

Generalized Bisimulation Metrics^{*}

Konstantinos Chatzikokolakis^{1,2}, Daniel Gebler³,
Catuscia Palamidessi^{4,2}, and Lili Xu^{2,5}

¹ CNRS

² LIX, Ecole Polytechnique

³ VU University Amsterdam

⁴ INRIA

⁵ Institute of Software, Chinese Academy of Science

Abstract. The bisimilarity pseudometric based on the Kantorovich lifting is one of the most popular metrics for probabilistic processes proposed in the literature. However, its application in verification is limited to linear properties. We propose a generalization of this metric which allows to deal with a wider class of properties, such as those used in security and privacy. More precisely, we propose a family of metrics, parametrized on a notion of distance which depends on the property we want to verify. Furthermore, we show that the members of this family still characterize bisimilarity in terms of their kernel, and provide a bound on the corresponding metrics on traces. Finally, we study the case of a metric corresponding to differential privacy. We show that in this case it is possible to have a dual form, easier to compute, and we prove that the typical constructs of process algebra are non-expansive with respect to this metrics, thus paving the way to a modular approach to verification.

1 Introduction

Originally proposed in the seminal works of Desharnais et al. [17,18], the bisimilarity pseudometric based on the Kantorovich lifting has become very popular in the process algebra community. One reason of this success is that, when dealing with probabilistic processes, metrics are more suitable than equivalences, since the latter are not robust wrt small variation of probabilities. Another important reason is that, thanks to the dual presentation of the Kantorovich lifting in terms of the mass transportation problem, the metric can be computed using linear programming algorithms [4,7,2]. Furthermore, this metric is an extension of probabilistic bisimilarity, in the sense that two states have distance 0 if and only if they are bisimilar. In fact, the metric also shares with bisimilarity the fact of being based on a similar coinductive definition. More precisely, it is defined as the greatest fixpoint of a transformation that has the same structure

^{*} This work has been partially supported by the project ANR-12-IS02-001 PACE, the project ANR-11-IS02-0002 LOCALI, the INRIA Equipe Associée PRINCESS, the INRIA Large Scale Initiative CAPPRIS, and the EU grant 295261 MEALS.

as the one used for bisimilarity.¹ This allows to transfer some of the concepts and methods that have been extensively explored in process algebra, and to use lines of reasoning which the process algebra community is familiar with. Along the same lines, a nice property of the Kantorovich bisimilarity pseudometric is that the standard operators of process algebra are not expansive wrt it. This can be seen as a generalization of the result that bisimulation is a congruence, and can be used in a similar way, for compositional reasoning and verification.

Last but not least, the Kantorovich bisimilarity metric provides a bound on the corresponding distance on probabilistic traces [12] (corresponding in the sense that the definition is based on the same Kantorovich lifting). This means that it can be used to verify certain probabilistic properties on traces. More specifically, it can be used to verify properties that are expressed in terms of difference between probabilities of sets of traces. These properties are linear, in the sense that the difference increases linearly wrt variations on the distributions.

Many properties, however, such as several privacy and security ones, are not linear. This is the case of the popular property of differential privacy [19], which is expressed in terms of ratios of probabilities. In fact, there are processes that have small Kantorovich distance, and which are not ϵ -differentially private for any finite ϵ . Another example are the properties used in quantitative information flow, which involve logarithmic functions on probabilities.

The purpose of this work is to generalize the Kantorovich lifting to obtain a family of metrics suitable for the verification of a wide class of properties, following the principles that:

- i. the metrics of this family should depend on a parameter related to the class of properties (on traces) that we wish to verify,
- ii. each metric should provide a bound on the corresponding metric on traces,
- iii. the kernel of these metric should correspond to probabilistic bisimilarity,
- iv. the general construction should be coinductive,
- v. the typical process-algebra operators should be non-expansive,
- vi. it should be feasible to compute these metrics.

