

ε -Distance via Lévy-Prokhorov Lifting

Josée Desharnais 

Laval University, Québec, Canada

Ana Sokolova 

University of Salzburg, Austria

Abstract

The most studied and accepted pseudometric for probabilistic processes is one based on the Kantorovich distance between distributions. It comes with many theoretical and motivating results, in particular it is the fixpoint of a given functional and defines a functor on (complete) pseudometric spaces. It is also the foundation for a categorical lifting of pseudometrics.

Other notions of behavioural pseudometrics have also been proposed, one of them (ε -distance) based on ε -bisimulation. ε -Distance has the advantages that it is intuitively easy to understand, it relates systems that are conceptually close (for example, an imperfect implementation is close to its specification), and it comes equipped with a natural notion of ε -coupling. Finally, this distance is easy to compute.

We show that ε -distance is also the greatest fixpoint of a functional and provides a functor. The latter is obtained by replacing the Kantorovich distance in the lifting functor with the Lévy-Prokhorov distance. In addition, we show that ε -couplings and ε -bisimulations have an appealing coalgebraic characterization.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Probabilistic computation; Theory of computation $\rightarrow$ Categorical semantics; Theory of computation $\rightarrow$ Program verification

Keywords and phrases Lévy-Prokhorov metric, behavioural distance, epsilon-bisimulation, reactive probabilistic transition systems, discrete labelled Markov processes, coalgebraic epsilon-(bi)simulation

Digital Object Identifier 10.4230/LIPIcs.CSL.2026.26

Related Version *Full Version:* <https://arxiv.org/abs/2507.10732>

Funding *Josée Desharnais:* Work funded by NSERC grant.

Acknowledgements This work was partly done during a sabbatical of Josée Desharnais at the University of Salzburg, as well as during Dagstuhl Seminar 24432 and the Bellairs Workshop on Quantitative Reasoning 2025. We thank these venues, the organizers, and the participants for providing a perfect working environment. We also thank Matteo Mio and Franck van Breugel for some fruitful and motivating discussions, and the anonymous reviewers for their insightful and constructive suggestions.

1 Introduction

Probabilistic systems [28, 38, 48, 7, 8, 44] and their behaviour have been the object of study of over thirty years in the area of formal verification and analysis of systems. They are used to represent uncertainty, incomplete information, as well as randomized behaviour. One important direction, in order to prove that systems behave the same, is the study of behavioural equivalences: these identify states with (exactly) the same behaviour. Behavioural equivalences are an elegant way to analyse and compare behaviour of systems, they have nice foundations in concurrency theory [41, 60] as well as elegant abstract generalizations in the theory of coalgebras [20, 49, 32, 36, 33, 35, 13, 14]. However, as was already observed in [29], in a non-exact world, e.g., when the probabilities in the model are approximate and not exactly known, or estimated by sampling, or when small differences should not be considered the same as large differences, behavioural equivalences may be too strong.



© Josée Desharnais and Ana Sokolova;

licensed under Creative Commons License CC-BY 4.0

34th EACSL Annual Conference on Computer Science Logic (CSL 2026).

Editors: Stefano Guerrini and Barbara König; Article No. 26; pp. 26:1–26:24

Leibniz International Proceedings in Informatics




LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

One solution is to employ a distance, actually a pseudometric, providing a quantitative notion of how much states in a probabilistic model differ from one another, or how close they are to each other – distance zero corresponding to bisimulation equivalence. The study of pseudometrics [22, 58, 24, 18, 59, 56, 26, 54, 55] and quantitative theories [39, 40, 5, 6, 9, 42, 16, 43] has been a fruitful one in the past decade(s).

Historically, the impulse to define a behavioral distance between probabilistic processes became imperative for those with a *continuous state space and continuous probability distributions*, called Labelled Markov Processes (LMPs). It started with the idea of defining a real-valued logic, as proposed by Kozen [37], to express properties satisfied by processes: Taking the supremum over the differences of the values of formulas was a natural way of defining a pseudometric, see Desharnais et al. [22, 24]. The authors later on observed that this distance was the greatest fixpoint of some functional [25] and, in between and later, van Breugel and Worrell [58, 59] and van Breugel et al. [56], proved that this gives a functor (monad) that is a lifting of the distribution functor (monad) using the Kantorovich metric on distributions. Moreover, the distance is the one obtained using final coalgebra semantics. Since then, this pseudometric has been studied further in different contexts and algorithms have been proposed to compute it [53, 5, 6]. The distance comes with many theoretical and motivating results, and has also been extended by Baldan et al. [9] and Sprunger et al. [51, 52] to a generic Kantorovich-style (and dually Wasserstein) lifting of arbitrary functors parametric in certain evaluation maps – the two liftings coincide on distributions. With all these seals of approval, the Kantorovich behavioural distance has imposed itself as the one to use and study. One aspect that it lacks is an intuitive accompanying notion of approximate bisimulation.

Approximate bisimulations provide another strategy to circumvent the non-robustness of bisimulation equivalences. They have been studied in non-probabilistic systems [61, 30, 51], and in probabilistic systems [26, 54, 27, 1, 2, 10, 50]. We focus on the notion of ε -bisimulation defined in [26]. One aim of such approximate bisimulations is to unite the best of both worlds: give a relational structure for reasoning about systems and at the same time define a distance. States related by an ε -bisimulation are at distance at most ε from one another. The advantages of such a distance are its intuitive nature and its ability to relate systems that are almost the same in structure, as the next example will illustrate. Finally, the distance can be computed relatively easily.

For comparison with the Kantorovich behavioural distance, consider the imperfect channel (when $\varepsilon > 0$) depicted in Figure 1. The ε -distance of this channel to a perfect channel (one with $\varepsilon = 0$) is ε , and it provides a simple example that imperfect implementation can be



■ **Figure 1** A simple channel that fails to send a token with probability ε .

considered close to its specification. On the other hand, the Kantorovich distance (without discount) gives distance 1 to the pair of perfect and imperfect channels, which reflects the fact that the two channels have a very different behavior on the long run – see Example 1 for more details. Interestingly, ε -distance has an almost build-in notion of continuity: Two channels with close values of ε are close in the ε -distance as well.

However, the approximate equivalences had no categorical formulation attached and it was not clear whether the associated distance shares some of the generic good properties identified for the Kantorovich distance. Our main motivation for this work was to investigate whether we can remedy this situation. Indeed, we provide the missing abstract characterization of ε -bisimulation in two directions: by obtaining a fixpoint characterization of the induced distance and by giving a coalgebraic view following Aczel-Mendler coalgebraic bisimilarity. The fixpoint characterization is the main result of the paper. For this, we show how to replace the Kantorovich metric on distributions with the *Lévy-Prokhorov metric* on distributions, and show that the distance obtained is indeed the ε -distance. It is remarkable how well the Lévy-Prokhorov distance fits into the definition of ε -bisimulation: We show that ε -distance is the greatest fixpoint of a suitable functional, and that the lifting also lifts the discrete distributions functor to a functor on pseudometric spaces. We also show a result that has no matching for the Kantorovich functional: any fixpoint distance of the Lévy-Prokhorov functional defines an ε -bisimulation. Applications of our results may take different routes: On the practical side, the distance could be computed by an iterative algorithm given a suitable fixpoint theorem. On the theoretical side, the fixpoint characterization may open the way for studying new problems, one example being ε -bisimulations up-to as an acceleration to computing ε -bisimilarity, along the lines of [12].

Our observation may also open a new way of studying other pseudometrics on probabilistic systems that would be constructed from other basic distances on distributions.

Unlike in the case of the Kantorovich lifting, the Lévy Prokhorov lifting does not yield a monad on the category of pseudometric spaces, see Section 3.4. As a consequence, we could not follow the path of [56] and reproduce their results on the final coalgebra. Instead, we take a different way and describe ε -(bi)simulations as coalgebraic (bi)simulations, following and extending the framework of Hasuo [32, 33] and the original results of Hughes and Jacobs [34]. We only present the necessary abstraction needed to discuss ε -(bi)simulations on discrete time Markov chains and labelled Markov processes, generalizations of these notions are possible and we plan to elaborate on them in follow-up work.

All proofs omitted for space reasons are in the appendix.

2 The objects of interest

Our development concerns discrete labelled Markov processes (LMPs) also called reactive probabilistic transition systems. They consist of a set of states S and transitions $\tau_a : S \rightarrow \mathcal{DS}$ labelled with actions a of some set $\mathcal{A}$, where $\mathcal{DS}$ is the set of discrete sub-distributions on S :

$$\mathcal{DS} = \{\varphi : S \rightarrow [0, 1] \mid \sum_{s \in S} \varphi(s) \leq 1\}.$$

Note that the size of the state set need not be restricted (for the sums to be defined), as the sum is defined as $\sum_{x \in S} \varphi(x) = \sup_{\text{finite } X \subseteq S} (\sum_{x \in X} \varphi(x))$. It is easy to prove that discrete probability distributions always have a countable support, i.e., they assign non-zero probability to at most countably many elements of S .

For $s, t \in S$ and $a \in \mathcal{A}$, the value $\tau_a(s)(t)$ (also written $\tau_a(s, t)$) encodes the probability of jumping from s to t when action a is taken. Figure 1 shows such a system with actions *send* and *ack*. The first transition is a subprobability: the missing probability represents the probability of failure. If $X \subseteq S$, we write $\tau_a(s)(X)$ for $\sum_{t \in X} \tau_a(s)(t)$. Coalgebraically, LMPs, for a set of actions or labels $\mathcal{A}$, are arrows $\tau : S \rightarrow (\mathcal{DS})^{\mathcal{A}}$. We write $\tau_a(s)$ for the subdistribution $\tau(s)(a)$. In case of just one label, LMPs are (discrete time) Markov chains,

i.e., coalgebras of the functor $\mathcal{D}$. The reason for working with subdistributions is that they come equipped with a non-trivial pointwise order. Bisimilarity of LMPs is the classical Larsen-Skou [38] bisimilarity.

We denote by $\mathbf{PMet}_1$ the category of 1-bounded pseudometric spaces. It has as objects pseudometric spaces, which are pairs $\langle X, d \rangle$ of a set X and a function $d: X \times X \rightarrow [0, 1]$ which satisfies $d(x, x) = 0$, is symmetric: $d(x, y) = d(y, x)$ and satisfies the triangular inequality: $d(x, y) + d(y, z) \geq d(x, z)$ for all $x, y, z \in X$. As arrows, $\mathbf{PMet}_1$ has non-expansive maps $f: \langle X, d_X \rangle \rightarrow \langle Y, d_Y \rangle$, i.e., functions $f: X \rightarrow Y$ satisfying $d_Y(f(x_1), f(x_2)) \leq d_X(x_1, x_2)$. A metric or distance is a pseudometric that additionally satisfies $d(x, y) = 0$ iff $x = y$. Of particular interest for LMPs are pseudometrics whose kernel is bisimilarity, that is, $d(x, y) = 0$ iff x and y are bisimilar. As we only deal with pseudometrics here, we often use the word distance or metric as a shorthand for pseudometric. Following [59, 56], we will use the opposite pointwise order¹ on distances defined by:

$$d_1 \sqsubseteq d_2 \iff \forall x, y \in X : d_1(x, y) \geq d_2(x, y).$$

However, whenever it is clearer, we may use the direct pointwise order too: $d_1 \geq d_2 \Leftrightarrow d_1 \sqsubseteq d_2$.

2.1 The Kantorovich pseudometric

The observation that bisimulation is too strong in the context of probabilistic systems has led to the idea of defining a pseudometric that would give zero distance to bisimilar states. Initially, this pseudometric was defined using a real-valued logic. A set $\mathcal{F}$ of functionals were defined from states of LMPs to $[0, 1]$, with the following syntax, mimicking logical formulas [24]:

$$f := 1 \mid \inf(f_1, f_2) \mid 1 - f \mid f \ominus q \mid \langle a \rangle f \text{ with } q \in [0, 1] \cap \mathbb{Q}.$$

We omit the semantics, but the next example will give a taste of it. A distance emerged naturally as

$$d_K(s, t) = \sup_{f \in \mathcal{F}} |f(s) - f(t)|.$$

► **Example 1.** As an example, we can look back at the channel of Figure 1. The maximum difference over functionals between c_ϵ and c_γ , with $\epsilon, \gamma > 0$ is over the functional $\langle \text{send} \rangle \langle \text{ack} \rangle 1$, which evaluates to $1 - \epsilon$ and $1 - \gamma$, respectively, and thus yields the distance $d_K(c_\epsilon, c_\gamma) = |\epsilon - \gamma|$. Between c_ϵ and c_0 , the functional $(\langle \text{send} \rangle \langle \text{ack} \rangle)^n 1$ evaluates to $(1 - \epsilon)^n$ on c_ϵ , and to 1 on c_0 . So the supremum results in $d_K(c_\epsilon, c_0) = 1$.

