

Compositional metric reasoning with Probabilistic Process Calculi

Abstract. Probabilistic process calculi are algebraic theories to specify and verify probabilistic concurrent systems. Bisimulation metric is a fundamental semantic notion that defines the behavioral distance of probabilistic processes. We study which operators of probabilistic process calculi allow for compositional reasoning with respect to bisimulation metric semantics. Moreover, we characterize the distance between probabilistic processes composed by standard process algebra operators.

1 Introduction

Process algebra is undoubtedly one of the most successful formalism to specify and verify concurrent systems. Processes are described as algebraic terms with a formal semantics provided by either operational semantics (e.g. process graph as labelled transition system), denotational semantics (e.g. decorated trace models), or axiomatic semantics (e.g. equational characterization of process equivalence). For many process algebras there are appropriate operational, denotational and axiomatic semantics available such that the respective equivalence notions coincide.

In this paper we study various probabilistic process algebras from the perspective of operational semantics. The operational semantics of a process term is a probabilistic nondeterministic transition system [12] with transitions derived from SOS rules in the probabilistic GSOS format [2, 11]. The SOS rules specify for each process combinator the set of transitions that processes combined by that process combinator can perform. The probabilistic GSOS format is expressive enough for all standard probabilistic process algebra operators.

Behavioral equivalences equate processes that are indistinguishable to any external observer. The most prominent example is bisimulation equivalence [10, 12] which provides a well-established theory of the behavior of probabilistic nondeterministic transition systems. However, bisimulation equivalence is too sensitive to the exact probabilities of transitions. The slightest perturbation of the probabilities can destroy bisimilarity. Bisimulation metric [8] provides a robust semantics for probabilistic processes. It is the quantitative analogue to bisimulation equivalence and assigns to each pair of processes a distance which measures the proximity of their quantitative properties. The distances form a pseudometric with bisimilar processes at distance 0.

In order to specify and verify systems in a compositional manner, it is necessary that the behavioral semantics is compatible with all operators of the language that describe these systems. For behavioral equivalence semantics there is the common agreement that compositional reasoning requires that the considered behavioral equivalence is a congruence wrt. all operators. On the other hand, for behavioral metric semantics there are several proposals of properties that process combinator should satisfy in order to facilitate compositional reasoning. Most prominent examples are non-expansiveness [8]

and non-extensiveness [1]. We discuss these compositionality criteria and propose continuity¹ as the most natural property of process combinators to facilitate compositional reasoning wrt. behavioral pseudometrics. Continuity of a binary process combinator f ensures that for any bisimulation distance ϵ (understood as the admissible tolerance from the operational behavior of the composed process $f(p_1, p_2)$) there are non-zero bisimulation distances δ_1 and δ_2 (understood as the admissible tolerance from the operational behavior of the to be combined processes) such that the distance between $f(p_1, p_2)$ and $f(p'_1, p'_2)$ is at most ϵ whenever p'_1 (resp. p'_2) is in distance of at most δ_1 from p_1 (resp. at most δ_2 from p_2).

Our first main result is that the non-recursive fragments of standard probabilistic process algebras, i.e. process algebras including operators for (nondeterministic and probabilistic variants of) sequential, alternative and parallel composition, allow for compositional reasoning wrt. the compositionality criteria of non-expansiveness (and hence also wrt. continuity). Moreover, recursive operators like Kleene-star iteration and π -calculus bang replication allow for compositional reasoning wrt. the compositionality criteria of continuity (but not wrt. non-expansiveness).

Our second main result is an upper bound on the bisimulation distance between composed processes. The distance between composed processes is a function of the distances between its parts in terms of the ‘degree of continuity’ that the process combinator ensures.

2 Preliminaries

We consider transition systems with process terms as states and a transition relation from states to distributions inductively defined by means of SOS rules. Process terms are inductively defined from base processes that are composed by process combinators. The SOS rules are syntax-driven inference rules that define the behavior of complex processes in terms of the behavior of their components. Technical details may be found in [5, 11].

A bisimulation metric is the quantitative analogue to the relational notion of bisimulation equivalence [7, 8]. A 1-bounded pseudometric d is a bisimulation metric if for all states s and t each transition from s can be mimicked by a transition from t with the same label such that the distance between their derivatives is at most $d(s, t)$.