In this paper we have achieved the first four desiderata. For the last two, so far we have studied the particular case of the multiplicative variant of the Kantorovich metric, which is based on the notion of distance used in the definition of differential privacy. We were able to find a dual form of the lifting, which allows to reduce the problem of its computation to a linear optimization problem solvable with standard algorithms. We have also proved that several typical process-algebra operators are non-expansive, and we have given explicitly the expression of the bound. For some of them we were able to prove this result in a general form, i.e., non-expansiveness wrt all the metrics of the family, and with the bound represented by the same expression.

¹ In the original definition the Kantorovich bisimilarity pseudometric was defined as the greatest fixpoint, but such definition requires using the reverse order on metrics. More recently, authors tend to use the natural order, and define the bisimilarity metric as the least fixpoint, see [12,1,2]. Here we follow the latter approach.

As an example of application of our method, we show of to instantiate our construction to obtain the multiplicative variant of the Kantorovich metric, and how to use it to verify the property of differential privacy.

All proofs are given in the report version of this paper [11].

Related Work. Bisimulation metrics based on the standard Kantorovich distance have been used in various applications, such as systems biology [25], games [9], planning [13] and security [8]. We consider in this paper discrete state spaces. Bisimulation metrics on uncountable state spaces have been explored in [18]. We define bisimulation metrics as fixed point of an appropriate functor. Alternative characterizations were provided in terms of coalgebras [6] and real-valued modal logics [18]. The formulation of the Kantorovich lifting as primal and dual linear program is due to [5].

Verification of differential privacy has been itself an active area of research. Prominent approaches based on formal methods are those based on type systems [22] and logical formulations [3]. Earlier papers [26,27] define a bisimulation distance, which however suffered from the fact that the respective kernel relation (states in distance 0) does not fully characterize probabilistic bisimilarity.

2 Preliminaries

2.1 Labelled Concurrent Markov Chains

Given a set X , we denote by $Prob(X), Disc(X)$ the set of all and discrete probability measures over X respectively; the support of a measure μ is defined as $supp(\mu) = \{x \in X \mid \mu(x) > 0\}$. A *labelled concurrent Markov chain* (henceforth LCMC) \mathcal{A} is a tuple (S, A, D) where S is a countable set of *states*, A is a countable set of *action labels*, and $D \subseteq S \times A \times Disc(S)$ is a *transition relation*. We write $s \xrightarrow{a} \mu$ for $(s, a, \mu) \in D$.

An *execution* α is a (possibly infinite) sequence $s_0 a_1 s_1 a_2 s_2 \dots$ of alternating states and labels, such that for each $i : s_i \xrightarrow{a_{i+1}} \mu_{i+1}$ and $\mu_{i+1}(s_{i+1}) > 0$. We use $lstate(\alpha)$ to denote the last state of a finite execution α . We use $Exec^*(\mathcal{A})$ and $Exec(\mathcal{A})$ to represent the set of finite executions and of all executions of \mathcal{A} , respectively. A *trace* is a sequence of labels in $A^* \cup A^\omega$ obtained from executions by removing the states. We use $[\]$ to represent the empty trace, and \frown to concatenate two traces.

A *labelled Markov chain* (henceforth LMC) \mathcal{A} is a *fully probabilistic* LCMC, namely a LCMC where from each state of \mathcal{A} there is at most one transition available. We denote by $L(s)$ and $\pi(s)$ the label and distribution of the unique transition starting from s (if any).

In a LMC \mathcal{A} , a state s of \mathcal{A} induces a probability measure over traces as follows. The basic measurable events are the cones of finite traces, where the cone of a finite trace \mathbf{t} , denoted by $C_{\mathbf{t}}$, is the set $\{\mathbf{t}' \in A^* \cup A^\omega \mid \mathbf{t} \leq \mathbf{t}'\}$, where \leq is the standard prefix preorder on sequences. The probability induced by s on a cone $C_{\mathbf{t}}$, denoted by $\Pr[s \triangleright C_{\mathbf{t}}]$, is defined recursively as follows:

$$\Pr[s \triangleright C_t] = \begin{cases} 1 & \text{if } t = [] \\ 0 & \text{if } t = a \hat{\sim} t' \text{ and } a \neq L(s) \\ \sum_{s_i} \mu(s_i) \Pr[s_i \triangleright C_{t'}] & \text{if } t = a \hat{\sim} t' \text{ and } s \xrightarrow{a} \mu \end{cases} \quad (1)$$

