}{2} : c, \frac{1}{2} : c\}\} \xrightarrow{\mathcal{N}} \{\{\frac{1}{2} : d_1, \frac{1}{2} : d_2\}\}.$$

The final step is possible because $\{\{\frac{1}{2} : c, \frac{1}{2} : c\}\}$ is not collapsed to $\{\{1 : c\}\}$.

The definition of the reduction relation $\xrightarrow{\mathcal{M}}$ has several consequences. First, any multidistribution μ is “reducible” in $\xrightarrow{\mathcal{M}}$, even if μ consists only of terminal objects. Second, if μ is the subconvex combination of multidistributions v_i ($i \in I$), then any reduct of μ is given by the subconvex combination of some reducts of v_i . Further, this observation carries over to reduction sequences. Let us extend subconvex combination pointwise to sequences of multidistributions, i.e., $\vec{\mu} = \bigoplus_{i \in I} p_i \cdot (\vec{v}_i)_{i \in I}$ if $\mu_n = \bigoplus_{i \in I} p_i \cdot (v_{i,n})_{i \in I}$ for all $n \in \mathbb{N}$, where $\vec{\mu} = (\mu_n)_{n \in \mathbb{N}}$ and $\vec{v}_i = (v_{i,n})_{n \in \mathbb{N}}$ for all $i \in I$.

Lemma 3. $\vec{\mu} \in \text{red}_{\rightarrow} (\bigoplus_{i \in I} p_i \cdot v_i)$ if and only if there exists $\vec{\rho}_i \in \text{red}_{\rightarrow} (v_i)$ for each $i \in I$ such that $\vec{\mu} = \bigoplus_{i \in I} p_i \cdot \vec{\rho}_i$.

Proof. The “if” direction is obvious. For the “only if” direction, consider an arbitrary $\vec{\mu} \in \text{red}_{\rightarrow} (\bigoplus_{i \in I} p_i \cdot v_i)$. We inductively construct $\rho_{i,n}$ so that $(\rho_{i,n})_{n \in \mathbb{N}} \in \text{red}_{\rightarrow} (v_i)$. For the base case we define $\rho_{i,0} = v_i$. For the inductive case, assume by induction hypothesis that $\mu_n = \bigoplus_{i \in I} p_i \cdot \rho_{i,n} \rightarrow \mu_{n+1}$. An induction on the derivation of this step yields $\rho_{i,n+1}$ such that $\rho_{i,n} \rightarrow \rho_{i,n+1}$ for every $i \in I$. This concludes the proof. \square

Finally, we remark that PARSs constitute a generalization of ARSs: an ARS can be seen as a PARS whose right-hand sides are *singleton distributions*, i.e., $\{\{1 : b\}\}$ for some b , and the non-probabilistic reduction is simulated via the relation $\{\{1 : \cdot\}\} \rightarrow \{\{1 : \cdot\}\}$. Only a little care is needed as $\{\{1 : a\}\}$ reduces to \emptyset when a is terminal.

Proposition 1. Let \rightarrow be an ARS and define \leftrightarrow as the smallest PARS such that $a \leftrightarrow \{\{1 : b\}\}$ if $a \rightarrow b$. Then $\{\{1 : a\}\} \leftrightarrow \mu$ iff either $a \rightarrow b$ and $\mu = \{\{1 : b\}\}$ for some b , or a is a normal form in \rightarrow and $\mu = \emptyset$.

Proof. For $\{\{1 : a\}\} \leftrightarrow \mu$ only the first two rules of Definition 5 are effective. Then the claim directly follows. \square

3.2. Notions of probabilistic termination

An ARS \rightarrow is called *terminating* if it does not give rise to an infinite sequence $a_1 \rightarrow a_2 \rightarrow \dots$. In a probabilistic setting, our interest is whether infinite sequences occur with non-zero probability or not. This notion is defined using our notation as follows.

Definition 6 (Almost Sure Termination; AST). A PARS \rightarrow is said to be *almost surely terminating (AST)* if for any reduction sequence it holds that $\lim_{n \rightarrow \infty} |\mu_n| = 0$.

Intuitively, $|\mu_n|$ is the probability of having a reduct in the n th step, so its tendency towards zero indicates that infinite runs occur with zero probability.

Example 5 (Example 1, Revisited). The PARS $\xrightarrow[\mathcal{W}_p]$ is AST for $p \geq \frac{1}{2}$, whereas it is not for $p < \frac{1}{2}$ (cf. [31]). Note that although $\xrightarrow[\mathcal{W}_{\frac{1}{2}}]$ is AST, the expected number of steps needed to reach a terminal is infinite.

As motivated by Bournez and Garnier [6], it is also interesting and important to further ensure that the *expected length of runs* is finite. We define the notion by means of the following concise definitions. In Section 3.4 we show that our definition is equivalent to the corresponding notion defined via stochastic processes.

Definition 7 (Expected Derivation Length). Let \rightarrow be a PARS and $\vec{\mu} = (\mu_n)_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}(\mu)$. We define the *expected derivation length* $\text{edl}(\vec{\mu}) \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ in this multidistribution reduction sequence $\vec{\mu}$ by

$$\text{edl}(\vec{\mu}) := \sum_{n \geq 1} |\mu_n|.$$

Bournez and Garnier [6] introduced the notion of *positive almost sure termination (PAST)*, which can be formulated as follows.

Definition 8 (Positive AST; PAST). A PARS \rightarrow is said to be *positively almost surely terminating (PAST)* iff for any reduction $\vec{\mu} \in \text{red}_{\rightarrow}(a)$ starting from any $a \in A$, $\text{edl}(\vec{\mu})$ is finite.

One should be careful to notice that the above definition does not ensure a bound on the expected length of *all runs* starting from a given $a \in A$. This phenomenon is already visible in the non-probabilistic setting.

Example 6. Consider the ARS $\xrightarrow{\omega}$ over $\mathbb{N} \cup \{a\}$ given by

$$a \xrightarrow{\omega} n \quad n + 1 \xrightarrow{\omega} n$$

for all $n \in \mathbb{N}$. Every reduction in $\xrightarrow{\omega}$ is finite, thus $\xrightarrow{\omega}$ seen as a PARS is PAST. However, there is no upper bound on the length of reductions starting from a .

In this example, although $\xrightarrow{\omega}$ is terminating, it is not *finitely branching*: a admits infinitely many one-step reducts. In the non-probabilistic setting, a finitely branching terminating ARS admits a bound on the length of derivation for each starting element. In the probabilistic setting, however, PAST does not ensure such bounds even if the system is finitely branching³:

Example 7. Consider the PARS $\xrightarrow{\infty}$ over $\mathbb{N} \cup \{a_n \mid n \in \mathbb{N}\}$ defined by

$$a_n \xrightarrow{\infty} \left\{ \left\{ \frac{1}{2} : a_{n+1}; \frac{1}{2} : 0 \right\} \right\} \quad a_n \xrightarrow{\infty} \left\{ \{1 : 2^n \cdot n\} \right\} \quad n + 1 \xrightarrow{\infty} \left\{ \{1 : n\} \right\}$$

for all $n \in \mathbb{N}$. Obviously $\xrightarrow{\infty}$ is finitely branching. We show that $\xrightarrow{\infty}$ is PAST as follows. Observe that every sequence $\vec{\mu} \in \text{red}_{\xrightarrow{\infty}}(a_n)$ has one of the following two forms:

1. The first rule is fired infinitely often:

$$\vec{\mu} = \left\{ \{1 : a_n\} \right\} \xrightarrow{\infty} \left\{ \left\{ \frac{1}{2} : a_{n+1}; \frac{1}{2} : 0 \right\} \right\} \xrightarrow{\infty} \left\{ \left\{ \frac{1}{4} : a_{n+2}; \frac{1}{4} : 0 \right\} \right\} \xrightarrow{\infty} \dots$$

In this case, $\text{edl}(\vec{\mu}) = 1 + \frac{1}{2} + \frac{1}{4} + \dots = \sum_{n \in \mathbb{N}} \frac{1}{2^n} = 2$.

2. The second rule is fired after m applications of the first:

$$\vec{\mu} = \left\{ \{1 : a_n\} \right\} \xrightarrow{\infty}^m \left\{ \left\{ \frac{1}{2^m} : a_{n+m}; \frac{1}{2^m} : 0 \right\} \right\} \xrightarrow{\infty} \left\{ \left\{ \frac{1}{2^m} : k \right\} \right\} \xrightarrow{\infty} \left\{ \left\{ \frac{1}{2^m} : 0 \right\} \right\},$$

where $k = 2^{n+m} \cdot (n + m)$. In this case, $\text{edl}(\vec{\mu}) = 1 + \frac{1}{2} + \dots + \frac{1}{2^m} + \frac{1}{2^m} \cdot k \geq m$.

The second class of derivations witnesses that, given a_n , there is no bound $b_n < \infty$ such that $\text{edl}(\vec{\mu}) \leq b_n$ for any $\vec{\mu} \in \text{red}_{\xrightarrow{\infty}}(a_n)$.

Therefore we introduce a stronger notion, which ensures a bound on the expected derivation length for each starting element. The definition is based on a natural extension of *derivation height* [22] from complexity analysis of term rewriting.

Definition 9 (Strong AST; SAST). For a PARS $\rightarrow \subseteq A \times \mathcal{M}(A)$, the *expected derivation height* $\text{edh}_{\rightarrow}(\mu) \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ of $\mu \in \mathcal{M}_{\leq 1}(A)$ is defined by

$$\text{edh}_{\rightarrow}(\mu) := \sup \left\{ \text{edl}(\vec{v}) \mid \vec{v} \in \text{red}_{\rightarrow}(\mu) \right\}.$$

For $a \in A$, we write $\text{edh}_{\rightarrow}(a)$ for $\text{edh}_{\rightarrow}(\{1 : a\})$. We say \rightarrow is *strongly almost surely terminating* (SAST) iff $\text{edh}_{\rightarrow}(a)$ of every $a \in A$ is finite.