Later on, it was proven that the functionals could be any non-expansive maps and that this distance is the distance obtained via a final coalgebra construction for a functor (monad) that is a lifting of the Distribution functor (monad) to pseudometric spaces using the Kantorovich metric, by van Breugel and Worrell [59] and van Breugel et al. [56, 57].

This distance is a fixpoint of the functional on distances on the states of an LMP $\tau : S \rightarrow (\mathcal{D}S)^A$

$$\Delta_K(d)(s_0, s_1) = \sup_{a \in A} \delta_K^d(\tau_a(s_0), \tau_a(s_1)) \quad (1)$$

where δ_K^d is the well-known Kantorovich distance between two distributions. We do not need its definition here, the interested reader can find it in, e.g., [56].

¹ One reason for this choice is to obtain the distance as the greatest, rather than least, fixpoint – just like bisimilarity is the greatest fixpoint of a suitable functional.

The distance on distributions is just a parameter in Equation (1); changing the Kantorovich distance δ_K^d to another distance on distributions and looking for a fixpoint will give us a new distance on states. In fact we are interested in the question formulated the other way around: find a distance on distributions for which the fixpoint is the ε -distance d^* , defined in the next section. Similarly, the functional Δ_K was discovered after the behavioural distance d_K was introduced. However, it could have been done the other way around, choosing a distance on distributions, and looking for the resulting distance on states, as is done by Baldan et al. [9]. We discuss this in more detail again in Section 3 below.

2.2 ε -Bisimulation

Bisimulation, being too strong for the comparison of states in quantitative systems, has also led to relaxing the definition of simulation and bisimulation itself to approximate relations. As mentioned in the introduction, there are a few of those definitions of approximate bisimulations, but we are interested in the following.

► **Definition 2** ([26]). *Let $\tau : S \rightarrow (\mathcal{DS})^A$ be an LMP and let $\varepsilon \in [0, 1]$. A relation $R \subseteq S \times S$ is an ε -simulation if whenever sRt , then for all $a \in A$, for all $X \subseteq S$*


$$\tau_a(s)(X) \leq \tau_a(t)(R(X)) + \varepsilon, \quad \text{where } R(X) = \{y \mid \exists x \in X : xRy\}.$$

If R is symmetric, it is an ε -bisimulation. A state s is ε -simulated by state t , written $s \prec_\varepsilon t$, if sRt for some ε -simulation R . If sRt for R ε -bisimulation, we write $s \sim_\varepsilon t$ and we say that s and t are ε -bisimilar.

As expected, ordinary (bi)simulation on LMPs [38, 11] is simply 0-(bi)simulation. This definition has an extension to nondeterministic and probabilistic finite systems [54].

The operation $R(X)$ on a set X is what restricts this work to discrete distributions. In general, if X is measurable, we may not have $R(X)$ measurable. Working in analytic spaces, as was done before for LMPs [23] may solve this issue, but we leave it for future work. In particular, one could ask $R(X)$ to be measurable whenever X is, as done in [10], but many results are not proven in that case, like the logical characterisation, or Theorem 5 below.

► **Example 3.** Consider the following example, with $\gamma \in (0, 1]$.



The relation $R = \{(s, t), (s_1, t_1)\}$ is an ε -simulation for $\varepsilon = \gamma$, and so is $R \cup R^{-1}$. Hence, $s \prec_\gamma t$, $t \prec_\gamma s$ and $s \sim_\gamma t$. However, the situation is different for $\varepsilon < \gamma$: for example, taking $\varepsilon = 0$, we observe that $s \prec_0 t$ (with the relation $R \cup \{(s_2, t_1)\}$). However, $t \not\prec_0 s$, because a relation R' relating t and s would have to also relate t_1 to s_1 and to s_2 , if $\gamma < 1$, and t_1 to s_2 if $\gamma = 1$. Indeed, taking $X = \{t_1\}$, we need $\tau_a(t)(X) \leq \tau_a(s)(R'(X)) + 0$, but for this to be the case, we would need $\tau_a(s)(R'(X)) = 1$ and hence R' must include the pairs (t_1, s_1) and (t_1, s_2) if $\gamma < 1$, and otherwise must include (t_1, s_2) if $\gamma = 1$. So, in any case $(t_1, s_2) \in R'$, but the pair (t_1, s_2) cannot be related by any ε -simulation relation R' for $\varepsilon < 1$ (which is the case here as $\varepsilon = 0$), as, for example, taking again $X = \{t_1\}$, we would have $\tau_a(t_1)(X) = 1$ but $\tau_a(s_2)(R'(X)) = 0$. Hence, also $s \not\prec_0 t$.

In contrast to the case of 0-bisimilarity, two-way ε -similarity is not ε -bisimilarity [26, 54]. From the notion of ε -bisimilarity, a pseudometric arises naturally by taking the infimum.

► **Definition 4** (The ε -distance d^* [26]). Let $\tau : S \rightarrow (\mathcal{DS})^A$ be an LMP. The pseudometric d^* on $S \times S$, called ε -distance, is defined as follows:

$$\begin{aligned} d^* : S \times S &\rightarrow [0, 1] \\ (s, t) &\mapsto \inf \{ \varepsilon \in [0, 1] \mid s \sim_\varepsilon t \}. \end{aligned}$$

The function d^* is a pseudometric on the states of the considered LMP: the triangle inequality comes from the well-known fact, see e.g. [26], that $s \sim_{\varepsilon_1} u$ and $u \sim_{\varepsilon_2} t$ imply $s \sim_{\varepsilon_1 + \varepsilon_2} t$. (The same property holds for $\prec_\varepsilon$ as well.) Indeed, $\prec_\varepsilon$ and $\sim_\varepsilon$ are not transitive. Instead, these relations are *entourages* that form a *uniform structure*, or *uniformity*, cf. [15].

The kernel of d^* , that is, the set of pairs at zero distance, is $\sim_0$, the bisimilarity relation on LMPs.

► **Theorem 5** ([26]). $d^*(s, t) = 0 \Leftrightarrow s \sim_0 t$.

One of the nice properties of this distance is its relatively easy computability, especially regarding computation “by hand”. It suffices to define an ε -bisimulation to obtain an upper bound to the distance between states. In Example 3, we have $d^*(s, t) = \gamma$, and, indeed, we easily see that the probabilities of s and t seen as processes are within γ . Also here, the Kantorovich distance between these states is 1. One way to see that is by using the functional $\langle a \rangle^n 1$, which evaluates to $(1 - \varepsilon)^n$ on s , and to 1 on t . So $d_K(s, t) = 1$. Similarly, as already noted in the introduction, for the channel in Figure 1, we can prove that $d^*(c_\varepsilon, c_\gamma) = |\varepsilon - \gamma| = d_K(c_\varepsilon, c_\gamma)$, for $\varepsilon, \gamma > 0$, but $d^*(c_\varepsilon, c_0) = \varepsilon \neq d_K(c_\varepsilon, c_0)$. In fact, ε -distance is incomparable to both the discounted and the undiscounted Kantorovich distance, as argued in [26]. That paper also gives a straightforward polynomial algorithm to compute d^* . The computation of the Kantorovich distance, discounted or not, has been the subject of many papers, which also led to polynomial algorithms [53].

3 The Lévy-Prokhorov distance lifting

We will now explain in details what we think should be a behavioral pseudometric on LMPs. These details may not be new, or may appear obvious to the expert reader, but we find it useful to spell them out here. When looking for a distance $d^\dagger$ on the states of an LMP, one observes that states are both targets of distributions and (a set of) distributions themselves (since they are *defined* by their outgoing transition distributions): that is, in Example 3, the states s and t can be viewed as distributions over the set $\{s_1, s_2, t_1\}$. As there are already a few distances on distributions that were studied outside computer science and concurrency theory, they give a starting point for distances on states: we could say, given such a distance δ :

$$d^\dagger(s, t) := \sup_{a \in A} \delta(\tau_a(s), \tau_a(t)), \quad \text{for } s, t \in S. \quad (2)$$

We will use the letter d for distances on states, and the symbol δ for distances on distributions. Analysing the options of this equation in full generality is outside the scope of the current paper but some distances, such as δ_K^d (and the one we introduce below δ_{LP}^d) have a particularity. They are defined using a parameter d , a basic distance on the space where the distributions are defined. This is desirable for concurrency theory because distributions on states that are different but close (or even bisimilar!) should also be close: so a starting distance d on the state space is of key importance to account for these similarities between states/processes. Note that a few papers [30, 19, 51] start from distances on the states that account for extra information (e.g. a distance on the labels/observations that are attached to states) – but this goes beyond the behaviour of states that we want to capture here (although they constitute an interesting line of work to extend our method).

A needed property for $d^\dagger$ to be a *behavioural* distance, as also pointed out in [9] for δ_K^d , is that it accords with the starting distance d used in δ^d , as follows.

$$d^\dagger(s, t) = \sup_{a \in \mathcal{A}} \delta^{d^\dagger}(\tau_a(s), \tau_a(t)), \quad \text{for } s, t \in S. \quad (3)$$

This says that the distance between states viewed as simple members of the space (the left hand part of the equality), is the same as their distance when viewed as processes, that is, their outgoing transition distributions. So this really says that $d^\dagger$ treats states according to their behaviour. Technically, this is saying that we are looking for a distance fixpoint $d^\dagger$ of the functional

$$\Delta(d)(s, t) := \sup_{a \in \mathcal{A}} \delta^d(\tau_a(s), \tau_a(t)), \quad \text{for } s, t \in S. \quad (4)$$

Of course, by taking $\delta^d := \delta_K^d$, this instantiates exactly to Equation (1), the fixpoint property of the Kantorovich behavioural distance, observed by [25, 59] and explicitly proven in [57]. One of our goals in this work was to find a similar functional that would have the ε -distance d^* as its greatest fixed point. That is, we are seeking a suitable distance δ^d on distributions.

► **Remark 6.** Another property that one might consider natural and one might expect from such a distance δ^d is that the distance between the Dirac distributions on states is the same as the starting distance between them. That is, for a distance d on states, one might expect

$$d(s, t) = \delta^d(1_s, 1_t), \quad \text{for } s, t \in S, \quad (5)$$

where we write 1_s for the Dirac distribution on $s \in S$. This is an interesting property of the lifting δ^d – it is a stronger version of the non-expansiveness of the unit of the distribution monad, showing that the unit is an isometry, and we will return to it in Section 3.4 below where we show that it holds for the Lévy-Prokhorov distance, for any starting distance d on states. It also holds for the (undiscounted) Kantorovich distance – this is easy to prove using the Wasserstein formulation of the Kantorovich distance, whereas it does not hold for the total variation distance [43]. Therefore, this property is neither unique to the Lévy-Prokhorov distance that we are interested in, nor it is “behavioural” in the sense that we explained so far: it does not take into account the transition structure τ of the states. For a somewhat behavioural explanation, note that combining it with Equation (2), this condition implies that a pair of states s', t' each having a single outgoing transition as a Dirac on state s and t respectively, would get the same distance as s and t . In particular, such a distance does not discount the future.

3.1 The Lévy-Prokhorov distance on distributions

While examining equation (3) and trying to make d^* fit into it as $d^\dagger$ we came up with the following distance on distributions. Only afterwards, we discovered that the distance was actually known as the Lévy-Prokhorov lifting of a distance d to distributions. In this section we introduce the distance and provide an example that illustrates it.

► **Definition 7** (Lévy-Prokhorov distance [45]). *Let $\langle S, d \rangle$ be a pseudometric space; we endow $\mathcal{DS}$ with the pseudometric (or distance) $\delta_{LP}^d : \mathcal{DS} \times \mathcal{DS} \rightarrow [0, 1]$, defined as*

$$\delta_{LP}^d(\mu_0, \mu_1) = \inf\{\varepsilon \mid \forall X \subseteq S : \mu_i(X) \leq \mu_{1-i}(X_\varepsilon^d) + \varepsilon, \text{ for } i = 0, 1\},$$


where $X_\varepsilon^d = \{y \mid \exists x \in X : d(x, y) < \varepsilon\}$.

We call δ_{LP}^d the Lévy-Prokhorov (LP, for short) distance.

The strict inequality in the definition of X_ε^d is necessary for the proof of Theorem 11.