Definition 1 (Bisimulation metric). *A 1-bounded pseudometric d on $T(\Sigma)$ is a bisimulation metric if for all $t, t' \in T(\Sigma)$ with $d(t, t') < 1$, if $t \xrightarrow{a} \pi$ then there exists a transition $t' \xrightarrow{a} \pi'$ with $\mathbf{K}(d)(\pi, \pi') \leq d(t, t')$ and $\mathbf{K}(d): \Delta(T(\Sigma)) \times \Delta(T(\Sigma)) \rightarrow [0, 1]$ the Kantorovich pseudometric.*

The smallest bisimulation metric is denoted by \mathbf{d} .

We introduce as running example a probabilistic process algebra that comprises many of the probabilistic CCS [2, 6] process combinators. Let Σ_{PA} be the signature

¹ Kim G. Larsen proposed continuity as the canonical property of process combinators in order to facilitate compositional metric reasoning (discussion at the Dagstuhl seminar 14041: ‘Quantitative Models: Expressiveness, Analysis, and New Applications’).

$$\begin{array}{c}
\frac{a. \bigoplus_{i=1}^n [p_i]x_i \xrightarrow{a} \sum_{i=1}^n p_i \delta(x_i)}{\quad} \quad \frac{x \xrightarrow{a} \mu \quad a \neq \surd}{x; y \xrightarrow{a} \mu; \delta(y)} \quad \frac{x \xrightarrow{\surd} \mu \quad y \xrightarrow{a} \nu}{x; y \xrightarrow{a} \nu} \\
\\
\frac{x \xrightarrow{a} \mu}{x + y \xrightarrow{a} \mu} \quad \frac{y \xrightarrow{a} \nu}{x + y \xrightarrow{a} \nu} \quad \frac{x \xrightarrow{a} \mu \quad y \not\xrightarrow{a}}{x +_p y \xrightarrow{a} \mu} \quad \frac{x \not\xrightarrow{a} \quad y \xrightarrow{a} \nu}{x +_p y \xrightarrow{a} \nu} \quad \frac{x \xrightarrow{a} \mu \quad y \xrightarrow{a} \nu}{x +_p y \xrightarrow{a} \mu \oplus_p \nu} \\
\\
\frac{x \xrightarrow{a} \mu \quad y \xrightarrow{a} \nu}{x | y \xrightarrow{a} \mu | \nu} \quad \frac{x \xrightarrow{a} \mu}{x^{n+1} \xrightarrow{a} \mu; \delta(x^n)} \quad \frac{x \xrightarrow{a} \mu}{x^\omega \xrightarrow{a} \mu; \delta(x^\omega)}
\end{array}$$

Table 1: Standard probabilistic process combinators

with the following operators: i) 0 (stop process); ii) a family of n -ary prefix operators $a.([p_1]_-\oplus\cdots\oplus[p_n]_-)$ with $a \in A$, $n \geq 1$, $p_1, \dots, p_n \in (0, 1]$ and $\sum_{i=1}^n p_i = 1$; iii) binary operators $_-$; $_-$ (sequential composition), $_-$ $_-$ (alternative composition), $_-$ $+_p$ $_-$ (probabilistic alternative composition), $_-$ $|$ $_-$ (synchronous parallel composition), iv) $_-$ ^{n} (finite iteration), and v) $_-$ ^{ω} (infinite iteration). The PTSS $P_{PA} = (\Sigma_{PA}, A, R_{PA})$ is given by the rules R_{PA} in Table 1. The probabilistic prefix operator expresses that the term $a.([p_1]t_1 \oplus \cdots \oplus [p_n]t_n)$ can perform action a and evolves to process t_i with probability p_i . The sequential composition and the alternative composition are as usual. The synchronous parallel composition $t | t'$ describes the simultaneous evolution of processes t_1 and t_2 . The probabilistic alternative composition $t +_p t'$ evolves to the probabilistic choice between the evolution of process t (with probability p) and the evolution of process t' (with probability $1 - p$) for actions which can be performed by both processes. For actions which only one of the summands can perform, the probabilistic alternative composition $t +_p t'$ behaves just like the nondeterministic alternative composition $t + t'$. The finite iteration t^n (resp. infinite iteration t^ω) of process t expresses that t is performed n times (resp. infinitely often).

3 Specification of Compositional Process Combinators

SOS researchers developed over the last decades numerous rule formats that allow for compositional reasoning wrt. various behavioral equivalences (e.g. [3–5, 9, 11]). We develop rule and specification formats that allow for compositional reasoning wrt. bisimulation metric. More precisely, for each compositionality property we provide a specification format such that the specified process combinators satisfy the respective compositionality property.