Notice that on the class of *deterministic PARSs*, i.e., when $a \rightarrow \mu_1$ and $a \rightarrow \mu_2$ implies $\mu_1 = \mu_2$, SAST coincides with PAST. In the nondeterministic case, on the other hand, the notions of PAST and SAST do not coincide. SAST guarantees that the expected length of reductions is bounded, taking a demonic view on nondeterminism. As highlighted in Example 7 such a bound does not necessarily exist for systems that are PAST. This holds even for non-probabilistic ARSs, as Example 6 demonstrates.

Proposition 2. For every PARS \rightarrow , the following implications hold:

$$\rightarrow \text{ is SAST} \implies \rightarrow \text{ is PAST} \implies \rightarrow \text{ is AST}$$

We conclude the section with the following natural property of expected derivation height. For a (multi)distribution μ over real numbers, the *expected value* of μ is defined by $\mathbb{E}(\mu) := \sum_{p:a \in \mu} p \cdot a$. As a function $f : A \rightarrow \mathbb{R}$ is naturally generalized to $f : \mathcal{M}_{\leq 1}(A) \rightarrow \mathcal{M}_{\leq 1}(\mathbb{R})$, for $\mu \in \mathcal{M}_{\leq 1}(A)$ we have $\mathbb{E}(f(\mu)) = \sum_{p:a \in \mu} p \cdot f(a)$.

³ We are grateful to the anonymous reviewer of an earlier version of this work [3] who pointed us to this example.

Lemma 4. For every PARS \rightarrow , $\text{edh}_{\rightarrow}(\mu) = \mathbb{E}(\text{edh}_{\rightarrow}(\mu))$.

Proof. More generally, we have

$$\begin{aligned}
\text{edh}_{\rightarrow}\left(\bigoplus_{i \in I} p_i \cdot \mu_i\right) &= \sup \left\{ \sum_{n \in \mathbb{N}} |v_n| \mid (v_n)_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}\left(\bigoplus_{i \in I} p_i \cdot \mu_i\right) \right\} \\
&\quad (\text{Lemma 3}) \\
&= \sup \left\{ \sum_{n \in \mathbb{N}} |\bigoplus_{i \in I} p_i \cdot v_{n,i}| \mid \forall i \in I. (v_{n,i})_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}(\mu_i) \right\} \\
&\quad (\text{Lemma 2}) \\
&= \sup \left\{ \sum_{i \in I} p_i \cdot \sum_{n \in \mathbb{N}} |v_{n,i}| \mid \forall i \in I. (v_{n,i})_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}(\mu_i) \right\} \\
&= \sum_{i \in I} p_i \cdot \sup \left\{ \sum_{n \in \mathbb{N}} |v_n| \mid (v_n)_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}(\mu_i) \right\} \\
&= \sum_{i \in I} p_i \cdot \text{edh}_{\rightarrow}(\mu_i).
\end{aligned}$$

The main claim then follows, since

$$\mu = \bigoplus_{p:a \in \mu} p \cdot \{\{1 : a\}\} \quad \text{and} \quad \sum_{p:a \in \mu} p \cdot \text{edh}_{\rightarrow}(\{\{1 : a\}\}) = \mathbb{E}(\text{edh}_{\rightarrow}(\mu)). \quad \square$$

3.3. Proving probabilistic termination

A popular way of proving termination of non-probabilistic systems is via *embedding* reductions into a well-founded set $(B, >)$ [29]. The latter can be seen also as a terminating ARS, giving rise to the following general statement.

Proposition 3. An ARS $\rightarrow \subseteq A \times A$ is terminating if and only if there exists a mapping $f : A \rightarrow B$ and a terminating ARS $> \subseteq B \times B$ such that $a \rightarrow b$ implies $f(a) > f(b)$.

The function f is often called a *ranking function*, typically B is \mathbb{N} , and $>$ is the standard order on \mathbb{N} . Notice that the “only if” direction is trivial by taking the identity as f and \rightarrow as $>$. We give an analog of Proposition 3 for probabilistic systems.

Definition 10 (*Embedding*). Let $\rightarrow \subseteq A \times \mathcal{M}(A)$ and $\sqsupset \subseteq B \times \mathcal{M}(B)$ be two PARSs. We say that a mapping $f : A \rightarrow B$ is an *embedding of \rightarrow into \sqsupset* , if $a \rightarrow \mu$ implies $f(a) \sqsupset f(\mu)$.

With Theorem 1 we will show that such embeddings give a sound and complete method for proving SAST. The proof of this theorem is based on the following simulation result.

Lemma 5. Let $f : A \rightarrow B$ be an embedding of a PARS \rightarrow into a PARS \sqsupset . If $\mu \rightarrow \nu$ then $f(\mu) \sqsupset f(\nu) \uplus \xi$ for some multidistribution ξ .

Proof. We proceed by induction on the derivation of $\mu \rightarrow \nu$.

- If $\mu = \{\{1 : a\}\} \rightarrow \emptyset = \nu$ with $a \in \text{NF}_{\rightarrow}$, then the claim is satisfied by any ξ with $f(\mu) \sqsupset \xi$. Note that by the definition of $\rightarrow^{\mathcal{M}}$, such a ξ always exists.
- If $\mu = \{\{1 : a\}\} \rightarrow \nu$ with $a \rightarrow \nu$, then the assumption gives $f(a) \sqsupset f(\nu)$ and thus $f(\mu) = \{\{1 : f(a)\}\} \sqsupset f(\nu)$.
- Finally, if $\mu = \bigoplus_{i \in I} p_i \cdot \mu_i \rightarrow \bigoplus_{i \in I} p_i \cdot \nu_i = \nu$ with $\mu_i \rightarrow \nu_i$, then the claim is an easy consequence of the induction hypothesis. \square

Lemma 6. Let $f : A \rightarrow B$ be an embedding of a PARS \rightarrow into a PARS \sqsupset . For every reduction sequence $(\mu_n)_{n \in \mathbb{N}} \in \mathbb{N} \in \text{red}_{\rightarrow}(\mu)$ there exists a reduction sequence $(\nu_n)_{n \in \mathbb{N}} \in \text{red}_{\sqsupset}(f(\mu))$ such that $f(\mu_n) \sqsupset \nu_n$ for all $n \in \mathbb{N}$. In particular, $\text{edl}((\mu_n)_{n \in \mathbb{N}}) \leq \text{edl}((\nu_n)_{n \in \mathbb{N}})$.

Proof. Given $(\mu_n)_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}(\mu)$, we inductively construct the desired $\nu_0, \nu_1, \nu_2, \dots$ as follows. We set $\nu_0 = f(\mu_0)$. Suppose $f(\mu_n) \sqsubseteq \nu_n$, and thus $\nu_n = f(\mu_n) \uplus \xi$ for some multidistribution ξ . Lemma 5 yields $f(\mu_n) \sqsupseteq f(\mu_{n+1}) \uplus \rho$ for some ρ . Take an arbitrary ξ' with $\xi \sqsupseteq \xi'$ and define $\nu_{n+1} = f(\mu_{n+1}) \uplus \rho \uplus \xi'$. We have

$$\nu_n = f(\mu_n) \uplus \xi \sqsupseteq f(\mu_{n+1}) \uplus \rho \uplus \xi' = \nu_{n+1}$$

as desired. From this, $\text{edl}((\mu_n)_{n \in \mathbb{N}}) \leq \text{edl}((\nu_n)_{n \in \mathbb{N}})$ follows as we have $|\mu_n| \leq |\nu_n|$ for every $n \in \mathbb{N}$. \square

Theorem 1. A PARS $\rightarrow \subseteq A \times \mathcal{M}(A)$ is SAST if and only if there exists an embedding $f : A \rightarrow B$ into a PARS $\sqsupseteq \subseteq B \times \mathcal{M}(B)$ which is SAST. Moreover, in this case $\text{edh}_{\rightarrow}(a) \leq \text{edh}_{\sqsupseteq}(f(a))$ for any $a \in A$.

Proof. The “only if” direction is trivial by taking \rightarrow as \sqsupseteq and the identity as f . Concerning the “if” direction, suppose that f is an embedding of \rightarrow into \sqsupseteq and \sqsupseteq is SAST. Consider an arbitrary $a \in A$. By Lemma 6, for every $\vec{\mu} \in \text{red}_{\rightarrow}(a)$ there is a “longer” reduction $\vec{\nu} \in \text{red}_{\sqsupseteq}(f(a))$ in the sense $\text{edl}(\vec{\mu}) \leq \text{edl}(\vec{\nu})$. Consequently,

$$\begin{aligned} \text{edh}_{\rightarrow}(a) &= \sup\{\text{edl}(\vec{\mu}) \mid \vec{\mu} \in \text{red}_{\rightarrow}(a)\} \\ &\leq \sup\{\text{edl}(\vec{\nu}) \mid \vec{\nu} \in \text{red}_{\sqsupseteq}(f(a))\} \\ &= \text{edh}_{\sqsupseteq}(f(a)) < \infty. \quad \square \end{aligned}$$

Probabilistic ranking functions [6] or ranking supermartingales [9,14] are instances of Theorem 1. To demonstrate this, we introduce the following PARS over $\mathbb{R}_{\geq 0}$, which is canonically SAST.

Definition 11 (Probabilistic Ranking Function). Let $\epsilon > 0$. We define the ARS $[\geq \epsilon +] \subseteq \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$ by

$$a [\geq \epsilon +] b :\iff a \geq \epsilon + b,$$

and the PARS $[\geq \epsilon + \mathbb{E}] \subseteq \mathbb{R}_{\geq 0} \times \mathcal{M}(\mathbb{R}_{\geq 0})$ by

$$a [\geq \epsilon + \mathbb{E}] \mu :\iff a \geq \epsilon + \mathbb{E}(\mu).$$

We call an embedding $f : A \rightarrow \mathbb{R}_{\geq 0}$ of a PARS $\rightarrow \subseteq A \times \mathcal{M}(A)$ into $[\geq \epsilon + \mathbb{E}]$ a *probabilistic ranking function* for \rightarrow .