In the following example we define simple probability measures that help illustrate the need for the extra “ $+\varepsilon$ ” in this definition and the need for the “ ε ball around X ”. At first sight, δ_{LP}^d looks very much like the total variation distance, but this “ ε ball around X ” makes it very different.

► **Example 8.** Consider the set of states $S = \{x_\gamma \mid \gamma \in [0, 1]\}$ and a distance on these states given by $d(x_\gamma, x_\xi) = |\gamma - \xi|$, for $\gamma, \xi \in [0, 1]$. In a probabilistic transition system, one could imagine that x_γ has an a -loop to itself with probability $1 - \gamma$, $\gamma \in [0, 1]$, as depicted below (adding transitions out of the states to help see their differences as processes in a transition system). We now define a family of distributions on these states, and we picture them below as states with an outgoing transition without label (one could imagine an a -label). Let $\nu_\gamma = \gamma 1_{x_1} + (1 - \gamma) 1_{x_\gamma}$, for $\gamma \in [0, \frac{1}{2}]$, with 1_x the Dirac measure on x . The distributions ν_γ and ν_0 are illustrated below. For a fixed γ , these two distributions are actually non-zero on a three-state space $S' = \{x_0, x_1, x_\gamma\}$.



We show that $\delta_{LP}^d(\nu_\gamma, \nu_0) = \gamma$. Let $1 - \gamma > \varepsilon > \gamma$. We show that all such ε satisfy the inequalities in the definition of δ_{LP}^d and any $\varepsilon \leq \gamma$ does not, so the distance, being the infimum over all such values, is indeed γ . In particular, we need to check that for all $X \subseteq S$, $\nu_\gamma(X) \leq \nu_0(X_\varepsilon^d) + \varepsilon$ and $\nu_0(X) \leq \nu_\gamma(X_\varepsilon^d) + \varepsilon$. The cases $X = S$ and $X = \emptyset$ are trivial, since the distributions are full. For the sets X of size two, we have

For $X := \{x_0, x_1\}$:	$\nu_\gamma(X) = \gamma$	$\leq 1 + \varepsilon$	$= \nu_0(S) + \varepsilon$	$= \nu_0(X_\varepsilon^d) + \varepsilon$
	$\nu_0(X) = 1$	$\leq 1 + \varepsilon$	$= \nu_\gamma(S) + \varepsilon$	$= \nu_\gamma(X_\varepsilon^d) + \varepsilon$
For $X := \{x_1, x_\gamma\}$:	$\nu_\gamma(X) = 1$	$\leq 1 + \varepsilon$	$= \nu_0(S) + \varepsilon$	$= \nu_0(X_\varepsilon^d) + \varepsilon$
	$\nu_0(X) = 0$	$\leq 1 + \varepsilon$	$= \nu_\gamma(S) + \varepsilon$	$= \nu_\gamma(X_\varepsilon^d) + \varepsilon$
For $X := \{x_0, x_\gamma\}$:	$\nu_\gamma(X) = 1 - \gamma$	$\leq 1 + \varepsilon$	$= \nu_0(\{x_0, x_\gamma\}) + \varepsilon$	$= \nu_0(X_\varepsilon^d) + \varepsilon$
	$\nu_0(X) = 1$	$\leq (1 - \gamma) + \varepsilon$	$= \nu_\gamma(\{x_0, x_\gamma\}) + \varepsilon$	$= \nu_\gamma(X_\varepsilon^d) + \varepsilon$

The last inequality is satisfied thanks to the room given by the “ $+\varepsilon$ ”. For sets of size one, omitting the checks where the probability is zero on the left-hand side of the inequality:

For $X := \{x_1\}$:	$\nu_\gamma(X) = \gamma$	$\leq 0 + \varepsilon$	$= \nu_0(\{x_1\}) + \varepsilon$	$= \nu_0(X_\varepsilon^d) + \varepsilon$
For $X := \{x_\gamma\}$:	$\nu_\gamma(X) = 1 - \gamma$	$\leq 1 + \varepsilon$	$= \nu_0(\{x_\gamma, x_0\}) + \varepsilon$	$= \nu_0(X_\varepsilon^d) + \varepsilon$
For $X := \{x_0\}$:	$\nu_0(X) = 1$	$\leq 1 - \gamma + \varepsilon$	$= \nu_\gamma(\{x_\gamma, x_0\}) + \varepsilon$	$= \nu_\gamma(X_\varepsilon^d) + \varepsilon$

In the one but last inequality, it is crucial that x_γ be within ε of x_0 , which includes x_0 in the ε ball around x_γ . With X in place of X_ε^d , the inequality would not be satisfied, and not even with $X_{\gamma/2}^d$, a set where states at distance zero are included (like bisimilar states – of which there are none here), because we would obtain:

$$\nu_\gamma(\{x_\gamma\}) = 1 - \gamma \not\leq 0 + \varepsilon = \nu_0(\{x_\gamma\}_{\gamma/2}^d) + \varepsilon.$$

The last inequality also illustrates this, and also the need for the “ $+\varepsilon$ ”, as the ε -ball itself is not enough for the inequality to be satisfied. Moreover, for a pair of distributions ν_γ, ν_ξ in our family one can similarly prove that $\delta_{LP}^d(\nu_\gamma, \nu_\xi) = |\gamma - \xi|$.

Finally, we note that for $\varepsilon \leq \gamma$ we have $X_\varepsilon^d = X$ for any $X \subseteq S$ and hence, e.g., for $X = \{x_\gamma\}$, as already noted above we have $\nu_\gamma(X) = 1 - \gamma \not\leq 0 + \varepsilon = \nu_0(X_\varepsilon^d) + \varepsilon$ as $\varepsilon \leq \gamma < \frac{1}{2}$.

One can notice that when viewed as processes, ν_γ and ν_0 are the same as s and t of Example 3. Technically, it is rather $\tau_a(s)$ and $\tau_a(t)$ of course.

Once one has seen the Lévy-Prokhorov distance on distributions it seems not surprising that it has some link with the ε -distance. However, our surprise was the other way around, when we realized that the distance on distributions δ^d that we had to concoct to make d^* a fixed point of Δ in Equation (4) already existed.

3.2 The Lévy-Prokhorov distance on LMPs

We are now ready to define the functional on distances over LMPs that we were looking for. We expect a fixpoint of this functional to be a behavioral distance, but more importantly we expect d^* to be the greatest fixpoint of it.

► **Definition 9.** Let $\langle S, d \rangle$ be a pseudometric space, on which an LMP $\tau : S \rightarrow (\mathcal{D}S)^A$ is defined. Let s_0 and s_1 be states. We define the functional Δ_{LP} as

$$\begin{aligned} \Delta_{LP}(d)(s_0, s_1) &= \sup_{a \in A} \delta_{LP}^d(\tau_a(s_0), \tau_a(s_1)) \\ &= \sup_{a \in A} \inf\{\varepsilon \mid (\forall A \subseteq S : \tau_a(s_i)(A) \leq \tau_a(s_{1-i})(A_\varepsilon^d) + \varepsilon, \quad i = 0, 1)\}. \end{aligned}$$

Hence, for a fixed pseudometric space $\langle S, d \rangle$, Δ_{LP} applied to d gives a new pseudometric on S . Clearly, this can be seen as a functional on $\mathbf{PMet}_1$, mapping $\langle S, d \rangle$ to $(S, \Delta_{LP}(d))$.

Before proving that d^* is a fixpoint, we show how the supremum over actions can safely be replaced by inserting a universal quantification on actions inside the set, as will be useful in Proposition 15.

► **Lemma 10.** Let $\langle S, d \rangle$ be a pseudometric space, on which an LMP $\tau : S \rightarrow (\mathcal{D}S)^A$ is defined. Let s_0 and s_1 be states. Then

$$\Delta_{LP}(d)(s_0, s_1) = \inf\{\varepsilon \mid (\forall a \in A, \forall A \subseteq S : \tau_a(s_i)(A) \leq \tau_a(s_{1-i})(A_\varepsilon^d) + \varepsilon, \quad i = 0, 1)\}.$$

One remarkable property of prefixpoints of this functional is that they define ε -bisimulations. Of course the theorem is also true for fixpoints.

► **Theorem 11.** Let $\tau : S \rightarrow (\mathcal{D}S)^A$ be an LMP with $\langle S, d \rangle$ a pseudometric space. If d is a prefixpoint of Δ_{LP} , then for any $\varepsilon > 0$, $R_\varepsilon := \{(s, t) \in S \times S \mid d(s, t) < \varepsilon\}$, is an ε -bisimulation.

Proof. Let d be a prefixpoint, i.e., let $\Delta_{LP}(d) \leq d$. Let $R_\varepsilon = \{(s, t) \mid d(s, t) < \varepsilon\}$. We show it is an ε -bisimulation. Let $s_0 R_\varepsilon s_1$, i.e., $d(s_0, s_1) < \varepsilon$, and also $\Delta_{LP}(d)(s_0, s_1) < \varepsilon$, because d is a prefixpoint. Let $a \in A$; then $\delta_{LP}^d(\tau_a(s_0), \tau_a(s_1)) < \varepsilon$. Because δ_{LP}^d is an infimum, there is some $\gamma < \varepsilon$ such that for $i = 1, 2$ and $X \subseteq S$

$$\begin{aligned} \tau_a(s_i)(X) &\leq \tau_a(s_{1-i})(X_\gamma^d) + \gamma && \text{by definition of } \delta_{LP} \\ &\leq \tau_a(s_{1-i})(X_\varepsilon^d) + \varepsilon && \text{because } X_\gamma^d \subseteq X_\varepsilon^d \text{ and } \gamma \leq \varepsilon. \\ &= \tau_a(s_{1-i})(R_\varepsilon(X)) + \varepsilon && \text{by definition of } R_\varepsilon(X). \end{aligned}$$

So R_ε is an ε -bisimulation, as wanted. ◀

The proof of this theorem needs that X_ε^d be defined with a *strict* inequality. It is necessary for the existence of γ (and for the last equality).

This theorem is nice in itself but it also has an important corollary, that d^* is greater than or equal to all fixpoints of Δ_{LP} .

► **Corollary 12.** d^* is greater than or equal to all fixpoints, i.e.,

$$\Delta_{LP}(d) = d \implies d \sqsubseteq d^*.$$

Proof. Let d be a fixpoint of Δ_{LP} and let $d(s, t) = \varepsilon$. By the previous theorem, $s \sim_\gamma t$ for all $\gamma > \varepsilon$. Since $d^*(s, t) = \inf\{\varepsilon \mid s \sim_\varepsilon t\}$, we obtain $d^*(s, t) \leq \varepsilon = d(s, t)$, as wanted. ◀

With the next proposition, we provide another example of a fixpoint of Δ_{LP} : the exact bisimilarity distance, which we denote by $d_\sim$. It is equal to 0 if the states are bisimilar, otherwise it is 1.

► **Proposition 13.** The exact bisimilarity distance $d_\sim$ is a fixpoint of Δ_{LP} .

Proof. We first prove that bisimilar states s, t satisfy $\Delta_{LP}(d_\sim)(s, t) = 0$. Consider R to be bisimilarity. Because R is a bisimulation, we have that for all $s_0 R s_1$, $a \in \mathcal{A}$ and $X \subseteq S$,

$$\tau_a(s_0)(X) \leq \tau_a(s_1)(R(X)) = \tau_a(s_1)(X_\gamma^{d_\sim}),$$

for any $\gamma \in (0, 1)$, because $X_\gamma^{d_\sim}$ is just the smallest R -closed set that contains X . A symmetric argument applies with 0 and 1 interchanged. Consequently, $\Delta_{LP}(d_\sim)(s_0, s_1) = 0$, as wanted.

Conversely, if $\Delta_{LP}(d_\sim)(s_0, s_1) = 0$, then there is a sequence of $\gamma_n \in (0, 1)$ that converges to zero, for which we have

$$\tau_a(s_0)(X) \leq \tau_a(s_1)(X_{\gamma_n}^{d_\sim}) + \gamma_n = \tau_a(s_1)(R(X)) + \gamma_n.$$

Since the property is monotone, recall also (6), for all $\gamma \in (0, 1)$ we have

$$\tau_a(s_0)(X) \leq \tau_a(s_1)(X_\gamma^{d_\sim}) + \gamma = \tau_a(s_1)(R(X)) + \gamma.$$

So $\tau_a(s_0)(X) \leq \tau_a(s_1)(R(X))$, as wanted. ◀

► **Proposition 14.** Δ_{LP} is monotonic.

We can now prove that the distance defined by ε -bisimulation is a fixpoint of Δ_{LP} .