In essence, SOS rules specifying non-extensive, non-expansive or Lipschitz-continuous process combinators ensure that combined processes replicate its parts only a limited number of times along their evolution. In detail, a rule specifies a p -non-extensive operator if at most one of the combined processes evolves, it specifies a non-expansive operator if only one instance of each of the combined processes evolves, and it specifies a Lipschitz-continuous operator if only one source and one derived instance of each of

the combined processes evolves. The essence of specifications of uniform-continuous process combinators is that both the replication of source processes as well as the depth of the derivation tree of each transition is finitely bounded.

We will see that the non-deterministic and probabilistic alternative composition is ∞ -non-extensive, that all variants of parallel composition are non-expansive but not p -non-extensive if $p > 1$, and that the operators of iteration and replication are Lipschitz-continuous.

4 Distance between probabilistic processes

The compositionality properties of process combinators allows us to derive an upper bound on the distance of composed processes.

The distance of non-recursively composed processes is as follows:

Proposition 1. *Let $s_i, t_i \in \mathcal{T}(\Sigma_{PA})$ be non-recursive nondeterministic probabilistic processes. Then*

- $\mathbf{d}(a. \bigoplus_{i=1}^n [p_i]s_i, a. \bigoplus_{i=1}^n [p_i]t_i) \leq \sum_{i=1}^n p_i \mathbf{d}(s_i, t_i)$
- $\mathbf{d}(s_1; s_2, t_1; t_2) \leq 1 - (1 - \mathbf{d}(s_1, t_1))(1 - \mathbf{d}(s_2, t_2))$
- $\mathbf{d}(s_1 + s_2, t_1 + t_2) \leq \max(\mathbf{d}(s_1, t_1), \mathbf{d}(s_2, t_2))$
- $\mathbf{d}(s_1 +_p s_2, t_1 +_p t_2) \leq p\mathbf{d}(s_1, t_1) + (1 - p)\mathbf{d}(s_2, t_2)$
- $\mathbf{d}(s_1 \mid s_2, t_1 \mid t_2) \leq 1 - (1 - \mathbf{d}(s_1, t_1))(1 - \mathbf{d}(s_2, t_2))$

The distance of recursively composed processes is as follows:

Proposition 2. *Let $s, t \in \mathcal{T}(\Sigma_{PA})$ be non-recursive nondeterministic probabilistic processes. Then*

- $\mathbf{d}(s^n, t^n) \leq 1 - (1 - \mathbf{d}(s, t))^n$
- $\mathbf{d}(s^\omega, t^\omega) \leq 1$

References

1. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: Computing behavioral distances, compositionally. In: Proc. MFCS'13, pp. 74–85. Springer (2013)
2. Bartels, F.: On Generalised Coinduction and Probabilistic Specification Formats. Ph.D. thesis, VU University Amsterdam (2004)
3. Bloom, B., Fokkink, W., van Glabbeek, R.J.: Precongruence formats for decorated trace semantics. ACM TOCL 5, 26–78 (2004)
4. Bloom, B., Istrail, S., Meyer, A.R.: Bisimulation can't be traced. J. ACM 42, 232–268 (1995)
5. D'Argenio, P.R., Lee, M.D.: Probabilistic transition system specification: Congruence and full abstraction of bisimulation. In: Proc. FoSSaCS'12. LNCS, vol. 7213, pp. 452–466. Springer (2012)
6. Deng, Y., Du, W.: Probabilistic barbed congruence. In: Proc. QAPL'07. ENTCS, vol. 190, pp. 185–203 (2007)
7. Deng, Y., Du, W.: Logical, metric, and algorithmic characterisations of probabilistic bisimulation. Tech. Rep. CMU-CS-11-110, CMU (March 2011)

8. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled markov processes. TCS 318(3), 323–354 (2004)
9. Groote, J.F.: Transition system specifications with negative premises. TCS 118(2), 263–299 (1993)
10. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. I&C 94, 1–28 (1991)
11. Lee, M.D., Gebler, D., D’Argenio, P.R.: Tree rules in probabilistic transition system specifications with negative and quantitative premises. In: Proc. EXPRESS/SOS’12. EPTCS, vol. 89, pp. 115–130 (2012)
12. Segala, R.: Modeling and Verification of Randomized Distributed Real-Time Systems. Ph.D. thesis, MIT (1995)