Intuitively, a mapping f is a probabilistic ranking function for \rightarrow if the value of f decreases by ϵ in expectation whenever there is a reduction in \rightarrow . Lemma 9 will confirm that $[\geq \epsilon + \mathbb{E}]$ is SAST, and hence, probabilistic ranking functions are sound for proving SAST, by Theorem 1. Moreover, we will also see that they are complete. Towards the soundness result, we analyze reductions of $[\geq \epsilon + \mathbb{E}]$, in particular, how they evolve in expectation. We start with the analysis of a single reduction.

Lemma 7. If $\mu [\geq \epsilon + \mathbb{E}]^{\mathcal{M}} \nu$, then $\mathbb{E}(\mu) \geq \epsilon \cdot |\nu| + \mathbb{E}(\nu)$.

Proof. We prove the claim by induction on the derivation of $\mu [\geq \epsilon + \mathbb{E}]^{\mathcal{M}} \nu$.

- Suppose $\mu = \{\{1 : a\}\}$ and $a \in \text{NF}_{[\geq \epsilon + \mathbb{E}]}$, that is, $a < \epsilon$. Then $\nu = \emptyset$ and $\mathbb{E}(\mu) \geq 0 = \epsilon \cdot |\nu| + \mathbb{E}(\nu)$ since $\mathbb{E}(\emptyset) = |\emptyset| = 0$.
- Suppose $\mu = \{\{1 : a\}\}$ and $a [\geq \epsilon + \mathbb{E}] \nu$, that is, $a \geq \epsilon + \mathbb{E}(\nu)$. As $|\nu| = 1$ we have $\mathbb{E}(\mu) = a \geq \epsilon \cdot |\nu| + \mathbb{E}(\nu)$.
- Suppose $\mu = \biguplus_{i \in I} p_i \cdot \mu_i$, $\nu = \biguplus_{i \in I} p_i \cdot \nu_i$, and $\mu_i [\geq \epsilon + \mathbb{E}]^{\mathcal{M}} \nu_i$ for every $i \in I$. The induction hypothesis gives $\mathbb{E}(\mu_i) \geq \epsilon \cdot |\nu_i| + \mathbb{E}(\nu_i)$. Thus,

$$\mathbb{E}(\mu) = \sum_{i \in I} p_i \cdot \mathbb{E}(\mu_i) \geq \sum_{i \in I} p_i \cdot (\epsilon \cdot |\nu_i| + \mathbb{E}(\nu_i)) = \epsilon \cdot \sum_{i \in I} p_i \cdot |\nu_i| + \sum_{i \in I} p_i \cdot \mathbb{E}(\nu_i) = \epsilon \cdot |\nu| + \mathbb{E}(\nu). \quad \square$$

Lemma 8. For every $\vec{\mu} = (\mu_n)_{n \in \mathbb{N}} \in \text{red}_{[\geq \epsilon + \mathbb{E}]}(\mu_0)$, $\mathbb{E}(\mu_0) \geq \epsilon \cdot \text{edl}(\vec{\mu})$.

Proof. We first show $\mathbb{E}(\mu_m) \geq \epsilon \cdot \sum_{i=m+1}^n |\mu_i|$ for every $n \geq m$, by induction on $n - m$. The base case is trivial, so let us consider the inductive step. By Lemma 7 and the induction hypothesis we get

$$\begin{aligned} \mathbb{E}(\mu_m) &\geq \epsilon \cdot |\mu_{m+1}| + \mathbb{E}(\mu_{m+1}) \\ &\geq \epsilon \cdot |\mu_{m+1}| + \epsilon \cdot \sum_{i=m+2}^n |\mu_i| = \epsilon \cdot \sum_{i=m+1}^n |\mu_i|. \end{aligned}$$

By fixing $m = 0$, we conclude that the sequence $(\epsilon \cdot \sum_{i=1}^n |\mu_i|)_{n \geq 1}$ is bounded by $\mathbb{E}(\mu_0)$, and so is its limit $\epsilon \cdot \sum_{i \geq 1} |\mu_i| = \epsilon \cdot \text{edl}(\vec{\mu})$. \square

Lemma 9. *If $\epsilon > 0$ then $\text{edh}_{[\geq \epsilon + \mathbb{E}]}(a) \leq \frac{a}{\epsilon}$ for every $a \in \mathbb{R}_{\geq 0}$, and thus the PARS $[\geq \epsilon + \mathbb{E}]$ is SAST.*

Proof. By Lemma 8, we have $\text{edl}(\vec{\mu}) \leq \frac{\mathbb{E}(\langle 1, a \rangle)}{\epsilon} = \frac{a}{\epsilon}$ for every $\vec{\mu} \in \text{red}_{[\geq \epsilon + \mathbb{E}]}(a)$. Hence, $\text{edh}_{[\geq \epsilon + \mathbb{E}]}(a) \leq \frac{a}{\epsilon}$, concluding that $[\geq \epsilon + \mathbb{E}]$ is SAST. \square

Theorem 2. *A PARS \rightarrow is SAST if and only if there is a probabilistic ranking function f for \rightarrow , i.e., f embeds \rightarrow into $[\geq \epsilon + \mathbb{E}]$ for some $\epsilon > 0$. Moreover, in this case $\text{edh}_{\rightarrow}(a) \leq \frac{f(a)}{\epsilon}$ for any $a \in A$.*

Proof. The “if” direction follows from Theorem 1 and Lemma 9. Concerning the “only if” direction, suppose that \rightarrow is SAST. Thus, we have $\text{edh}_{\rightarrow} : A \rightarrow \mathbb{R}_{\geq 0}$. Consider $a \rightarrow \mu$. Then we have $|\mu| = 1$ and

$$\begin{aligned} \text{edh}_{\rightarrow}(a) &= \sup \{ \text{edl}(\vec{\mu}) \mid \vec{\mu} \in \text{red}_{\rightarrow}(a) \} \\ &\geq \sup \{ |\mu| + \text{edl}(\vec{\mu}) \mid \vec{\mu} \in \text{red}_{\rightarrow}(\mu) \} \\ &= |\mu| + \sup \{ \text{edl}(\vec{\mu}) \mid \vec{\mu} \in \text{red}_{\rightarrow}(\mu) \} \\ &= |\mu| + \text{edh}_{\rightarrow}(\mu) && \text{(Lemma 4)} \\ &= |\mu| + \mathbb{E}(\text{edh}_{\rightarrow}(\mu)), \end{aligned}$$

concluding $\text{edh}_{\rightarrow}(a) \geq 1 + \mathbb{E} \text{edh}_{\rightarrow}(\mu)$. Thus edh_{\rightarrow} is an embedding of \rightarrow into $[\geq 1 + \mathbb{E}]$, i.e., a ranking function. \square

3.4. Relation to formulation by Bournez and Garnier

The dynamics of probabilistic systems are commonly defined as stochastic processes, so that the n th random variable represents the n th reduct. Bournez and Garnier [6] follow this approach. In this section, we establish a precise correspondence between their formulation and ours. In particular, we show that the corresponding notions of AST and PAST coincide.

We assume familiarity with *stochastic processes*, see e.g. [34]. We briefly fix central notions and notations. A *measurable space* is a tuple (Ω, Σ) consisting of a set Ω and a *sigma-algebra* Σ on it, i.e., Σ is a collection of subsets of Ω that contains the empty set and is closed under complement and countable unions. A *probability space* is a triple $(\Omega, \Sigma, \mathbb{P})$ with (Ω, Σ) a measurable space and $\mathbb{P} : \Sigma \rightarrow \mathbb{R}_{\geq 0}$ a countable additive function such that $\mathbb{P}(\Omega) = 1$. Given $A, B \in \Sigma$ with $\mathbb{P}(B) > 0$, the *conditional probability* $\mathbb{P}(A|B)$ is defined as $\mathbb{P}(A \cap B)/\mathbb{P}(B)$. The law of total probability states that $\mathbb{P}(A) = \sum_i \mathbb{P}(A \cap B_i)$ for any $A \in \Sigma$ and finite or countable partition $(B_i)_{i \in \mathbb{N}} \in \Sigma$ of Ω .

A *random variable* X over a countable set A in a probability space $(\Omega, \Sigma, \mathbb{P})$ is a *measurable function* $X : \Omega \rightarrow A$, that is, $X^{-1}(a) := \{\omega \mid X(\omega) = a\} \in \Sigma$ for all $a \in A$. We follow the usual conventions concerning random variables, in particular, with $\mathbb{P}(X = a)$ we denote the probability $\mathbb{P}(X^{-1}(a))$. The *probability distribution* of X is the probability distribution over A that assigns to every $a \in A$ the probability $\mathbb{P}(X = a)$. The *expected value* of a random variable X over $\mathbb{N} \cup \{\infty\}$ is defined by $\mathbb{E}(X) := \sum_{n \in \mathbb{N} \cup \{\infty\}} \mathbb{P}(X = n) \cdot n$, where $0 \cdot \infty = 0$ and $p \cdot \infty = \infty$ for $p > 0$.