► **Proposition 15.** d^* is a fixpoint of Δ_{LP} , and hence it is the greatest fixpoint.

At this point we remark again that even if it seems not surprising that the Lévy-Prokhorov metric on distributions has some link with the ε -distance, going from the basic metric on distributions δ_{LP} to a meaningful distance on LMPs is not direct (one has to find the fixed-point properties).

3.3 The Lévy-Prokhorov lifting of the subdistribution functor to $\mathbf{PMet}_1$

Using the LP distance lifting, we define a functor $\mathcal{D}_{LP}$ on $\mathbf{PMet}_1$, the category of 1-bounded pseudometric spaces with nonexpansive functions, as follows: On objects $\langle X, d \rangle$, we have

$$\mathcal{D}_{LP}\langle X, d \rangle = (\mathcal{D}X, \delta_{LP}^d)$$

and on morphisms $f: \langle X_0, d_0 \rangle \rightarrow \langle X_1, d_1 \rangle$ we set $\mathcal{D}_{LP}f = \mathcal{D}f$, that is

$$\mathcal{D}_{LP}f: \mathcal{D}_{LP}\langle X_0, d_0 \rangle \rightarrow \mathcal{D}_{LP}\langle X_1, d_1 \rangle \text{ with } \varphi \mapsto \lambda y. \varphi(f^{-1}(\{y\})).$$

► **Proposition 16.** $\mathcal{D}_{LP}$ is a functor on $\mathbf{PMet}_1$.

Just like for the Kantorovich-lifting of the distribution functor, we can prove that the Lévy-Prokhorov lifting of the distribution functor is locally nonexpansive, in the next proposition. For this reason, note that given pseudometric spaces $\langle X, d_X \rangle$ and $\langle Y, d_Y \rangle$ in $\mathbf{PMet}_1$, the hom-set $X \rightarrow Y$ of all nonexpansive maps from X to Y carries a metric defined by

$$d_{X \rightarrow Y}(f_1, f_2) = \sup_{x \in X} d_Y(f_1(x), f_2(x)).$$

► **Proposition 17.** *The functor $\mathcal{D}_{LP}$ is locally nonexpansive, that is, for $f_1, f_2 \in X \rightarrow Y$*

$$\delta_{\mathcal{D}_{LP}X \rightarrow \mathcal{D}_{LP}Y}(\mathcal{D}_{LP}f_1, \mathcal{D}_{LP}f_2) \leq d_{X \rightarrow Y}(f_1, f_2)$$

with $\delta_{\mathcal{D}_{LP}X \rightarrow \mathcal{D}_{LP}Y}$ being the metric on the hom-set $\mathcal{D}_{LP}\langle X, d_X \rangle \rightarrow \mathcal{D}_{LP}\langle Y, d_Y \rangle$.

For the following property, recall that a map $f: X \rightarrow Y$ for (pseudo)metric spaces $\langle X, d_X \rangle$ and $\langle Y, d_Y \rangle$ is an isometry if and only if for all $x, y \in X$, $d_Y(f(x), f(y)) = d_X(x, y)$.

► **Proposition 18.** *The functor $\mathcal{D}_{LP}$ preserves isometries, i.e., if $f: X \rightarrow Y$ is an isometry for (pseudo)metric spaces $\langle X, d_X \rangle$ and $\langle Y, d_Y \rangle$, then so is $\mathcal{D}_{LP}f$.*

As a consequence, from the results of [31], we obtain that the lifting functor $\mathcal{D}_{LP}$ is *Kantorovich*, i.e., it has *codensity lifting* which means that there is a set of predicate liftings of type $\mathcal{D}[0, 1] \rightarrow [0, 1]$ such that codensity lifting gives rise to $\mathcal{D}_{LP}$. This has already been remarked in [31, Example 5.5.2]. However, the exact set of predicate liftings needed to give rise to the LP-lifting is not provided.

3.4 The lifted functor $\mathcal{D}_{LP}$ is not a monad lifting of $\mathcal{D}$

Behavioural distances have been axiomatized within the line of work on quantitative equational theories [39]. It was therefore a natural question for us whether the ε -distance, or the Lévy-Prokhorov distance itself can be given a quantitative axiomatization. This is what we briefly investigate in this section.

Mio et al. [43, Lemma 7.2, Theorem 7.7(2)] have shown that:

- If a monad on metric spaces is axiomatizable, then (just by being a monad on metric spaces) it has nonexpansive unit and multiplication.
 - A monad on metric spaces that is a lifting of a monad on **Sets** is axiomatizable with a quantitative theory. Being a lifting means that it acts on objects and arrows in metric spaces in the same way it does in **Sets** and the unit and multiplication are nonexpansive.
- The Kantorovich lifting of $\mathcal{D}$ has been axiomatized, and the total variation distance has been shown non-axiomatizable (as the unit is not nonexpansive). In the case of $\mathcal{D}_{LP}$ we can see that the unit is nonexpansive, but the multiplication is not, and hence $\mathcal{D}_{LP}$ is not a lifting of the monad $\mathcal{D}$ to $\mathbf{PMet}_1$ and the Lévy-Prokhorov distance on distributions is not axiomatizable with a quantitative theory, at least not with the standard multiplication.


Recall the definitions of η and μ for the subdistribution monad $\mathcal{D}$:

$$\eta_X(x) = 1_x, \text{ the Dirac distribution at } x; \text{ and } \mu_X(\Phi)(x) = \sum_{\varphi} \Phi(\varphi) \cdot \varphi(x).$$

► **Lemma 19.** *The unit η of $\mathcal{D}$ is nonexpansive with respect to the Lévy-Prokhorov distance. Moreover, it is an isometry, i.e., $\delta_{LP}^d(1_x, 1_y) = d(x, y)$.*

The unit of the distribution monad is also an isometry with respect to the Kantorovich distance (without discount), as is easy to prove using the duality with the Wasserstein distance. However, the multiplication μ of the monad $\mathcal{D}$ is not nonexpansive, as the following example shows.

► **Example 20.** Consider a , b and c in the figure below as distributions over $\{\perp, \bullet\}$.



The ambient distance is $d(\perp, \bullet) = 1$, we omit it from the notation: we write δ_{LP} for δ_{LP}^d . Consider the distributions $\varphi := 1_a$ and $\psi := (1 - \varepsilon)1_b + \varepsilon 1_c$ over these distributions, as pictured. We collect several useful facts:

- $\delta_{LP}(a, b) = \varepsilon$ and $\delta_{LP}^{\delta_{LP}}(\varphi, \psi) = \varepsilon$.
- $\mu\varphi = a$ and $\mu\psi = (1 - \varepsilon)^2 \perp + (1 - \varepsilon)\varepsilon \bullet + \varepsilon \bullet = (1 - \varepsilon)^2 \perp + (2 - \varepsilon)\varepsilon \bullet$.
- $\{\perp\}_\gamma = \{\perp\}$ for any $\gamma \leq \varepsilon$ and hence $\mu\varphi(\{\perp\}) = 1$ and $\mu\psi(\{\perp\}_\gamma) = \mu\psi(\{\perp\}) = (1 - \varepsilon)^2$.

For $\gamma < \varepsilon(2 - \varepsilon)$, we have $1 \not\leq (1 - \varepsilon)^2 + \gamma$ which therefore yields

$$\mu\varphi(\{\perp\}) = 1 \not\leq (1 - \varepsilon)^2 + \gamma = \mu\psi(\{\perp\}_\gamma) + \gamma.$$

This shows that $\delta_{LP}(\mu\varphi, \mu\psi) \geq \varepsilon(2 - \varepsilon) > \varepsilon = \delta_{LP}^{\delta_{LP}}(\varphi, \psi)$ and hence μ is not nonexpansive.

This result is not surprising, as the binding operator μ really is a multiplication, and so it accords well with the Kantorovich metric, which does multiply the probabilities of distributions in play. We do not know yet if another operator could help express the functor $\mathcal{D}_{LP}$ as a monad lifting.

4 ε -(Bi)simulations, coalgebraically

The Kantorovich behavioural distance has a coalgebraic characterization via a final coalgebra semantics. Whether the ε -distance is obtained by finality is still open, but the results on the Kantorovich distance do not seem to directly apply here. However, as the distance is defined via ε -bisimulations and bisimilarity, it is natural to see whether a coalgebraic semantics arising from generalizing coalgebraic bisimulations and bisimilarity is in place. The answer to this question is positive and we present the necessary observations in this section. An abstract (in this case coalgebraic) characterization allows for, on the theoretical side, deeper and clearer understanding of the notion under study, and, on the practical side, generalizations. We could use this generality to define notions of ε -bisimulations for other types of (probabilistic) systems.

In a nutshell, in this section we show that ε -simulation and ε -bisimulation have a span-diagram characterization, in a way similar to coalgebraic simulation [34, 32, 33] and Aczel-Mendler coalgebraic bisimulation [3, 35, 47]. The development is using a notion of ε -coupling defined in [54]. For simplicity, we ignore the labels in this section, i.e., we assume there is a single label.

To start with, we recall the basic notions related to coalgebraic (bi)simulation, formulated for the functor $\mathcal{D}$ that we are mainly interested in in this work. Recall the definition of the sub-distribution functor $\mathcal{D}$ on **Sets**: It maps a set X to $\mathcal{D}X$, the set of discrete subdistributions on X and on arrows $f: X \rightarrow Y$, $\mathcal{D}f: \mathcal{D}X \rightarrow \mathcal{D}Y$ is $\mathcal{D}f(\varphi)(y) = \varphi(f^{-1}(\{y\}))$.

Coalgebras of the functor $\mathcal{D}$ are Markov chains, formally they are pairs (X, c) of a carrier set X and transition map $c: X \rightarrow \mathcal{D}X$. We often just refer to the coalgebra by the transition map c . Given two such coalgebras, $c: X \rightarrow \mathcal{D}X$ and $d: Y \rightarrow \mathcal{D}Y$ a coalgebra homomorphism from (X, c) to (Y, d) is a map $h: X \rightarrow Y$ making the left diagram below commute.

$$\begin{array}{ccc}
X & \xrightarrow{h} & Y \\
c \downarrow & & \downarrow d \\
\mathcal{D}X & \xrightarrow{\mathcal{D}h} & \mathcal{D}Y
\end{array}
\qquad
\begin{array}{ccccc}
X & \xleftarrow{\pi_1} & R & \xrightarrow{\pi_2} & Y \\
c \downarrow & & b \downarrow & & \downarrow d \\
\mathcal{D}X & \xleftarrow{\mathcal{D}\pi_1} & \mathcal{D}R & \xrightarrow{\mathcal{D}\pi_2} & \mathcal{D}Y
\end{array}$$

► **Definition 21.** Given two coalgebras $c: X \rightarrow \mathcal{D}X$ and $d: Y \rightarrow \mathcal{D}Y$, a coalgebraic bisimulation is a relation $R \subseteq X \times Y$ such that there exists a coalgebra structure $b: R \rightarrow \mathcal{D}R$ making the two projections $\pi_1: R \rightarrow X$ and $\pi_2: R \rightarrow Y$ coalgebra homomorphisms, i.e., making the right span diagram above commute.

It is well known ([20, 49]) that this corresponds to the standard notion of probabilistic bisimulation ([38]). For the case of continuous distributions, taking a cospan instead of a span avoids the need for analytic spaces [17]. Back to the discrete case, one way to define the coalgebra structure b is using the notion of a coupling: If for all $(x, y) \in R$, there is a coupling β for $\mu = c(x)$ and $\nu = d(y)$, then setting $b(x, y) = \beta$ provides the needed transition structure. Recall that $\beta \in \mathcal{D}R$ is a *coupling* of $\mu \in \mathcal{D}X$ and $\nu \in \mathcal{D}Y$ if its marginals are μ and ν , respectively:

$$\sum_{y \in Y} \beta(x, y) = \mu(x), \quad \sum_{x \in X} \beta(x, y) = \nu(y).$$

Note that this definition of a coupling, fitting the definition of coalgebraic bisimulation for the subprobability distribution functor, ensures that a coupling of two subdistributions exists only if they have the same total mass. One could also define couplings by adding a dummy element and hence viewing a subdistribution as a full distribution, which provides a more general notion. However, these details are not relevant for what we are really interested in, which are ε -bisimulations.