A *stochastic process* on A is an infinite sequence $\vec{X} = (X_n)_{n \in \mathbb{N}}$ of random variables over A all defined on some probability space $(\Omega, \Sigma, \mathbb{P})$. A *stopping time* with respect to \vec{X} is a random variable S , taking values in $\mathbb{N} \cup \{\infty\}$, with the property that for each $n \in \mathbb{N} \cup \{\infty\}$, the occurrence or non-occurrence of the event $S = n$ depends only on the values of X_0, \dots, X_n . Every stopping time S satisfies

$$\mathbb{E}(S) := \sum_{n=1}^{\infty} n \cdot \mathbb{P}(S = n) = \sum_{n=1}^{\infty} \mathbb{P}(S \geq n). \quad (1)$$

An instance of a stopping time is the *first hitting time* with respect to a set $H \subseteq A$, which is defined as $\tau_H(\omega) := \min\{n \mid X_n(\omega) \in H\}$ for all $\omega \in \Omega$, where $\min \emptyset = \infty$.

For the remaining of the section, we fix a PARS \rightarrow on A such that all right-hand sides in \rightarrow are distributions. In order to define reduction sequences of \rightarrow as stochastic processes, first, all nondeterministic choices are resolved by fixing a *strategy*, also called *policy*.

Definition 12 ([6]). An n -step *history* is a nonempty finite sequence $a_{0..n} = (a_0, \dots, a_n) \in A^+$. We say $a_{0..n}$ is (non)terminal iff a_n is. A *strategy* is a mapping $\phi : A^+ \rightarrow \mathcal{D}(A)$ that satisfies $a_n \rightarrow \phi(a_{0..n})$ whenever a_n is nonterminal. We say $a_{0..n}$ is *realizable under ϕ* iff $\phi(a_{0..i})(a_{i+1}) > 0$ for every $0 \leq i < n$.

Given an initial element or distribution, a strategy completely defines a stochastic process that corresponds to a particular reduction in the PARS. As [6] does not allow subdistributions, in the following we fix a special symbol $\perp \notin A$ to denote termination.

Definition 13 (Stochastic Reduction, [6]). A stochastic reduction under a strategy ϕ is a sequence $\vec{X} = (X_n)_{n \in \mathbb{N}}$ of random variables over $A \uplus \{\perp\}$ such that

$$\begin{aligned} \mathbb{P}(X_{n+1} = \perp \mid X_n = \perp) &= 1; \\ \mathbb{P}(X_{n+1} = \perp \mid X_n = a) &= 1 && \text{if } a \in \text{NF}_{\rightarrow}; \\ \mathbb{P}(X_{n+1} = \perp \mid X_n = a) &= 0 && \text{if } a \notin \text{NF}_{\rightarrow}; \\ \mathbb{P}(X_{n+1} = a \mid X_n = a_n, \dots, X_0 = a_0) &= \phi(a_{0..n})(a), \end{aligned}$$

whenever $a_{0..n}$ is a nonterminal realizable history under ϕ . The *initial distribution* of \vec{X} is given by the probability distribution of X_0 .

As an immediate consequence of the law of total probability, by Definition 13 any stochastic reduction $\vec{X} = (X_n)_{n \in \mathbb{N}}$ satisfies

$$\mathbb{P}(X_n = a_n) = \sum_{a_{0..n} \in A^{n+1}} \mathbb{P}(X_0 = a_0, \dots, X_n = a_n). \quad (2)$$

Notice that a_i ranges over A in this equality, i.e., $a_i \neq \perp$. As \perp signals termination, the derivation length of a stochastic reduction is given by the *first hitting time* to \perp .

Definition 14 (AST and PAST of [6]). For $\vec{X} = (X_n)_{n \in \mathbb{N}}$ define the random variable $T := \min\{n \in \mathbb{N} \mid X_n = \perp\}$ where $\min \emptyset = \infty$ by convention. The random variable T is called the *stopping time* of \vec{X} . A PARS \rightarrow is *stochastically AST* if for every stochastic reduction \vec{X} in \rightarrow , $\mathbb{P}(T = \infty) = 0$. We say \rightarrow is *stochastically PAST* if for every stochastic reduction \vec{X} in \rightarrow starting from $a \in A$, $\mathbb{E}(T) < \infty$.

We will now see that stochastic (P)AST coincides with (P)AST. To this end, we first clarify the correspondence of stochastic reductions and multidistribution reduction sequences. Given a stochastic reduction $\vec{X} = (X_n)_{n \in \mathbb{N}}$, for each $n \in \mathbb{N}$ we define the random variable $X_{0..n}$ ranging over $(A \cup \{\perp\})^{n+1}$ by

$$\mathbb{P}(X_{0..n} = a_{0..n}) = \mathbb{P}(X_0 = a_0, \dots, X_n = a_n),$$

where a_0, \dots, a_n range over $A \cup \{\perp\}$. Now consider projecting to the last element in $X_{0..n}$. We define the following multidistribution over A :

$$X_n^{\mathcal{M}} := \{\{p : a_n \mid p = \mathbb{P}(X_{0..n} = a_{0..n}) > 0\}\}.$$

The following two lemmas state the one-to-one correspondence between multidistribution reduction sequences and stochastic reductions in each direction.

Lemma 10 (Stochastic Reductions to Reductions). Let $\vec{X} = (X_n)_{n \in \mathbb{N}}$ be the stochastic reduction in \rightarrow under strategy ϕ . Then $X_n^{\mathcal{M}} \xrightarrow{\mathcal{M}} X_{n+1}^{\mathcal{M}}$ for every $n \in \mathbb{N}$, i.e., $(X_n^{\mathcal{M}})_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}(X_0)$.

Proof.

$$\begin{aligned} X_n^{\mathcal{M}} &= \{\{p : a_n \mid p = \mathbb{P}(X_{0..n} = a_{0..n}) > 0\}\} \\ &\rightarrow^{\mathcal{M}} \left[\bigoplus \right] \{p \cdot \phi(a_{0..n}) \mid p = \mathbb{P}(X_{0..n} = a_{0..n}) > 0, a_n \notin \text{NF}_{\rightarrow}\} \\ &= \{\{q : a_{n+1} \mid q = \mathbb{P}(X_{0..n} = a_{0..n}) \cdot \phi(a_{0..n})(a_{n+1}) > 0, a_n \notin \text{NF}_{\rightarrow}\}\} \\ &= \{\{q : a_{n+1} \mid q = \mathbb{P}(X_{0..n} = a_{0..n}, X_{n+1} = a_{n+1}) > 0\}\} \\ &= \{\{q : a_{n+1} \mid q = \mathbb{P}(X_{0..n+1} = a_{0..n+1}) > 0\}\} \\ &= X_{n+1}^{\mathcal{M}}. \quad \square \end{aligned}$$

Lemma 11 (Reduction to Stochastic Reductions). Let $\vec{\mu} = (\mu_n)_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}(d)$ for some distribution d . There exists a strategy ϕ which induces a stochastic reduction $(X_n)_{n \in \mathbb{N}}$ such that $\mu_n = X_n^{\mathcal{M}}$ for all $n \in \mathbb{N}$.

Proof. First, we inductively define $\phi_n : A^n \rightarrow A \rightarrow \mathbb{R}_{\geq 0}$ such that

$$\mu_n = \{\{\phi_n(a_{0..n-1})(a_n) : a_n \mid \phi_n(a_{0..n-1})(a_n) > 0\}\}. \quad (3)$$

Then by defining $\phi(a_{0..n}) = \phi_{n+1}(a_{0..n})$ we obtain the desired strategy. Since $d = \mu_0$ is a distribution, we define $\phi_0(a) := d(a)$. To obtain ϕ_{n+1} from ϕ_n , suppose by induction hypothesis that (3) holds. Let

$$I := \{a_{0..n} \mid \phi_n(a_{0..n-1})(a_n) > 0, a_n \notin \text{NF}_{\rightarrow}\}.$$

As $\mu_n \xrightarrow{\mathcal{M}} \mu_{n+1}$, for each $a_{0..n} \in I$ there is $a_n \rightarrow d_{a_{0..n}}$ such that

$$\begin{aligned} \mu_{n+1} &= \bigoplus_{a_{0..n} \in I} \phi_n(a_{0..n-1})(a_n) \cdot d_{a_{0..n}} \\ &= \{\{\phi_n(a_{0..n-1})(a_n) \cdot d_{a_{0..n}}(a) : a \mid \phi_{n+1}(a_{0..n})(a) > 0\}\}. \end{aligned}$$

Hence, defining

$$\phi_{n+1}(a_{0..n}) := \begin{cases} \phi_n(a_{0..n-1})(a_n) \cdot d_{a_{0..n}} & \text{if } (a_{0..n}) \in I, \\ 0 & \text{otherwise,} \end{cases}$$

yields the desired ϕ_{n+1} . \square

Now we relate the expected derivation length of multidistribution reduction sequences and the expected stopping time of stochastic reductions. For a multidistribution $\mu \in \mathcal{M}_{\leq 1}(A)$, we define the distribution $\bar{\mu} \in \mathcal{D}(A \cup \{\perp\})$ by $\bar{\mu}(a) := \sum_{p:a \in \mu} p$ for $a \in A$ and $\bar{\mu}(\perp) := 1 - |\mu|$. We have the following natural correspondence between $X_n^{\mathcal{M}}$ and X_n .

Lemma 12. $\overline{X_n^{\mathcal{M}}}(a) = \mathbb{P}(X_n = a)$.

Proof. The lemma follows from the definition of $\overline{(\cdot)}$ and Equation (2). \square

Lemma 13. Let $\vec{X} = (X_n)_{n \in \mathbb{N}}$ be the stochastic derivation in \rightarrow under strategy ϕ from the initial distribution d and let T denote its stopping time. The following two properties hold:

1. $\mathbb{P}(T \geq n) = |X_n^{\mathcal{M}}|$ for every $n \in \mathbb{N}$.
2. $\mathbb{P}(T = \infty) = \lim_{n \rightarrow \infty} |X_n^{\mathcal{M}}|$.

Proof. Concerning the first property, we have

$$\mathbb{P}(T \geq n) = \mathbb{P}(X_n \in A) = \sum_{a \in A} \mathbb{P}(X_n = a) = \sum_{a \in A} \overline{X_n^{\mathcal{M}}}(a) = |X_n^{\mathcal{M}}|,$$

for all $n \in \mathbb{N}$, where the penultimate equation follows from Lemma 12. As we have

$$\mathbb{P}(T = \infty) = \lim_{n \rightarrow \infty} \mathbb{P}(T \geq n),$$

the second property follows from the first. \square

Now we arrive at the main theorem of this section.

Theorem 3. A PARS \rightarrow is (P)AST if and only if it is stochastically (P)AST.

Proof. We consider the “if” direction first. Suppose \rightarrow is AST. Lemma 10 translates an arbitrary stochastic derivation $\vec{X} = (X_n)_{n \in \mathbb{N}}$ in \rightarrow to a reduction $(X_n^{\mathcal{M}})_{n \in \mathbb{N}} \in \text{red}_{\rightarrow}(a)$, for which we have $\mathbb{P}(T = \infty) = \lim_{n \rightarrow \infty} |\mu_n| = 0$ by Lemma 13(2). Hence, \rightarrow is stochastically AST. If \rightarrow is moreover PAST, using Lemma 13(1) and (1) we get

$$\infty > \text{edl} \left((X_n^{\mathcal{M}})_{n \in \mathbb{N}} \right) = \sum_{n \geq 1} |X_n^{\mathcal{M}}| = \sum_{n \geq 1} \mathbb{P}(T \geq n) = \mathbb{E}(T),$$

where we tacitly employ $\mathbb{P}(T = \infty) = 0$. The “only if” direction is proven dually, using Lemma 11. \square

Finally, we relate our probabilistic ranking functions (Theorem 2) to the following formulation by Bournez and Garnier [6].

Proposition 4 ([6]). A PARS \rightarrow over A is PAST if there exists a function $V : A \rightarrow \mathbb{R}$, with $\inf_{a \in A} V(a) > -\infty$, and $\epsilon > 0$ such that for all $a \in A$, for all μ with $a \rightarrow \mu$, the drift in a according to μ defined by

$$\Delta_{\mu} V(a) := \sum_{p:b \in \mu} p \cdot V(b) - V(a),$$

satisfies

$$\Delta_{\mu} V(a) \leq -\epsilon. \quad (4)$$

Formulated differently, condition (4) states that $V(a) \geq \epsilon + \mathbb{E}(V(\mu))$ holds whenever $a \rightarrow \mu$. Further, $\inf_{a \in A} V(a) > -\infty$ means that $f(a) := V(a) - \inf_{a \in A} V(a)$ gives a mapping $f : A \rightarrow \mathbb{R}_{\geq 0}$. Hence, f is a probabilistic ranking function in the sense of Theorem 2.

Proposition 5. Let \rightarrow be a PARS. There exists $V : A \rightarrow \mathbb{R}$ satisfying Proposition 4 if and only if there is a probabilistic ranking function for \rightarrow .

As a consequence, Proposition 4 can be strengthened from PAST to SAST. Moreover, the method is not only sound but also *complete* for SAST. Bournez and Garnier claimed with [6, Theorem 3] that Proposition 4 is complete for proving finitely branching PARSs to be PAST. Example 7 is a counterexample to this claim, depicting a finitely branching system that is PAST but not SAST.

4. Probabilistic term rewriting

We now formulate probabilistic term rewriting, and then lift the interpretation method for term rewriting to the probabilistic case.

We briefly recap notions from term rewriting; see [5] for more details. A *signature* F is a set of *function symbols*, each associated with a natural number called *arity*. The set $T(F, V)$ of *terms* over a signature F and a set V of variables (disjoint with F) is the least set such that $x \in T(F, V)$ if $x \in V$ and $\mathfrak{f}(t_1, \dots, t_n) \in T(F, V)$ whenever $\mathfrak{f} \in F$ is of arity n and $t_1, \dots, t_n \in T(F, V)$. A *substitution* is a mapping $\sigma : V \rightarrow T(F, V)$, which is extended homomorphically to terms. We write $t\sigma$ instead of $\sigma(t)$. A *context* is a term $C \in T(F, V \cup \{\square\})$ containing exactly one occurrence of a special variable \square . With $C[t]$ we denote the term obtained by replacing \square in C with t .

We extend substitutions and contexts to multidistributions over terms as follows: $\mu\sigma := \{\{p_i : t_i\sigma \mid i \in I\}\}$ and $C[\mu] := \{\{p_i : C[t_i] \mid i \in I\}\}$ for $\mu = \{\{p_i : t_i \mid i \in I\}\}$.

Definition 15 (*Probabilistic Term Rewriting*). A *probabilistic rewrite rule* is a pair of $l \in T(F, V)$ and $\mu \in \mathcal{M}(T(F, V))$, written $l \rightarrow \mu$. A *probabilistic term rewrite system* (PTRS) \mathcal{R} is a (typically finite) set of probabilistic rewrite rules. We write $\xrightarrow{\mathcal{R}}$ for the least PARS such that $C[l\sigma] \xrightarrow{\mathcal{R}} C[\mu\sigma]$ for every probabilistic rewrite rule $l \rightarrow \mu \in \mathcal{R}$, context C , and substitution σ . We say a PTRS \mathcal{R} is AST/PAST/SAST if $\xrightarrow{\mathcal{R}}$ is.

Note here that we deviate from [6]: there $\xrightarrow{\mathcal{R}}$ is defined by $C[l\sigma] \xrightarrow{\mathcal{R}} C[\overline{\mu\sigma}]$ for each $l \rightarrow \mu \in \mathcal{R}$ in our notation, where $\overline{\mu\sigma}$ is the distribution corresponding to the multidistribution $\mu\sigma$. Notice that even if μ is a distribution over terms, $\mu\sigma$ is in general a (proper) multidistribution; e.g., consider $\{\frac{1}{2} : x; \frac{1}{2} : y\}\sigma$ with $x\sigma = y\sigma$. On the other hand, Bournez and Garnier [6] demand that right-hand sides of PARSs to be distributions. Consequently, the extra operation $\bar{\cdot}$ is necessary, in order for $\xrightarrow{\mathcal{R}}$ to be a PARS of their definition. One consequence of this difference is that we allow strategies to know which rule was applied even if the resulting term was identical, while in the setting of Bournez and Garnier [6] strategies can be defined only in terms of the resulting term.

Example 8. The random walk of Example 1 can be modeled by a PTRS consisting of a single rule $s(x) \rightarrow \{\{p : x; 1 - p : s(s(x))\}\}$.

To rewrite a term, there are typically multiple choices of a subterm to reduce (i.e., *redexes*). For instance, $s(\mathfrak{f}(s(0)))$ has two redexes in the above PTRS, and consequently two possible reducts:

$$\{\{p : \mathfrak{f}(s(0)); 1 - p : s(\mathfrak{f}(s(0)))\}\} \quad \text{and} \quad \{\{p : s(\mathfrak{f}(0)); 1 - p : s(\mathfrak{f}(s(s(0))))\}\}.$$

4.1. Interpretation methods for proving SAST

We now generalize the *interpretation method* for term rewrite systems to the probabilistic setting. The following notion is standard.

Definition 16 (*F-Algebra, cf. [38]*). An *F-algebra* \mathcal{X} on a non-empty carrier set X specifies the *interpretation* $\varepsilon_{\mathcal{X}} : X^n \rightarrow X$ of each function symbol $\varepsilon \in F$ of arity n . We say \mathcal{X} is *monotone* with respect to a binary relation (an ARS) $\succ \subseteq X \times X$ if $x \succ y$ implies $\varepsilon_{\mathcal{X}}(\dots, x, \dots) \succ \varepsilon_{\mathcal{X}}(\dots, y, \dots)$ for every $\varepsilon \in F$. Given an *assignment* $\alpha : V \rightarrow X$, the interpretation of a term is defined as follows:

$$\llbracket t \rrbracket_{\mathcal{X}}^{\alpha} := \begin{cases} \alpha(t) & \text{if } t \in V, \\ \varepsilon_{\mathcal{X}}(\llbracket t_1 \rrbracket_{\mathcal{X}}^{\alpha}, \dots, \llbracket t_n \rrbracket_{\mathcal{X}}^{\alpha}) & \text{if } t = \varepsilon(t_1, \dots, t_n). \end{cases}$$

We write $s \succ_{\mathcal{X}} t$ iff $\llbracket s \rrbracket_{\mathcal{X}}^{\alpha} \succ \llbracket t \rrbracket_{\mathcal{X}}^{\alpha}$ for every assignment α .

Proposition 6 (*cf. [38]*). A TRS \mathcal{R} is *terminating* if and only if there exists an *F-algebra* \mathcal{X} which is *monotone* with respect to a *terminating relation* \succ and satisfies $\mathcal{R} \subseteq \succ_{\mathcal{X}}$.

In essence, this proposition is a consequence of Proposition 3: the imposed conditions witness that the interpretation embeds the ARS $\xrightarrow{\mathcal{R}}$ underlying \mathcal{R} into the well-founded order \succ . Conversely, completeness can be proven by taking $\xrightarrow{\mathcal{R}}$ as the well-founded order and for \mathcal{X} the *term algebra* \mathcal{T} , an *F-algebra* on terms such that $\varepsilon_{\mathcal{T}}(t_1, \dots, t_n) := \varepsilon(t_1, \dots, t_n)$. Note that in the term algebra, assignments are substitutions, and $\llbracket t \rrbracket_{\mathcal{T}}^{\sigma} = t\sigma$.