We will use a relaxed notion of ε -coupling [54] to define ε -(bi)simulation in analogy with the coalgebraic definition above. Before we recall those, let us first mention coalgebraic simulations, due to [34, 32, 33] with a coalgebraic formulation for LMPs already in [21]. For this, note that an ordered functor F is a functor with an order on each object FX . The original definition in [34] uses preorders as minimal requirement, and also implies that for any map $f: X \rightarrow Y$, Ff preserves the order, i.e., is monotone – a condition that is not used in our results. The subdistribution functor $\mathcal{D}$ is ordered, e.g., by pointwise order. This order becomes trivial in case of the distribution functor which is the reason why we work with subdistributions here. Coalgebraic simulation can be defined for ordered functors, here we recall the definition in the special case of the subdistribution functor $\mathcal{D}$. All notions involve *lax* and *oplax* morphism.

► **Definition 22.** A lax homomorphism from (X, c) to (Y, d) is a morphism $l: X \rightarrow Y$ with the property that $d \circ l \sqsubseteq \mathcal{D}l \circ c$, i.e., it makes the left lax diagram below commute. An oplax homomorphism from (X, c) to (Y, d) is a lax homomorphism for the dual/opposite order, that is a morphism $o: X \rightarrow Y$ that makes the middle diagram below commute:

$$\begin{array}{ccc}
X & \xrightarrow{l} & Y \\
c \downarrow & \sqsubseteq & \downarrow d \\
\mathcal{D}X & \xrightarrow{\mathcal{D}l} & \mathcal{D}Y
\end{array}
\qquad
\begin{array}{ccc}
X & \xrightarrow{o} & Y \\
c \downarrow & \sqsupseteq & \downarrow d \\
\mathcal{D}X & \xrightarrow{\mathcal{D}o} & \mathcal{D}Y
\end{array}
\qquad
\begin{array}{ccccc}
X & \xleftarrow{\pi_1} & R & \xrightarrow{\pi_2} & Y \\
c \downarrow & \sqsubseteq & b \downarrow & \sqsubseteq & \downarrow d \\
\mathcal{D}X & \xleftarrow{\mathcal{D}\pi_1} & \mathcal{D}R & \xrightarrow{\mathcal{D}\pi_2} & \mathcal{D}Y
\end{array}$$

The order on distributions is defined pointwise: For a set A and $\varphi, \psi \in \mathcal{D}A$, we have $\varphi \sqsubseteq \psi$ iff for all $a \in A$, $\varphi(a) \leq \psi(a)$.

Note that, equivalently, $\varphi \sqsubseteq \psi$ iff for all $B \subseteq A$, $\varphi(B) \leq \psi(B)$, where the right-to-left implication follows from instantiating on singleton subsets and the left-to-right implication follows from the additivity of distributions.

Initially, generic coalgebraic simulations have been studied for ordered functors, by Hughes and Jacobs [34], using a notion of lax relation lifting. On the other hand, Hasuo [32, 33] discovered (generalizations of) forward, backward, and hybrid simulations in Kleisli categories and moreover showed in [33] that Hughes-Jacobs simulations are a special case. While Kleisli categories are appealing for soundness results, and monads suitable for traces come equipped with an order, we may safely ignore all Kleisli aspects here. We now recall the following notions of simulations [32], formulated for general categories of coalgebras (as long as the functor comes equipped with an order) where they amount to different names for lax/oplax morphisms, for the sake of making the connection to [32] explicit.

► **Definition 23.** A forward simulation from (X, c) to (Y, d) is a lax homomorphism from (Y, d) to (X, c) . A backward simulation from (X, c) to (Y, d) is an oplax homomorphism. Combining the lax-commuting boxes leads to hybrid, forward-backward and backward-forward, simulations. In particular, Hughes-Jacobs simulation is a forward-backward simulation: It is a relation R on which there exists a coalgebra structure $b: R \rightarrow \mathcal{D}R$ making π_1 a forward simulation from (X, c) to (R, b) and π_2 a backward simulation from (R, b) to (Y, d) , depicted in the right diagram above.

Clearly, a symmetric simulation on a single system is a bisimulation.

In the rest of this section, we will show that ε -(bi)simulations can be depicted similarly. For this we will need bounded-lax-commutativity of morphisms as well as a notion of ε -coupling (which then directly gives an ε -relation-lifting). We will focus here on the functor $\mathcal{D}$ only. These notions can be generalized to generic coalgebras for functors with suitable structure. However, those observations are beyond the scope of this paper and we leave them for future work.

► **Definition 24.** An ε -lax homomorphism from (X, c) to (Y, d) is a morphism $l: X \rightarrow Y$ that makes the left ε -lax diagram below commute, i.e., $d \circ l \sqsubseteq_\varepsilon \mathcal{D}l \circ c$. An ε -oplax homomorphism from (X, c) to (Y, d) is a morphism $o: X \rightarrow Y$ with the property that $\mathcal{D}o \circ c \sqsubseteq_\varepsilon d \circ o$, i.e., it makes the right ε -lax diagram below commute.

$$\begin{array}{ccc} X & \xrightarrow{l} & Y \\ c \downarrow & \varepsilon \sqsubseteq & \downarrow d \\ \mathcal{D}X & \xrightarrow{\mathcal{D}l} & \mathcal{D}Y \end{array} \qquad \begin{array}{ccc} X & \xrightarrow{o} & Y \\ c \downarrow & \sqsubseteq_\varepsilon & \downarrow d \\ \mathcal{D}X & \xrightarrow{\mathcal{D}o} & \mathcal{D}Y \end{array}$$

The ε -order on distributions is defined by: For a set A and $\varphi, \psi \in \mathcal{D}A$, we have $\varphi \sqsubseteq_\varepsilon \psi$ iff for all $B \subseteq A$, $\varphi(B) \leq \psi(B) + \varepsilon$.

► **Remark 25.** Note that $\sqsubseteq_\varepsilon$ is not an order relation, namely it is not transitive. Note also that the definition of $\sqsubseteq_\varepsilon$ can not be expressed on elements (of A) only, as it is stronger than the property: for all $a \in A$, $\varphi(a) \leq \psi(a) + \varepsilon$. Namely, the role of ε here, just like in the definition of ε -(bi)simulation is *global*.

We next recall the notion of ε -coupling from [54], where it is called an ε -weight function.

► **Definition 26.** Let $R \subseteq X \times Y$. A distribution $\beta \in \mathcal{D}R$ is an ε -coupling for $\mu \in \mathcal{D}X$ and $\nu \in \mathcal{D}Y$ iff the following three conditions hold:

1. $\sum_{y \in Y} \beta(x, y) \leq \mu(x)$, for all $x \in X$
2. $\sum_{x \in X} \beta(x, y) \leq \nu(y)$, for all $y \in Y$
3. $\mu(X) \leq \sum_{x \in X, y \in Y} \beta(x, y) + \varepsilon = \sum_{(x, y) \in R} \beta(x, y) + \varepsilon$.

To be fully precise, we should consider $\beta \in \mathcal{D}(X \times Y)$ with support contained in R . We write $\beta \in \mathcal{D}R$ and yet sometimes sum over all $x \in X, y \in Y$, i.e., we identify $\beta \in \mathcal{D}R$ with a distribution in $\mathcal{D}(X \times Y)$ that assigns probability 0 to all pairs out of R and acts as β on R .

Before we proceed with the main observation of this section, we prove an auxiliary property that allows for rewriting the definition of ε -coupling.

► **Lemma 27.** *Assume that Condition 1. in Definition 26 holds for $\beta \in \mathcal{D}R$ with $R \subseteq X \times Y$ and $\mu \in \mathcal{D}X$. Then $\mu(X) \leq \sum_{x \in X, y \in Y} \beta(x, y) + \varepsilon$ (i.e., Condition 3. in Definition 26) is equivalent to the condition*

$$\mu(S) \leq \sum_{x \in S, y \in Y} \beta(x, y) + \varepsilon, \quad \text{for all } S \subseteq X.$$

This property again emphasises the global nature of ε in our situation. As a consequence, for an ε -coupling $\beta \in \mathcal{D}R$ of $\mu \in \mathcal{D}X$ and $\nu \in \mathcal{D}Y$, with $R \subseteq X \times Y$, for each $S \subseteq X$:

$$\sum_{x \in S, y \in Y} \beta(x, y) \leq \mu(S) \leq \sum_{x \in S, y \in Y} \beta(x, y) + \varepsilon.$$

► **Proposition 28.** *Let (X, c) and (Y, d) be two $\mathcal{D}$ -coalgebras. The following three properties are equivalent for $R \subseteq X \times Y$:*

1. *R is an ε -simulation.*
2. *For every $(x, y) \in R$, there is an ε -coupling $\beta \in \mathcal{D}R$ of $\mu = c(x)$ and $\nu = d(y)$.*
3. *The “ ε -lax-bounded” span of morphisms on the left below commutes.*

$$\begin{array}{ccc} X & \xleftarrow{\pi_1} R & \xrightarrow{\pi_2} Y \\ c \downarrow & \sqsupseteq \sqsubseteq_\varepsilon & b \downarrow \sqsubseteq \downarrow d \\ \mathcal{D}X & \xleftarrow{\mathcal{D}\pi_1} \mathcal{D}R & \xrightarrow{\mathcal{D}\pi_2} \mathcal{D}Y \end{array} \qquad \begin{array}{ccc} X & \xleftarrow{\pi_1} R & \xrightarrow{\pi_2} X \\ c \downarrow & \sqsupseteq \sqsubseteq_\varepsilon & b \downarrow \sqsubseteq \downarrow d \\ \mathcal{D}X & \xleftarrow{\mathcal{D}\pi_1} \mathcal{D}R & \xrightarrow{\mathcal{D}\pi_2} \mathcal{D}X \end{array}$$

Note that Condition 2. in Proposition 28 can be taken as a definition of ε -relation-lifting, in analogy to relation lifting for $\mathcal{D}$ being defined using the existence of a coupling for any pair of elements in R . This way was also taken in [34] with the definition of lax relation lifting.

As a consequence of Proposition 28, we immediately get that a relation $R \subseteq X \times X$ on the states of a $\mathcal{D}$ -coalgebras is an ε -bisimulation iff the right diagram above commutes. This diagram resembles the defining diagram of the recently introduced uncertain bisimulations [46], but requires a double property (for $\sqsubseteq_\varepsilon$ as well as for $\sqsubseteq^{op}$). A detailed comparison with this interesting work remains a future task.

As already mentioned, generalizing the notions of approximate (bi)simulations to generic coalgebras is an interesting direction for future work. A similar notion is being developed for cost automata [4].

5 Concluding remarks

We have shown that ε -bisimulations are closely connected with the Lévy-Prokhorov metric on (sub)probability distributions: The LP pseudometric yields a lifting of $\mathcal{D}$ to a functor on pseudometric spaces and induces coinductive behaviour distance on LMPs, which is the greatest fixpoint of a suitable functional and turns out to be exactly the distance induced by ε -bisimilarity. Remarkably, any fixpoint distance of that functional defines an ε -bisimulation. This is the first time that a distance on distributions other than the Kantorovich distance is used for characterizing a known behavioural distance.