We now generalize Proposition 6 for probabilistic TRSs via Theorem 1. For an *F-algebra* \mathcal{X} , we lift the interpretation of terms to multidistributions as follows:

$$\llbracket \{p_i : t_i \mid i \in I\} \rrbracket_{\mathcal{X}}^{\alpha} := \{p_i : \llbracket t_i \rrbracket_{\mathcal{X}}^{\alpha} \mid i \in I\}.$$

Definition 17 (*Probabilistic Monotone F-Algebra*). A *probabilistic monotone F-algebra* (\mathcal{X}, \sqsupset) is an *F-algebra* \mathcal{X} equipped with a relation $\sqsupset \subseteq X \times \mathcal{M}(X)$, such that $x \sqsupset \mu$ implies $\varepsilon_{\mathcal{X}}(\dots, x, \dots) \sqsupset \varepsilon_{\mathcal{X}}(\dots, \mu, \dots)$ for every $\varepsilon \in F$, where

$$\varepsilon_{\mathcal{X}}(\dots, \mu, \dots) := \{p : \varepsilon_{\mathcal{X}}(\dots, t, \dots) \mid p : t \in \mu\}.$$

We write $t \sqsupset_{\mathcal{X}} \mu$ iff $\llbracket t \rrbracket_{\mathcal{X}}^{\alpha} \sqsupset \llbracket \mu \rrbracket_{\mathcal{X}}^{\alpha}$ for every assignment $\alpha : V \rightarrow X$.

Notice that \sqsupset is a PARS over \mathcal{X} . As in the non-probabilistic case, it is closed under substitutions and contexts:

Lemma 14. Let (\mathcal{X}, \sqsupset) be a *probabilistic monotone F-algebra*. If $s \sqsupset_{\mathcal{X}} \mu$ then $\llbracket s\sigma \rrbracket_{\mathcal{X}}^{\alpha} \sqsupset \llbracket \mu\sigma \rrbracket_{\mathcal{X}}^{\alpha}$ and $\llbracket C[s] \rrbracket_{\mathcal{X}}^{\alpha} \sqsupset \llbracket C[\mu] \rrbracket_{\mathcal{X}}^{\alpha}$ for arbitrary α, σ , and C .

Proof. Let $\mu = \{p_i : t_i \mid i \in I\}$. Concerning the first property, define the assignment β by $\beta(x) = \llbracket x\sigma \rrbracket_{\mathcal{X}}^{\alpha}$ for every $x \in V$. By structural induction on t , one can verify $\llbracket t \rrbracket_{\mathcal{X}}^{\beta} = \llbracket t\sigma \rrbracket_{\mathcal{X}}^{\alpha}$ for any term t . Thus, from the assumption we get

$$\begin{aligned} \llbracket s\sigma \rrbracket_{\mathcal{X}}^{\alpha} &= \llbracket s \rrbracket_{\mathcal{X}}^{\beta} \sqsupset \llbracket \mu \rrbracket_{\mathcal{X}}^{\beta} = \{p_i : \llbracket t_i \rrbracket_{\mathcal{X}}^{\beta} \mid i \in I\} \\ &= \{p_i : \llbracket t_i\sigma \rrbracket_{\mathcal{X}}^{\alpha} \mid i \in I\} = \llbracket \mu\sigma \rrbracket_{\mathcal{X}}^{\alpha}. \end{aligned}$$

The second property is proven by induction on C , where the base case follows directly from the assumption, and the inductive step from monotonicity. \square

Theorem 4. A PTRS \mathcal{R} is SAST if and only if there exists a *probabilistic monotone F-algebra* (\mathcal{X}, \sqsupset) such that \sqsupset is SAST and $\mathcal{R} \subseteq \sqsupset_{\mathcal{X}}$.

Proof. For the “if” direction, let $\alpha : V \rightarrow X$ be an arbitrary assignment, which exists as X is non-empty. We show that $\llbracket \cdot \rrbracket_{\mathcal{X}}^{\alpha}$ is an embedding of $\xrightarrow{\mathcal{R}}$ into \sqsupset , so that the claim follows from Theorem 1. Consider $s \xrightarrow{\mathcal{R}} \mu$. Then we have $s = C[l\sigma]$ and $\mu = C[l\sigma]$ for some σ, C , and $l \rightarrow v \in \mathcal{R}$. By assumption we have $l \sqsupset_{\mathcal{X}} v$, and thus $\llbracket s \rrbracket_{\mathcal{X}}^{\alpha} \sqsupset \llbracket \mu \rrbracket_{\mathcal{X}}^{\alpha}$ by Lemma 14.

For the “only if” direction, it suffices to show that $(\mathcal{T}, \xrightarrow{\mathcal{R}})$ forms a *probabilistic monotone F-algebra*, which easily follows from Lemma 14. \square

4.2. Barycentric algebras

As probabilistic F -algebras are defined so generally, it is not yet clear how to search for ones for proving that a given PARS is SAST. Now we make one step towards finding probabilistic algebras, by imposing some conditions on (non-probabilistic) F -algebras, so that the relation \sqsupseteq can be defined from orderings which we are more familiar with. The following generalizes the PARS $[\geq \epsilon + \mathbb{E}]$ from Definition 11.

Definition 18 (*Barycentric Domain*). A *barycentric domain* is a set X where the *barycentric operation* $\mathbb{E} : \mathcal{M}(X) \rightarrow X$ is defined. From a binary relation \succ on X we define the PARS $[\succ \mathbb{E}] \subseteq X \times \mathcal{M}(X)$ by

$$x [\succ \mathbb{E}] \mu :\iff x \succ \mathbb{E}(\mu).$$

Of particular interest in this work will be the barycentric domains $\mathbb{R}_{\geq 0}$ and $\mathbb{R}_{> 0}^m$ with barycentric operations $\mathbb{E}(\{p_i \cdot a_i \mid i \in I\}) = \sum_{i \in I} p_i \cdot a_i$. We may write \mathbb{E}_X when we would like to clarify the domain X . The following generalizes standard notions from mathematics.

Definition 19 (*Concavity, Affinity*). Let $f : X \rightarrow Y$ be a function from and to barycentric domains. We say f is *concave* with respect to an order \succ on Y if $f(\mathbb{E}_X(\mu)) \succcurlyeq \mathbb{E}_Y(f(\mu))$ where \succcurlyeq is the reflexive closure of \succ . We say f is *affine* if it satisfies $f(\mathbb{E}_X(\mu)) = \mathbb{E}_Y(f(\mu))$.

Clearly, every affine function is concave.

Definition 20 (*Barycentric F-Algebra*). A *barycentric F-algebra* is a pair (\mathcal{X}, \succ) of an F -algebra \mathcal{X} on a barycentric domain X and an order \succ on X , such that for every $f \in F$, $f_{\mathcal{X}}$ is monotone and concave with respect to \succ in every argument.

Note that the following theorem claims soundness but not completeness, in contrast to Theorem 4.

Theorem 5. A PTRS \mathcal{R} is SAST if there exists a barycentric F -algebra (\mathcal{X}, \succ) such that $\mathcal{R} \subseteq [\succ \mathbb{E}]_{\mathcal{X}}$ and $[\succ \mathbb{E}]$ is SAST.

Proof. Due to Theorem 4, it suffices to show that $(\mathcal{X}, [\succ \mathbb{E}])$ is a probabilistic monotone F -algebra. To this end, suppose that $x \succ \mathbb{E}_X(\mu)$, and let $f \in F$. Since $f_{\mathcal{X}}$ is monotone and concave with respect to \succ , we have

$$f_{\mathcal{X}}(\dots, x, \dots) \succ f_{\mathcal{X}}(\dots, \mathbb{E}_X(\mu), \dots) \succcurlyeq \mathbb{E}_X(f_{\mathcal{X}}(\dots, \mu, \dots)). \quad \square$$

Extending a notion from Hirokawa and Moser [21], we say a relation $\succ \subseteq X \times X$ on a barycentric domain is *collapsible* if there is a concave function $f : X \rightarrow \mathbb{R}_{\geq 0}$ that embeds \succ into $[\geq \epsilon + \cdot]$.

Lemma 15. If $\succ \subseteq X \times X$ is collapsible then $[\succ \mathbb{E}]$ is SAST.

Proof. Suppose $x [\succ \mathbb{E}] \mu$, i.e., $x \succ \mathbb{E}_X(\mu)$. This implies

$$f(x) \geq \epsilon + f(\mathbb{E}_X(\mu)) \geq \epsilon + \mathbb{E}(f(\mu)),$$

for some $f : X \rightarrow \mathbb{R}_{\geq 0}$. Thus $f(x) [\geq \epsilon + \mathbb{E}] f(\mu)$, i.e., f is a probabilistic ranking function, and hence the lemma follows from Theorem 2. \square

The following is thus an immediate consequence of Theorem 1.

Corollary 1. A PTRS \mathcal{R} is SAST if there exists a barycentric F -algebra (\mathcal{X}, \succ) such that $\mathcal{R} \subseteq [\succ \mathbb{E}]_{\mathcal{X}}$ for a collapsible order \succ .

In the rest of this section we recast two popular interpretation methods, polynomial and matrix interpretations (over the reals), as barycentric F -algebras.

4.3. Polynomial interpretations

Polynomial interpretations were introduced (on natural numbers [28] and real numbers [30]) for the termination analysis of non-probabilistic rewrite systems. Various techniques for synthesizing polynomial interpretations (e.g., [16]) exist, and these techniques are easily applicable in our setting.

Definition 21 (*Polynomial Interpretation*). A polynomial interpretation is an F -algebra \mathcal{X} on $\mathbb{R}_{\geq 0}$ such that $\mathfrak{f}_{\mathcal{X}}$ is a polynomial for every $\mathfrak{f} \in F$. We say \mathcal{X} is *multilinear* if every $\mathfrak{f}_{\mathcal{X}}$ is of the following form with $c_V \in \mathbb{R}_{\geq 0}$:

$$\mathfrak{f}_{\mathcal{X}}(x_1, \dots, x_n) = \sum_{V \subseteq \{x_1, \dots, x_n\}} c_V \cdot \prod_{x_i \in V} x_i.$$

In order to use polynomial interpretations for probabilistic termination, multilinearity is necessary for satisfying the concavity condition.

Proposition 7. A PTRS \mathcal{R} is SAST if there exist $\epsilon > 0$ and a monotone multilinear polynomial interpretation \mathcal{X} such that $\mathcal{R} \subseteq [\geq \epsilon + \mathbb{E}]_{\mathcal{X}}$.

Proof. We show that $(\mathcal{X}, [\geq \epsilon + \cdot])$ forms a collapsible barycentric F -algebra, and thus Theorem 5 shows that \mathcal{R} is SAST. Collapsibility is trivial with $\mathbb{G}(x) = x$ and monotonicity is by assumption. Further, every multilinear polynomial is affine and thus concave in all variables. \square

An observation by Lucas [30] also holds in probabilistic case: To prove a finite PTRS \mathcal{R} SAST with polynomial interpretations, we do not have to find ϵ , but it is sufficient to check $l \text{ } [\gt \mathbb{E}]_{\mathcal{X}} \mu$ for all rules $l \rightarrow \mu \in \mathcal{R}$. Define $\epsilon_{l \rightarrow \mu} := \llbracket l \rrbracket_{\mathcal{X}}^{\alpha} - \mathbb{E}(\llbracket \mu \rrbracket_{\mathcal{X}}^{\alpha})$ for such α that $\alpha(x) = 0$. Then for any other α , we can show $\llbracket l \rrbracket_{\mathcal{X}}^{\alpha} - \mathbb{E}(\llbracket \mu \rrbracket_{\mathcal{X}}^{\alpha}) \geq \epsilon_{l \rightarrow \mu} > 0$. As \mathcal{R} is finite, we take $\epsilon := \min\{\epsilon_{l \rightarrow \mu} \mid l \rightarrow \mu \in \mathcal{R}\} > 0$.

Corollary 2. A finite PTRS \mathcal{R} is SAST if there exists a monotone multilinear polynomial interpretation \mathcal{X} such that $\mathcal{R} \subseteq [\gt \mathbb{E}]_{\mathcal{X}}$.

Example 9 (*Example 8 Revisited*). Consider again the PTRS consisting of the single rule $s(x) \rightarrow \{p : x; 1 - p : s(s(x))\}$. Define the polynomial interpretation \mathcal{X} by $0_{\mathcal{X}} := 0$ and $s_{\mathcal{X}}(x) := x + 1$. Then whenever $p > \frac{1}{2}$ we have

$$\llbracket s(x) \rrbracket_{\mathcal{X}}^{\alpha} = x + 1 > p \cdot x + (1 - p) \cdot (x + 2) = \mathbb{E}(\llbracket \{p : x; 1 - p : s(s(x)) \} \rrbracket_{\mathcal{X}}^{\alpha}).$$

Thus, when $p > \frac{1}{2}$ the PTRS is SAST by Corollary 2.

Corollary 2 constitutes a generalization of [6, Theorem 5]. The latter does not cover linear interpretations, since *context decrease* [6, Definition 8] demands e.g. $\llbracket \mathfrak{f}(t) \rrbracket_{\mathcal{X}}^{\alpha} - \llbracket \mathfrak{f}(t') \rrbracket_{\mathcal{X}}^{\alpha} \leq \llbracket t \rrbracket_{\mathcal{X}}^{\alpha} - \llbracket t' \rrbracket_{\mathcal{X}}^{\alpha}$ and thus excludes interpretations such as $\mathfrak{f}_{\mathcal{X}}(x) = 2x$.

4.4. Matrix interpretations

Matrix interpretations are introduced for the termination analysis of term rewriting [13]. Now we extend them for probabilistic term rewriting.

Definition 22 (*Matrix Interpretation*). A (real) matrix interpretation is an F -algebra \mathcal{X} on $\mathbb{R}_{\geq 0}^m$ such that for every $f \in F$, $\mathfrak{f}_{\mathcal{X}}$ is of the form

$$\mathfrak{f}_{\mathcal{X}}(\vec{x}_1, \dots, \vec{x}_n) = \sum_{i=1}^n C_i \cdot \vec{x}_i + \vec{c}, \quad (5)$$

where $\vec{c} \in \mathbb{R}_{\geq 0}^m$, and $C_i \in \mathbb{R}_{\geq 0}^{m \times m}$. The order $\gg_{\epsilon} \subseteq \mathbb{R}_{\geq 0}^m \times \mathbb{R}_{\geq 0}^m$ is defined by

$$\begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix} \gg_{\epsilon} \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix} : \iff x_1 \geq \epsilon + y_1 \text{ and } x_i \geq y_i \text{ for all } i = 2, \dots, m.$$

Monotonicity of matrix interpretations can be ensured by requiring (5) to satisfy $(C_i)_{1,1} \geq 1$ for all i , cf. [13]. It is easy to derive the following from Theorem 5:

Proposition 8. A PTRS \mathcal{R} is SAST if there exist $\epsilon > 0$ and a monotone matrix interpretation \mathcal{X} such that $\mathcal{R} \subseteq [\gg_{\epsilon} \mathbb{E}]_{\mathcal{X}}$.

Proof. The order \gg_{ϵ} is collapsible with $f((x_1, \dots, x_m)^T) = x_1$. It is well known that (5) is affine and thus concave. \square

As for polynomial interpretations, for finite systems we do not have to find ϵ . Below \gg is defined in the same manner as \gg_{ϵ} but replacing “ $\geq \epsilon +$ ” by $>$.

Corollary 3. A finite PTRS \mathcal{R} is SAST if there exists a monotone matrix interpretation \mathcal{X} such that $\mathcal{R} \subseteq [\gg \mathbb{E}]_{\mathcal{X}}$.

Example 10. Consider the PTRS consisting of the single probabilistic rule

$$a(a(x)) \rightarrow \{p : a(a(a(x))); 1 - p : a(b(a(x)))\}.$$

Consider the two-dimensional matrix interpretation

$$a_{\mathcal{X}}(\vec{x}) = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} \cdot \vec{x} + \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad b_{\mathcal{X}}(\vec{x}) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \cdot \vec{x}.$$

Then we have

$$\begin{aligned} \llbracket a(a(x)) \rrbracket_{\mathcal{X}}^{\alpha} &= \begin{bmatrix} x_1 + x_2 + 1 \\ 1 \end{bmatrix} \gg_{1-2p} \begin{bmatrix} x_1 + x_2 + 2p \\ 1 \end{bmatrix} \\ &= p \cdot \llbracket a(a(a(x))) \rrbracket_{\mathcal{X}}^{\alpha} + (1 - p) \cdot \llbracket a(b(a(x))) \rrbracket_{\mathcal{X}}^{\alpha} \end{aligned}$$

where $\alpha(x) = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$. Hence this PARS is SAST if $p < \frac{1}{2}$, by Corollary 3.

Note that the above example cannot be handled with polynomial interpretations, intuitively because monotonicity enforces the interpretation of the probable reducts $a(a(a(x)))$ and $a(b(a(x)))$ to be greater than that of the left-hand side $a(a(x))$. Generally, polynomial and matrix interpretations are incomparable in strength, since multilinear polynomials are not expressible in the form of (5), although linear ones are.

5. Implementation

We extended the termination prover NaTT [39] with a syntax for probabilistic rules, and implemented the probabilistic versions of polynomial and matrix interpretations as NaTT version 1.9.

The input format extends the WST format.⁴ A probabilistic rewrite rule is specified by

$$l \rightarrow w_1 : r_1 \quad || \quad \dots \quad || \quad w_n : r_n$$

indicating the probabilistic rewrite rule

$$l \rightarrow \left\{ \left\{ \frac{w_1}{w} : r_1, \dots, \frac{w_n}{w} : r_n \right\} \right\} \quad \text{with } w = \sum_{j=1}^n w_j. \quad (6)$$

The problem of finding interpretations is encoded as a *satisfiability modulo theory* (SMT) problem and solved by an SMT solver. We already have an implementation to encode that $\llbracket l \rrbracket_{\mathcal{X}}^{\alpha} > \llbracket r \rrbracket_{\mathcal{X}}^{\alpha}$ holds for arbitrary α , so we only need a little extension to encode

$$w \cdot \llbracket l \rrbracket_{\mathcal{X}}^{\alpha} > w_1 \cdot \llbracket r_1 \rrbracket_{\mathcal{X}}^{\alpha} + \dots + w_n \cdot \llbracket r_n \rrbracket_{\mathcal{X}}^{\alpha}$$

which expresses the desired orientation condition of a probabilistic rule (6).

Example 11. The bound depicted in Example 10 is found by our implementation in NaTT .

The following example deserves some attention.

Example 12. Consider the following encoding of [14, Fig. 1]:

$$\begin{aligned} ?(x) &\rightarrow \left\{ \left\{ \frac{1}{2} : ?(\text{s}(x)); \frac{1}{2} : \$(\text{g}(x)) \right\} \right\} & \$(0) &\rightarrow \{ \{ 1 : 0 \} \} \\ ?(x) &\rightarrow \{ \{ 1 : \$(\text{f}(x)) \} \} & \$(\text{s}(x)) &\rightarrow \{ \{ 1 : \$(x) \} \} \end{aligned}$$

describing a game where the player (strategy) can choose either to quit the game and ensure prize $\$(\text{f}(x))$, or to try a coin-toss which on success increments the score and on failure ends the game with consolation prize $\$(\text{g}(x))$.