References

- 1 Alessandro Abate. Approximation metrics based on probabilistic bisimulations for general state-space markov processes: A survey. *Electronic Notes in Theoretical Computer Science*, 297:3–25, 2013. Proceedings of the first workshop on Hybrid Autonomous Systems. doi:10.1016/j.entcs.2013.12.002.
- 2 Alessandro Abate, Marta Z. Kwiatkowska, Gethin Norman, and David Parker. Probabilistic model checking of labelled markov processes via finite approximate bisimulations. In Franck van Breugel, Elham Kashefi, Catuscia Palamidessi, and Jan Rutten, editors, *Horizons of the Mind. A Tribute to Prakash Panangaden - Essays Dedicated to Prakash Panangaden on the Occasion of His 60th Birthday*, volume 8464 of *Lecture Notes in Computer Science*, pages 40–58. Springer, 2014. doi:10.1007/978-3-319-06880-0_2.
- 3 Peter Aczel and Nax Mendler. A final coalgebra theorem. In *Proc. 3rd CTCS*, volume 389 of *LNCS*, pages 357–365. Springer, 1989. doi:10.1007/BFB0018361.
- 4 Pedro Azevedo de Amorim, Mayuko Kori, and Koko Muroya. A framework for coalgebraic reward-sensitive bisimulation, 2025.
- 5 Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, and Radu Mardare. On-the-fly computation of bisimilarity distances. *Log. Methods Comput. Sci.*, 13(2), 2017. doi:10.23638/LMCS-13(2:13)2017.
- 6 Giovanni Bacci, Giorgio Bacci, Kim G. Larsen, and Radu Mardare. On the metric-based approximate minimization of markov chains. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, volume 80 of *LIPIcs*, pages 104:1–104:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.ICALP.2017.104.
- 7 C. Baier. *On Algorithmic Verification Methods for Probabilistic Systems*. Habilitation Thesis, 1998.
- 8 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- 9 Paolo Baldan, Filippo Bonchi, Henning Kerstan, and Barbara König. Coalgebraic behavioral metrics. *Log. Methods Comput. Sci.*, 14(3), 2018. doi:10.23638/LMCS-14(3:20)2018.
- 10 Gaoang Bian and Alessandro Abate. On the relationship between bisimulation and trace equivalence in an approximate probabilistic context. In Javier Esparza and Andrzej S. Murawski, editors, *Foundations of Software Science and Computation Structures - 20th International Conference, FOSSACS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, volume 10203 of *Lecture Notes in Computer Science*, pages 321–337, 2017. doi:10.1007/978-3-662-54458-7_19.
- 11 Richard Blute, Josée Desharnais, Abbas Edalat, and Panangaden Panangaden. Bisimulation for labelled Markov processes. In *Proceedings of the Twelfth IEEE Symposium On Logic In Computer Science (LICS)*, Warsaw, Poland, 1997.
- 12 Filippo Bonchi, Barbara König, and Daniela Petrisan. Up-to techniques for behavioural metrics via fibrations. *Math. Struct. Comput. Sci.*, 33(4-5):182–221, 2023. doi:10.1017/S0960129523000166.
- 13 Filippo Bonchi, Alexandra Silva, and Ana Sokolova. The Power of Convex Algebras. In *CONCUR 2017*, volume 85, pages 23:1–23:18. LIPIcs, 2017. doi:10.4230/LIPIcs.CONCUR.2017.23.
- 14 Filippo Bonchi, Ana Sokolova, and Valeria Vignudelli. The theory of traces for systems with nondeterminism, probability, and termination. *Log. Methods Comput. Sci.*, 18(2), 2022. doi:10.46298/LMCS-18(2:21)2022.
- 15 Nicolas Bourbaki. *Elements of Mathematics*. Springer-Verlag, 1995. Original French edition published by MASSON, Paris in 1971.

- 16 Keri D'Angelo, Sebastian Gurke, Johanna Maria Kirss, Barbara König, Matina Najafi, Wojciech Rozowski, and Paul Wild. Behavioural metrics: Compositionality of the kantorovich lifting and an application to up-to techniques. In Rupak Majumdar and Alexandra Silva, editors, *35th International Conference on Concurrency Theory, CONCUR 2024, September 9-13, 2024, Calgary, Canada*, volume 311 of *LIPIcs*, pages 20:1–20:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.CONCUR.2024.20.
- 17 Vincent Danos, Josée Desharnais, François Laviolette, and Prakash Panangaden. Bisimulation and cocongruence for probabilistic systems. *Information and Computation*, 2005. Special issue for selected papers from CMCS04. 22 pages.
- 18 Luca de Alfaro, Marco Faella, and Mariëlle Stoelinga. Linear and branching metrics for quantitative transition systems. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings*, volume 3142 of *Lecture Notes in Computer Science*, pages 97–109. Springer, 2004. doi:10.1007/978-3-540-27836-8_11.
- 19 Luca de Alfaro and Rupak Majumdar. Quantitative solution of omega-regular games. In *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 675–683, New York, NY, USA, 2001. ACM. doi:10.1145/380752.380871.
- 20 Erik P. de Vink and Jan J. M. M. Rutten. Bisimulation for probabilistic transition systems: A coalgebraic approach. *Theor. Comput. Sci.*, 221(1-2):271–293, 1999. doi:10.1016/S0304-3975(99)00035-3.
- 21 Josée Desharnais. *Labelled Markov Processes*. PhD thesis, McGill University, 2000.
- 22 Josée Desharnais, Vineet Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labeled Markov processes. In Jos C. M. Baeten and S. Mauw, editors, *Proceedings of 10th International Conference on Concurrency Theory, Eindhoven, The Netherlands*, Lecture Notes in Computer Science, pages 258–273. Springer-Verlag, August 1999.
- 23 Josée Desharnais, Vineet Gupta, R. Jagadeesan, and P. Panangaden. Approximating continuous Markov processes. In *Proceedings of the 15th Annual IEEE Symposium On Logic In Computer Science, Santa Barbara, California, USA*, 2000. pp. 95-106.
- 24 Josée Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Metrics for labelled Markov processes. *Theoretical Computer Science*, 318(3):323–354, 2004. doi:10.1016/j.tcs.2003.09.013.
- 25 Josée Desharnais, R. Jagadeesan, Vineet Gupta, and P. Panangaden. The metric analogue of weak bisimulation for probabilistic processes. In *Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science (LICS02)*, pages 413–422, Copenhagen, Denmark, July 2002. IEEE Computer Society. doi:10.1109/LICS.2002.1029849.
- 26 Josée Desharnais, François Laviolette, and Mathieu Tracol. Approximate analysis of probabilistic processes: Logic, simulation and games. In *Fifth International Conference on the Quantitative Evaluation of Systems (QEST 2008), 14-17 September 2008, Saint-Malo, France*, pages 264–273. IEEE Computer Society, 2008. doi:10.1109/QEST.2008.42.
- 27 Alessandro D’Innocenzo, Alessandro Abate, and Joost-Pieter Katoen. Robust PCTL model checking. In Thao Dang and Ian M. Mitchell, editors, *Hybrid Systems: Computation and Control (part of CPS Week 2012), HSCC’12, Beijing, China, April 17-19, 2012*, pages 275–286. ACM, 2012. doi:10.1145/2185632.2185673.
- 28 A. Giacalone, P. Misra, and S. Prasad. Facile: A symmetric integration of concurrent and functional programming. In *LNCS 352: TAPSOFT 89*, 1989.
- 29 Alessandro Giacalone, Chi chang Jou, and Scott A. Smolka. Algebraic reasoning for probabilistic concurrent systems. In *Proc. IFIP TC2 Working Conference on Programming Concepts and Methods*, pages 443–458. North-Holland, 1990.
- 30 Antoine Girard and George J. Pappas. Approximate bisimulation relations for constrained linear systems. *Automatica*, 43(8):1307–1317, 2007. doi:10.1016/j.automatica.2007.01.019.

- 31 Sergey Goncharov, Dirk Hofmann, Pedro Nora, Lutz Schröder, and Paul Wild. Kantorovich functors and characteristic logics for behavioural distances. In Orna Kupferman and Pawel Sobocinski, editors, *Foundations of Software Science and Computation Structures - 26th International Conference, FoSSaCS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22-27, 2023, Proceedings*, volume 13992 of *Lecture Notes in Computer Science*, pages 46–67. Springer, 2023. doi:10.1007/978-3-031-30829-1_3.
- 32 Ichiro Hasuo. Generic forward and backward simulations. In *CONCUR 2006*, pages 406–420. LNCS 4137, 2006. doi:10.1007/11817949_27.
- 33 Ichiro Hasuo. Generic forward and backward simulations II: probabilistic simulation. In *CONCUR 2010*, pages 447–461. LNCS 6269, 2010. doi:10.1007/978-3-642-15375-4_31.
- 34 Jesse Hughes and Bart Jacobs. Simulations in coalgebra. *Theor. Comput. Sci.*, 327(1-2):71–108, 2004. doi:10.1016/J.TCS.2004.07.022.
- 35 Bart Jacobs. *Introduction to Coalgebra: Towards Mathematics of States and Observation*, volume 59 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2016. doi:10.1017/CB09781316823187.
- 36 Bart Jacobs, Alexandra Silva, and Ana Sokolova. Trace semantics via determinization. *J. Comput. Syst. Sci.*, 81(5):859–879, 2015. doi:10.1016/J.JCSS.2014.12.005.
- 37 Dexter Kozen. Semantics of probabilistic programs. *Journal of Computer and Systems Sciences*, 22:328–350, 1981. doi:10.1016/0022-0000(81)90036-2.
- 38 Kim G. Larsen and Arne Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991. doi:10.1016/0890-5401(91)90030-6.
- 39 Radu Mardare, Prakash Panangaden, and Gordon D. Plotkin. Quantitative algebraic reasoning. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 700–709. ACM, 2016. doi:10.1145/2933575.2934518.
- 40 Radu Mardare, Prakash Panangaden, and Gordon D. Plotkin. On the axiomatizability of quantitative algebras. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*, pages 1–12. IEEE Computer Society, 2017. doi:10.1109/LICS.2017.8005102.
- 41 R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
- 42 Matteo Mio, Ralph Sarkis, and Valeria Vignudelli. Combining nondeterminism, probability, and termination: Equational and metric reasoning. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–14. IEEE, 2021. doi:10.1109/LICS52264.2021.9470717.
- 43 Matteo Mio, Ralph Sarkis, and Valeria Vignudelli. Universal quantitative algebra for fuzzy relations and generalised metric spaces. *Log. Methods Comput. Sci.*, 20(4), 2024. doi:10.46298/LMCS-20(4:19)2024.
- 44 Prakash Panangaden. *Labelled Markov Processes*. Imperial College Press, 2009.
- 45 Yu. V. Prokhorov. Convergence of random processes and limit theorems in probability theory. *Theory of Probability & Its Applications*, 1(2):157–214, 1956. doi:10.1137/1101016.
- 46 Jurriaan Rot and Thorsten Wißmann. Bisimilar states in uncertain structures. In Paolo Baldan and Valeria de Paiva, editors, *10th Conference on Algebra and Coalgebra in Computer Science, CALCO 2023, Indiana University Bloomington, IN, USA, June 19-21, 2023*, volume 270 of *LIPIcs*, pages 12:1–12:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.CALCO.2023.12.
- 47 J.J.M.M. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 2000. doi:10.1016/S0304-3975(00)00056-6.
- 48 R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. In *Proc. Concur'94*, pages 481–496. LNCS 836, 1994.
- 49 A. Sokolova. *Coalgebraic Analysis of Probabilistic Systems*. PhD thesis, TU Eindhoven, 2005.

- 50 Timm Spork, Christel Baier, Joost-Pieter Katoen, Jakob Piribauer, and Tim Quatmann. A spectrum of approximate probabilistic bisimulations. In Rupak Majumdar and Alexandra Silva, editors, *35th International Conference on Concurrency Theory, CONCUR 2024, September 9-13, 2024, Calgary, Canada*, volume 311 of *LIPIcs*, pages 37:1–37:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, August 2024. doi:10.4230/LIPIcs.CONCUR.2024.37.
- 51 David Sprunger, Shin-ya Katsumata, Jérémy Dubut, and Ichiro Hasuo. Fibrational bisimulations and quantitative reasoning. In Corina Cirstea, editor, *Coalgebraic Methods in Computer Science*, pages 190–213, Cham, 2018. Springer International Publishing. doi:10.1007/978-3-030-00389-0_11.
- 52 David Sprunger, Shin-ya Katsumata, Jérémy Dubut, and Ichiro Hasuo. Fibrational bisimulations and quantitative reasoning: Extended version. *J. Log. Comput.*, 31(6):1526–1559, 2021. doi:10.1093/LOGCOM/EXAB051.
- 53 Qiyi Tang and Franck van Breugel. Algorithms to compute probabilistic bisimilarity distances for labelled markov chains. In *Proc. CONCUR 2017*, volume 85 of *LIPIcs*, pages 27:1–27:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.CONCUR.2017.27.
- 54 Mathieu Tracol, Josée Desharnais, and Abir Zhioua. Computing distances between probabilistic automata. In Mieke Massink and Gethin Norman, editors, *Proceedings Ninth Workshop on Quantitative Aspects of Programming Languages, QAPL 2011, Saarbrücken, Germany, April 1-3, 2011*, volume 57 of *EPTCS*, pages 148–162, 2011. doi:10.4204/EPTCS.57.11.
- 55 Franck van Breugel. Probabilistic bisimilarity distances. *ACM SIGLOG News*, 4(4):33–51, November 2017. doi:10.1145/3157831.3157837.
- 56 Franck van Breugel, Claudio Hermida, Michael Makkai, and James Worrell. Recursively defined metric spaces without contraction. *Theor. Comput. Sci.*, 380(1-2):143–163, 2007. doi:10.1016/j.tcs.2007.02.059.
- 57 Franck van Breugel, Babita Sharma, and James Worrell. Approximating a behavioural pseudometric without discount for probabilistic systems. In Helmut Seidl, editor, *FoSSaCS*, volume 4423 of *Lecture Notes in Computer Science*, pages 123–137. Springer, 2007. doi:10.1007/978-3-540-71389-0_10.
- 58 Franck van Breugel and James Worrell. Towards quantitative verification of probabilistic transition systems. In *ICALP '01: Proceedings of the 28th International Colloquium on Automata, Languages and Programming*, pages 421–432, London, UK, 2001. Springer-Verlag. doi:10.1007/3-540-48224-5_35.
- 59 Franck van Breugel and James Worrell. A behavioural pseudometric for probabilistic transition systems. *Theor. Comput. Sci.*, 331(1):115–142, February 2005. doi:10.1016/j.tcs.2004.09.035.
- 60 Rob J. van Glabbeek, Scott A. Smolka, and Bernhard Steffen. Reactive, generative and stratified models of probabilistic processes. *Inf. Comput.*, 121(1):59–80, 1995. doi:10.1006/inco.1995.1123.
- 61 Mingsheng Ying and Martin Wirsing. Approximate bisimilarity. In *AMAST '00: Proceedings of the 8th International Conference on Algebraic Methodology and Software Technology*, pages 309–322, London, UK, 2000. Springer-Verlag. doi:10.1007/3-540-45499-3_23.