When f and g can be bounded by linear polynomials, it is possible to automatically prove that the system is SAST. For instance, with rules for $\text{f}(x) = 2x$ and $\text{g}(x) = \lfloor \frac{x}{2} \rfloor$, NaTT (combined with the SMT solver z3 version 4.4.1) found the following polynomial interpretation proving SAST:

⁴ <https://www.lri.fr/~marche/tpdb/format.html>, accessed November 14, 2017.

$$\begin{array}{lll} ?_{\mathcal{X}}(x) = 7x + 11 & s_{\mathcal{X}}(x) = x + 1 & 0_{\mathcal{X}} = 1 \\ f_{\mathcal{X}}(x) = 3x + 1 & g_{\mathcal{X}}(x) = 2x + 1 & \$_{\mathcal{X}}(x) = 2x + 1. \end{array}$$

While the above example is merely an encoding of a procedural probabilistic program, PTRSs are well-suited to model probabilistic *functional* programs featuring datatypes such as list, trees, etc.

Example 13. The following PTRS encodes a probabilistic function `rlist` that samples a list over naturals, where elements and the length of the list follow a geometric distribution.

$$\begin{array}{l} \text{rlist}(xs) \rightarrow \{\{\frac{1}{2} : xs, \frac{1}{2} : \text{cons}(\text{rnat}(0), \text{rlist}(xs))\}\} \\ \text{rnat}(x) \rightarrow \{\{\frac{1}{2} : x, \frac{1}{2} : \text{rnat}(s(x))\}\} \end{array}$$

The system is SAST, as NaTT finds the following polynomial interpretation:

$$\begin{array}{lll} 0_{\mathcal{X}}() = 1 & s_{\mathcal{X}}(x) = x & \text{cons}_{\mathcal{X}}(x, xs) = x + xs \\ \text{rnat}_{\mathcal{X}}(x) = x + 1 & \text{rlist}_{\mathcal{X}}(xs) = x + 3. & \end{array}$$

Finally, we also remark that our methods cannot handle the case where the expected derivation length is not bounded by multilinear polynomials.

Example 14. Consider Example 12 with rules $f(x) \rightarrow \text{mul}(x, x)$ and other rules that demand $\text{mul}(x, x)$ to take reduction length x^2 . For any (multi)linear polynomial $f_{\mathcal{X}}$, it is impossible to have $f_{\mathcal{X}}(x) > x^2$. Hence we cannot prove the PTRS to be SAST by the multilinear polynomial interpretation method. The same holds for the matrix interpretation.

6. Conclusion

This is a study on how much of the classic interpretation-based techniques well known in term rewriting can be extended to *probabilistic* term rewriting, and to what extent they remain automatable. The obtained results are quite encouraging, although finding ways to combine techniques is crucial if one wants to capture a reasonably large class of systems, similarly to what happens in ordinary term rewriting [2]. Another future work includes clarifying the place of SAST in the arithmetical hierarchy, and extending our result for proving AST, not only SAST.

Acknowledgements

This work is partially supported by the ANR projects 14CE250005 ELICA and 16CE250011 REPAS, the FWF project Y757, the JSPS-INRIA bilateral joint research project ‘‘CRECOGI’’, the ERC Consolidator Grant DLV-818616 DIAPASoN, and JST ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603).

References

- [1] G. Agha, J. Meseguer, K. Sen, PMAude: Rewrite-based specification language for probabilistic object systems, *Electron. Notes Theor. Comput. Sci.* 153 (2) (2006) 213–239.
- [2] M. Avanzini, *Verifying Polytime Computability Automatically*, PhD Thesis, University of Innsbruck, 2013.
- [3] M. Avanzini, U. Dal Lago, A. Yamada, On probabilistic term rewriting, in: *Proc. of 14th FLOPS*, in: LNCS, vol. 10818, Springer, 2018, pp. 132–148.
- [4] M. Avanzini, M. Schaper, G. Moser, Modular runtime complexity analysis of probabilistic while programs, in: *Proc. of 10th DICE and 6th FOPARA*, 2019.
- [5] F. Baader, T. Nipkow, *Term Rewriting and All That*, Cambridge University Press, 1998.
- [6] O. Bournez, F. Garnier, Proving positive almost-sure termination, in: *Proc. of 16th RTA*, in: LNCS, vol. 3467, Springer, 2005, pp. 323–337.
- [7] O. Bournez, F. Garnier, Proving positive almost sure termination under strategies, in: *Proc. of 17th RTA*, in: LNCS, vol. 4098, Springer, 2006, pp. 357–371.
- [8] O. Bournez, C. Kirchner, Probabilistic rewrite strategies: applications to ELAN, in: *Proc. of 13th RTA*, 2002, pp. 252–266.
- [9] A. Chakarov, S. Sankaranarayanan, Probabilistic program analysis with martingales, in: *Proc. of 25th CAV*, in: LNCS, vol. 8044, Springer, 2013, pp. 511–526.
- [10] K. Chatterjee, H. Fu, A.K. Goharshady, Termination analysis of probabilistic programs through positivstellensatz’s, in: *Proc. of 28th CAV*, in: LNCS, vol. 9779, Springer, 2016, pp. 3–22.
- [11] U. Dal Lago, M. Zorzi, Probabilistic operational semantics for the lambda calculus, *RAIRO Theor. Inform. Appl.* 46 (3) (2012) 413–450.
- [12] K. De Leeuw, E.F. Moore, C.E. Shannon, N. Shapiro, Computability by probabilistic machines, *Autom. Stud.* 34 (1956) 183–198.
- [13] J. Endrullis, J. Waldmann, H. Zantema, Matrix interpretations for proving termination of term rewriting, *J. Autom. Reason.* 40 (3) (2008) 195–220.
- [14] L.M. Ferrer Fioriti, H. Hermans, Probabilistic termination: soundness, completeness, and compositionality, in: *Proc. of 42nd POPL*, ACM, 2015, pp. 489–501.
- [15] H. Fu, K. Chatterjee, Termination of nondeterministic probabilistic programs, in: *VMCAI*, in: LNCS, vol. 11388, Springer, 2019, pp. 468–490.
- [16] C. Fuhs, J. Giesl, A. Middeldorp, P. Schneider-Kamp, R. Thiemann, H. Zankl, SAT solving for termination analysis with polynomial interpretations, in: *Proc. of 10th SAT*, in: LNCS, vol. 4501, Springer, 2007, pp. 340–354.
- [17] J. Gill, Computational complexity of probabilistic Turing machines, *SIAM J. Comput.* 6 (4) (1977) 675–695.
- [18] I. Gnaedig, Induction for positive almost sure termination, in: *PPDP 2007*, ACM, 2007, pp. 167–178.

- [19] S. Goldwasser, S. Micali, Probabilistic encryption, *J. Comput. Syst. Sci.* 28 (2) (1984) 270–299.
- [20] N.D. Goodman, V.K. Mansinghka, D.M. Roy, K. Bonawitz, J.B. Tenenbaum, Church: a language for generative models, in: *Proc. of 24th UAI*, AUAI Press, 2008, pp. 220–229.
- [21] N. Hirokawa, G. Moser, Automated complexity analysis based on context-sensitive rewriting, in: *RTA-TLCA 2014*, in: LNCS, vol. 8560, 2014, pp. 257–271.
- [22] D. Hofbauer, C. Lautemann, Termination proofs and the length of derivations, in: *Proc. of 3rd RTA*, in: LNCS, vol. 355, Springer, 1989, pp. 167–177.
- [23] C. Jones, G.D. Plotkin, A probabilistic powerdomain of evaluations, in: *Proc. of 4th LICS*, 1989, pp. 186–195.
- [24] B.L. Kaminski, J.-P. Katoen, On the hardness of almost-sure termination, in: *MFCS 2015, Part I*, in: LNCS, vol. 9234, Springer, 2015, pp. 307–318.
- [25] B.L. Kaminski, J.-P. Katoen, C. Matheja, F. Olmedo, Weakest precondition reasoning for expected runtimes of randomized algorithms, *J. ACM* 65 (5) (2018) 30.
- [26] U. Dal Lago, C. Grellois, Probabilistic termination by monadic affine sized typing, in: *Proc. of 26th ESOP*, 2017, pp. 393–419.
- [27] U. Dal Lago, S. Martini, On constructor rewrite systems and the lambda calculus, *Log. Methods Comput. Sci.* 8 (3) (2012).
- [28] D. Lankford, Canonical Algebraic Simplification in Computational Logic, Technical Report ATP-25, University of Texas, 1975.
- [29] D.S. Lankford, On Proving Term Rewriting Systems are Noetherian, Technical Report MTP-3, Louisiana Technical University, 1979.
- [30] S. Lucas, Polynomials over the reals in proofs of termination: from theory to practice, *Inform. Theor. Appl.* 39 (3) (2005) 547–586.
- [31] A. McIver, C. Morgan, B.L. Kaminski, J.-P. Katoen, A new proof rule for almost-sure termination, *J. Proc. ACM Program. Lang.* 2 (POPL) (2018) 33.
- [32] R. Motwani, P. Raghavan, *Randomized Algorithms*, Cambridge University Press, 1995.
- [33] N.C. Ngo, Q. Carbonneaux, J. Hoffmann, Bounded expectations: resource analysis for probabilistic programs, in: *Proc. of 39th PLDI*, 2018, pp. 496–512.
- [34] M.L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, 1st edition, John Wiley & Sons, Inc., New York, NY, USA, 1994.
- [35] M.O. Rabin, Probabilistic automata, *Inf. Control* 6 (3) (1963) 230–245.
- [36] N. Saheb-Djahromi, Probabilistic LCF, in: *MFCS*, in: LNCS, vol. 64, Springer, 1978, pp. 442–451.
- [37] E.S. Santos, Probabilistic Turing machines and computability, *Proc. Am. Math. Soc.* 22 (3) (1969) 704–710.
- [38] Terese, *Term Rewriting Systems*, Cambridge Tracts in Theoretical Computer Science., vol. 55, Cambridge University Press, 2003.
- [39] A. Yamada, K. Kusakari, T. Sakabe, Nagoya termination tool, in: *RTA-TLCA 2014*, in: LNCS, vol. 8560, 2014, pp. 466–475.