A Proofs

► **Lemma 10.** *Let $\langle S, d \rangle$ be a pseudometric space, on which an LMP $\tau : S \rightarrow (\mathcal{DS})^A$ is defined. Let s_0 and s_1 be states. Then*

$$\Delta_{LP}(d)(s_0, s_1) = \inf\{\varepsilon \mid (\forall a \in \mathcal{A}, \forall A \subseteq S : \tau_a(s_i)(A) \leq \tau_a(s_{1-i})(A_\varepsilon^d) + \varepsilon, \quad i = 0, 1)\}.$$

Proof. For $a \in \mathcal{A}$ and $\varepsilon > 0$, we define the following predicate:

$$P(a, \varepsilon) := (\forall A \subseteq S : \tau_a(s_i)(A) \leq \tau_a(s_{1-i})(A_\varepsilon^d) + \varepsilon, \quad i = 0, 1).$$

This family of predicates satisfies

$$\text{if } \gamma > \varepsilon, P(a, \varepsilon) \Rightarrow P(a, \gamma). \quad (6)$$

We now prove that

$$\sup_{a \in \mathcal{A}} \inf\{\varepsilon \mid P(a, \varepsilon)\} = \inf\{\varepsilon \mid (\forall a \in \mathcal{A} : P(a, \varepsilon))\}$$

$\leq$: The following sequence of arguments shows this inequality.

$$\begin{aligned} \{\varepsilon \mid P(a_0, \varepsilon)\} &\supseteq \{\varepsilon \mid (\forall a \in \mathcal{A} : P(a, \varepsilon))\} && \text{for all } a_0 \in \mathcal{A}, \text{ so} \\ \inf\{\varepsilon \mid P(a_0, \varepsilon)\} &\leq \inf\{\varepsilon \mid (\forall a \in \mathcal{A} : P(a, \varepsilon))\} && \text{for all } a_0 \in \mathcal{A}, \text{ and so} \\ \sup_{a \in \mathcal{A}} \inf\{\varepsilon \mid P(a, \varepsilon)\} &\leq \inf\{\varepsilon \mid (\forall a \in \mathcal{A} : P(a, \varepsilon))\}. \end{aligned}$$

$\geq$: Let $\alpha = \sup_{a \in \mathcal{A}} \inf\{\varepsilon \mid P(a, \varepsilon)\}$. Then for all $a \in \mathcal{A}$, $\alpha \geq \inf\{\varepsilon \mid P(a, \varepsilon)\}$. Let $\gamma > \alpha$. Then we have

$$\begin{aligned} \text{for all } a \in \mathcal{A} \quad \gamma &\in \{\varepsilon \mid P(a, \varepsilon)\} && \text{by (6) and because } \gamma > \inf\{\varepsilon \mid P(a, \varepsilon)\} \\ \text{so} \quad \gamma &\in \bigcap_{a \in \mathcal{A}} \{\varepsilon \mid P(a, \varepsilon)\} \\ \text{so} \quad \gamma &\in \{\varepsilon \mid (\forall a \in \mathcal{A} : P(a, \varepsilon))\} \\ \text{so} \quad \gamma &\geq \inf\{\varepsilon \mid (\forall a \in \mathcal{A} : P(a, \varepsilon))\}. \end{aligned}$$

This implies $\alpha \geq \inf\{\varepsilon \mid (\forall a \in \mathcal{A} : P(a, \varepsilon))\}$ as well. Otherwise, suppose $\alpha < \iota := \inf\{\varepsilon \mid (\forall a \in \mathcal{A} : P(a, \varepsilon))\}$. Then $\alpha < (\alpha + \iota)/2 < \iota$, so the result above with $\gamma = (\alpha + \iota)/2$ yields $(\alpha + \iota)/2 \geq \iota$, a contradiction. $\blacktriangleleft$

The proof that Δ_{LP} is monotonic needs the following simple lemma.

► **Lemma 29.** *If $d_2 \sqsubseteq d_1$, we have $X_\varepsilon^{d_2} \subseteq X_\varepsilon^{d_1}$.*

Proof. Assume $d_1 \leq d_2$. Let $y \in X_\varepsilon^{d_2}$. Then there is some $x \in X$ such that $d_2(x, y) < \varepsilon$. This gives $d_1(x, y) \leq d_2(x, y) < \varepsilon$. So $y \in X_\varepsilon^{d_1}$. $\blacktriangleleft$

► **Proposition 14.** *Δ_{LP} is monotonic.*

Proof. Let $d_1 \leq d_2$. Let s_0, s_1 be states, let $a \in \mathcal{A}$ and define the two sets

$$\begin{aligned} E_1 &= \{\varepsilon \mid (\forall A \subseteq S : \tau_a(s_i, A) \leq \tau_a(s_{1-i}, A_\varepsilon^{d_1}) + \varepsilon, \quad i = 0, 1)\} \\ \text{and } E_2 &= \{\varepsilon \mid (\forall A \subseteq S : \tau_a(s_i, A) \leq \tau_a(s_{1-i}, A_\varepsilon^{d_2}) + \varepsilon, \quad i = 0, 1)\}. \end{aligned}$$

Then $E_2 \subseteq E_1$. Indeed, let $\varepsilon \in E_2$ and $A \subseteq S$. Then $\tau_a(s_i, A) \leq \tau_a(s_{1-i}, A_\varepsilon^{d_2}) + \varepsilon \leq \tau_a(s_{1-i}, A_\varepsilon^{d_1}) + \varepsilon$, because $A_\varepsilon^{d_2} \subseteq A_\varepsilon^{d_1}$ by Lemma 29. So $\inf E_1 \leq \inf E_2$ and hence $\Delta_{LP}(d_1) = \sup_{a \in \mathcal{A}} \inf E_1 \leq \sup_{a \in \mathcal{A}} \inf E_2 = \Delta_{LP}(d_2)$, as wanted. $\blacktriangleleft$

► **Proposition 15.** *d^* is a fixpoint of Δ_{LP} , and hence it is the greatest fixpoint.*

Proof. Let $s_0, s_1 \in S$. We want $\Delta_{LP}(d^*)(s_0, s_1) = d^*(s_0, s_1)$. Consider the following sets

$$\begin{aligned} E^\Delta &= \{\varepsilon \mid \forall a \in \mathcal{A} : (\forall X \subseteq S : \tau_a(s_i, X) \leq \tau_a(s_{1-i}, X_\varepsilon^{d^*}) + \varepsilon, \quad i = 0, 1)\} \\ \text{and } E^* &= \{\varepsilon \mid s_0 \sim_\varepsilon s_1\} \\ &= \{\varepsilon \mid \exists R \text{ an } \varepsilon\text{-bisimulation s.t. } s_0 R s_1, \text{ that is, for all } (t_0, t_1) \in R, \text{ we have} \\ &\quad \forall a \in \mathcal{A} : (\forall X \subseteq S : \tau_a(t_i, X) \leq \tau_a(t_{1-i}, R(X)) + \varepsilon, \quad i = 0, 1)\}. \end{aligned}$$

Then $\Delta(d^*)(s_0, s_1) = \inf E^\Delta$ by Lemma 10 and $d^*(s_0, s_1) = \inf E^*$. We prove $\inf E^\Delta = \inf E^*$.

$\geq$: We prove that $E^\Delta \subseteq E^*$. Let $\varepsilon \in E^\Delta$. Let $R_\varepsilon = \{(x, y) \mid x \sim_\varepsilon y\}$, that is, the biggest ε -bisimulation. We need to prove that $s_0 \sim_\varepsilon s_1$. For that we prove that $R_\varepsilon^+ := R_\varepsilon \cup \{(s_0, s_1)\}$ is an ε -bisimulation. Let $a \in \mathcal{A}$ and $X \subseteq S$. Then $X_\varepsilon^{d^*} \subseteq R_\varepsilon^+(X)$. Indeed, let $y \in X_\varepsilon^{d^*}$ then $d^*(y, x) < \varepsilon$ for some $x \in X$. So, because $d^*(y, x) = \inf\{e \mid y \sim_e x\}$, we have $y \sim_\gamma x$ for some $\gamma < \varepsilon$, and hence $y \sim_\varepsilon x$ as every γ -bisimulation is an ε -bisimulation. Hence $y \in R_\varepsilon^+(X)$. We only need to prove the bisimulation condition for the pair (s_0, s_1) . Then, for $i = 0, 1$ we have

$$\begin{aligned} \tau_a(s_i, X) &\leq \tau_a(s_{1-i}, X_\varepsilon^{d^*}) + \varepsilon && \text{by choice of } \varepsilon \\ &\leq \tau_a(s_{1-i}, R_\varepsilon^+(X)) + \varepsilon && \text{because } X_\varepsilon^{d^*} \subseteq R_\varepsilon^+(X). \end{aligned}$$

So $s_0 \sim_\varepsilon s_1$ and hence $\varepsilon \in E^*$.

$\leq$: Let $\varepsilon \in E^*$. We prove that for all $\gamma > \varepsilon$, we have $\gamma \in E^\Delta$. This will prove that $\inf E^\Delta \leq \varepsilon$. Since ε is arbitrary, this will give the result: $\inf E^\Delta \leq \inf E^*$.

So let $\gamma > \varepsilon$. Because $\varepsilon \in E^*$, we have $s_0 \sim_\varepsilon s_1$, and hence there is some ε -bisimulation R such that $s_0 R s_1$. We want $\gamma \in E^\Delta$, so let $a \in \mathcal{A}$ and $X \subseteq S$. We want $\tau_a(s_i, X) \leq \tau_a(s_{1-i}, X_\gamma^{d^*}) + \gamma$, $i = 0, 1$. First observe that $R(X) \subseteq X_\gamma^{d^*}$. Indeed, if $y \in R(X)$, then there is some $x \in X$ such that $y R x$. Because R is an ε -bisimulation, $d^*(y, x) \leq \varepsilon$, so $y \in X_\gamma^{d^*}$ (we cannot say $y \in X_\varepsilon^{d^*}$ because these sets are *open* balls). Then, for $i = 0, 1$:

$$\begin{aligned} \tau_a(s_i, X) &\leq \tau_a(s_{1-i}, R(X)) + \varepsilon, && \text{because } R \text{ is an } \varepsilon\text{-bisimulation} \\ &\leq \tau_a(s_{1-i}, X_\gamma^{d^*}) + \gamma, && \text{because } R(X) \subseteq X_\gamma^{d^*} \text{ and } \varepsilon < \gamma. \end{aligned}$$

So, $\gamma \in E^\Delta$, as wanted. We have proven that d^* is a fixpoint of Δ_{LP} , and it is the greatest by Proposition 12. $\blacktriangleleft$

► **Proposition 16.** $\mathcal{D}_{LP}$ is a functor on $\mathbf{PMet}_1$.

Proof. The only thing to prove is that $\mathcal{D}_{LP}f$ is nonexpansive. Let $\varphi_0, \varphi_1 \in \mathcal{D}_{LP}\langle X_0, d_0 \rangle$. Then (we omit the symmetric inequality for simplicity)

$$\begin{aligned} \delta_{LP}^{d_1}(\mathcal{D}_{LP}f(\varphi_0), \mathcal{D}_{LP}f(\varphi_1)) &= \inf\{\varepsilon \mid \forall B \subseteq X_1 : \varphi_0(f^{-1}(B)) \leq \varphi_1(f^{-1}(B_\varepsilon^{d_1})) + \varepsilon\} \\ &\stackrel{(*)}{\leq} \inf\{\varepsilon \mid \forall B \subseteq X_1 : \varphi_0(f^{-1}(B)) \leq \varphi_1((f^{-1}B)_\varepsilon^{d_0}) + \varepsilon\} \\ &= \inf\{\varepsilon \mid \forall A \in \{f^{-1}(B) \mid B \subseteq X_1\} : \varphi_0(A) \leq \varphi_1(A_\varepsilon^{d_0}) + \varepsilon\} \\ &\stackrel{(**)}{\leq} \inf\{\varepsilon \mid \forall A \subseteq X_0 : \varphi_0(A) \leq \varphi_1(A_\varepsilon^{d_0}) + \varepsilon\} \\ &= \delta_{LP}^{d_0}(\varphi_0, \varphi_1) \end{aligned}$$

where the marked inequalities are justified below. For the first one, marked with $(*)$, let

$$V = \{\varepsilon \mid \forall B \subseteq X_1 : \varphi_0(f^{-1}(B)) \leq \varphi_1((f^{-1}B)_\varepsilon^{d_0}) + \varepsilon\} \text{ and}$$

$$W = \{\varepsilon \mid \forall B \subseteq X_1 : \varphi_0(f^{-1}(B)) \leq \varphi_1(f^{-1}(B_\varepsilon^{d_1})) + \varepsilon\}.$$

We show that $V \subseteq W$. Let $\varepsilon \in V$. For every $B \subseteq X_1$, let us first show that $f^{-1}(B)_\varepsilon^{d_0} \subseteq f^{-1}(B_\varepsilon^{d_1})$. Indeed, take $x \in f^{-1}(B)_\varepsilon^{d_0}$ and let $a \in f^{-1}(B)$ be such that $d_0(x, a) < \varepsilon$. Then by nonexpansivity of f ,

$$d_1(f(x), f(a)) \leq d_0(x, a) < \varepsilon$$

and so $f(x) \in B_\varepsilon^{d_1}$, because $f(a) \in B$. So $x \in f^{-1}(B_\varepsilon^{d_1})$. Now combining this with the inequality condition in V we get

$$\varphi_0(f^{-1}(B)) \leq \varphi_1(f^{-1}(B)_\varepsilon^{d_0}) + \varepsilon \leq \varphi_1(f^{-1}(B_\varepsilon^{d_1})) + \varepsilon$$

and so $\varepsilon \in W$.

For the inequality marked with $(**)$, observe that if ε satisfies the inequality for all $A \subseteq X_0$ then it does for all $A \in \{f^{-1}(B) \mid B \subseteq X_1\}$. So the infimum is taken on a possibly bigger set on the left-hand side of the inequality, which concludes the argument. $\blacktriangleleft$

► **Proposition 17.** *The functor $\mathcal{D}_{LP}$ is locally nonexpansive, that is, for $f_1, f_2 \in X \rightarrow Y$*

$$\delta_{\mathcal{D}_{LP}X \rightarrow \mathcal{D}_{LP}Y}(\mathcal{D}_{LP}f_1, \mathcal{D}_{LP}f_2) \leq d_{X \rightarrow Y}(f_1, f_2)$$

with $\delta_{\mathcal{D}_{LP}X \rightarrow \mathcal{D}_{LP}Y}$ being the metric on the hom-set $\mathcal{D}_{LP}\langle X, d_X \rangle \rightarrow \mathcal{D}_{LP}\langle Y, d_Y \rangle$.

Proof. Let $f_1, f_2 \in X \rightarrow Y$. We have that

$$\delta_{\mathcal{D}_{LP}X \rightarrow \mathcal{D}_{LP}Y}(\mathcal{D}_{LP}f_1, \mathcal{D}_{LP}f_2) = \sup_{\varphi \in \mathcal{D}X} \delta_{LP}^{d_Y}(\mathcal{D}_{LP}f_1(\varphi), \mathcal{D}_{LP}f_2(\varphi)).$$

Let $\varphi \in \mathcal{D}X$ and recall that $\mathcal{D}_{LP}f_i(\varphi) = \varphi(f_i^{-1}(\cdot))$. We will show that

$$\delta_{LP}^{d_Y}(\varphi(f_1^{-1}(\cdot)), \varphi(f_2^{-1}(\cdot))) \leq d_{X \rightarrow Y}(f_1, f_2).$$

Let $\alpha = d_{X \rightarrow Y}(f_1, f_2) = \sup_{x \in X} d_Y(f_1(x), f_2(x))$. Let $B \subseteq Y$. We show $f_1^{-1}(B) \subseteq f_2^{-1}(B_\alpha^{d_Y})$

$$\begin{aligned} f_1^{-1}(B) &= \{x \mid \exists y \in B : y = f_1(x)\} \\ &= \{x \mid \exists y \in B : y = f_1(x) \text{ and } d_Y(y, f_2(x)) \leq \alpha\} \quad \text{since } d_Y(f_1(x), f_2(x)) \leq \alpha \\ &\subseteq \{x \mid \exists y \in B : d_Y(y, f_2(x)) \leq \alpha\} \\ &= \{x \mid f_2(x) \in B_\alpha^{d_Y}\} \\ &= f_2^{-1}(B_\alpha^{d_Y}). \end{aligned}$$

This implies that α satisfies $(\forall B \subseteq Y : \varphi(f_1^{-1}(B)) \leq \varphi(f_2^{-1}(B_\alpha^{d_Y})) + \alpha)$, and similarly with 1 and 2 inverted.

$$\begin{aligned} \delta_{LP}^{d_Y}(\varphi(f_1^{-1}(\cdot)), \varphi(f_2^{-1}(\cdot))) &= \inf\{\varepsilon \mid \forall B \subseteq Y : \varphi(f_1^{-1}(B)) \leq \varphi(f_2^{-1}(B_\varepsilon^{d_Y})) + \varepsilon\} \\ &\leq \alpha = d_{X \rightarrow Y}(f_1, f_2), \end{aligned}$$

as wanted. $\blacktriangleleft$

► **Lemma 19.** *The unit η of $\mathcal{D}$ is nonexpansive with respect to the Lévy-Prokhorov distance. Moreover, it is an isometry, i.e., $\delta_{LP}^d(1_x, 1_y) = d(x, y)$.*

Proof. We have

$$\begin{aligned} \delta_{LP}^d(\eta_X(x), \eta_X(y)) &= \delta_{LP}^d(1_x, 1_y) \\ &= \inf\{\varepsilon \mid \forall A \subseteq X : 1_x(A) \leq 1_y(A_\varepsilon^d) + \varepsilon, 1_y(A) \leq 1_x(A_\varepsilon^d) + \varepsilon\}. \end{aligned}$$

Let A be a subset of X . We distinguish four cases in which we consider the relevant inequalities:

$$1_x(A) \leq 1_y(A_\varepsilon^d) + \varepsilon \tag{7}$$

$$\text{and } 1_y(A) \leq 1_x(A_\varepsilon^d) + \varepsilon \tag{8}$$

1. $x \notin A, y \notin A$: Then $1_x(A) = 0$ and $1_y(A) = 0$, and any $\varepsilon \geq 0$ satisfies both (7) and (8).
2. $x \notin A, y \in A$: Here $1_x(A) = 0$ and $1_y(A) = 1$, and (7) holds always, but (8) holds for $\varepsilon > d(x, y)$ as then $x \in A_\varepsilon^d$ and hence $1_x(A_\varepsilon^d) = 1$, and might not hold for $\varepsilon \leq d(x, y)$, e.g. when $A = \{y\}$ and $\varepsilon < d(x, y) \leq 1$, as then $1_x(A_\varepsilon^d) = 0$.
3. $x \in A, y \notin A$: This case is dual to the second case. Here (8) always holds, but (7) holds for $\varepsilon > d(x, y)$ and might not hold for $\varepsilon \leq d(x, y)$ as, e.g., for $A = \{x\}$ we have $y \in A_\varepsilon^d$ if and only if $\varepsilon > d(x, y)$.
4. $x \in A, y \in A$: Then $1_x(A) = 1$ and $1_y(A) = 1$, as well as $1_x(A_\varepsilon^d) = 1_y(A_\varepsilon^d) = 1$, and again, as in the first case, any $\varepsilon \geq 0$ satisfies both (7) and (8).

Hence, both inequalities are satisfied for all A iff $\varepsilon > d(x, y)$, and therefore $\delta_{LP}^d(1_x, 1_y) = d(x, y)$. $\blacktriangleleft$

► **Lemma 27.** *Assume that Condition 1. in Definition 26 holds for $\beta \in \mathcal{DR}$ with $R \subseteq X \times Y$ and $\mu \in \mathcal{DX}$. Then $\mu(X) \leq \sum_{x \in X, y \in Y} \beta(x, y) + \varepsilon$ (i.e., Condition 3. in Definition 26) is equivalent to the condition*

$$\mu(S) \leq \sum_{x \in S, y \in Y} \beta(x, y) + \varepsilon, \quad \text{for all } S \subseteq X.$$

Proof. The right-to-left implication is immediate, as $X \subseteq X$. For the left-to-right implication, assume that for some $S \subseteq X$ we have

$$\mu(S) > \sum_{x \in S, y \in Y} \beta(x, y) + \varepsilon.$$

Since, by assumption, for all $x \in X$, and hence for all $x \in X \setminus S$, $\mu(x) \geq \sum_{y \in Y} \beta(x, y)$, we have

$$\begin{aligned} \mu(X) &= \mu(S) + \mu(X \setminus S) \\ &> \sum_{x \in S, y \in Y} \beta(x, y) + \varepsilon + \sum_{x \in X \setminus S} \mu(x) \\ &\geq \sum_{x \in S, y \in Y} \beta(x, y) + \sum_{x \in X \setminus S} \sum_{y \in Y} \beta(x, y) + \varepsilon \\ &= \sum_{x \in X, y \in Y} \beta(x, y) + \varepsilon. \end{aligned}$$

The property now follows by contraposition. $\blacktriangleleft$

► **Proposition 28.** *Let (X, c) and (Y, d) be two $\mathcal{D}$ -coalgebras. The following three properties are equivalent for $R \subseteq X \times Y$:*

1. R is an ε -simulation.
2. For every $(x, y) \in R$, there is an ε -coupling $\beta \in \mathcal{DR}$ of $\mu = c(x)$ and $\nu = d(y)$.
3. The “ ε -lax-bounded” span of morphisms on the left below commutes.

$$\begin{array}{ccc} X & \xleftarrow{\pi_1} R & \xrightarrow{\pi_2} Y \\ c \downarrow & \sqsupseteq \sqsubseteq \varepsilon & b \downarrow \sqsubseteq \downarrow d \\ \mathcal{D}X & \xleftarrow{\mathcal{D}\pi_1} \mathcal{D}R & \xrightarrow{\mathcal{D}\pi_2} \mathcal{D}Y \end{array} \qquad \begin{array}{ccc} X & \xleftarrow{\pi_1} R & \xrightarrow{\pi_2} X \\ c \downarrow & \sqsupseteq \sqsubseteq \varepsilon & b \downarrow \varepsilon \sqsupseteq \sqsubseteq \downarrow d \\ \mathcal{D}X & \xleftarrow{\mathcal{D}\pi_1} \mathcal{D}R & \xrightarrow{\mathcal{D}\pi_2} \mathcal{D}X \end{array}$$

26:24 ε -Distance via Lévy-Prokhorov Lifting

Proof. The equivalence of 1. and 2. has been shown in [54]. Let R be an ε -simulation. For $(x, y) \in R$, let $\mu = c(x)$ and $\nu = d(y)$. Define $b: R \rightarrow \mathcal{D}R$ by $(x, y) \mapsto \beta$, the epsilon coupling of μ and ν that exists.

Then, unfolding the definitions, we get:

- $\mathcal{D}\pi_1 \circ b \sqsubseteq c \circ \pi_1$ is equivalent to: for all $(x, y) \in R$, condition 1. from Definition 26 holds.
- $\mathcal{D}\pi_2 \circ b \sqsubseteq d \circ \pi_2$ is equivalent to: for all $(x, y) \in R$, condition 2. from Definition 26 holds.
- $c \circ \pi_1 \sqsubseteq_{\varepsilon} \mathcal{D}\pi_1 \circ b$ is equivalent to: for all $(x, y) \in R$, for all $S \subseteq X$,

$$\mu(S) \leq \sum_{x \in S, y \in Y} \beta(x, y) + \varepsilon$$

which by Lemma 27 is equivalent to condition 3. from Definition 26.

These facts yield the equivalence of 3. with 2. (and hence 1. as well). ◀