This probability measure is extended to arbitrary measurable sets in the σ -algebra of traces in the standard way. We write $\Pr[s \triangleright \sigma]$ to represent the probability induced by s on the set of traces σ .

2.2 Pseudometrics

A pseudometric is a relaxed notion of a normal metric in which distinct elements can have distance zero. We consider here a generalized notion where the distance can also be infinite, and we use $[0, +\infty]$ to denote the non-negative fragment of the real numbers \mathbb{R} enriched with $+\infty$. Formally, an (extended) pseudometric on a set X is a function $m : X^2 \rightarrow [0, +\infty]$ with the following properties: $m(x, x) = 0$ (reflexivity), $m(x, y) = m(y, x)$ (symmetry), and $m(x, y) \leq m(x, z) + m(z, y)$ (triangle inequality). A metric has the extra condition that $m(x, y) = 0$ implies $x = y$. Let \mathcal{M}_X denote the set of all pseudo-metrics on X with the ordering $m_1 \preceq m_2$ iff $\forall x, y. m_1(x, y) \leq m_2(x, y)$. It can be shown that (\mathcal{M}_X, \preceq) is a complete lattice with bottom element \perp such that $\forall x, y. \perp(x, y) = 0$ and top element \top such that $\forall x, y. \top(x, y) = \infty$.

The *ball* (wrt m) of radius r centered at $x \in X$ is defined as $B_r^m(x) = \{x' \in X : m(x, x') \leq r\}$. A point $x \in X$ is called *isolated* iff there exists $r > 0$ such that $B_r^m(x) = \{x\}$. The *diameter* (wrt m) of $A \subseteq X$ is defined as $\text{diam}_m(A) = \sup_{x, x' \in A} m(x, x')$. The *kernel* $\ker(m)$ is an equivalence relation on X defined as

$$(x, x') \in \ker(m) \quad \text{iff } m(x, x') = 0$$

3 A General Family of Kantorovich Liftings

We introduce here a family of liftings from pseudometrics on a set X to pseudometrics on $\text{Prob}(X)$. This family is obtained as a generalization of the Kantorovich lifting, in which the Lipschitz condition plays a central role.

Given two pseudometric spaces $(X, d_X), (Y, d_Y)$, we say that $f : X \rightarrow Y$ is 1-Lipschitz wrt d_X, d_Y iff $d_Y(f(x), f(x')) \leq d_X(x, x')$ for all $x, x' \in X$. We denote by $1\text{-Lip}[(X, d_X), (Y, d_Y)]$ the set of all such functions.

A function $f : X \rightarrow \mathbb{R}$ can be lifted to a function $\hat{f} : \text{Prob}(X) \rightarrow \mathbb{R}$ by taking its expected value. For discrete distributions (countable X) it can be written as:

$$\hat{f}(\mu) = \sum_{x \in X} \mu(x) f(x) \quad (2)$$

while for continuous distributions we need to restrict f to be measurable wrt the corresponding σ -algebra on X , and take $\hat{f}(\mu) = \int f d\mu$.

Given a pseudometric m on X , the *standard Kantorovich lifting* of m is a pseudometric $K(m)$ on $\text{Prob}(X)$, defined as:

$$K(m)(\mu, \mu') = \sup\{|\hat{f}(\mu) - \hat{f}(\mu')| : f \in 1\text{-Lip}[(X, m), (\mathbb{R}, d_{\mathbb{R}})]\}$$

where $d_{\mathbb{R}}$ denotes the standard metric on reals. For continuous distributions we implicitly take the sup to range over measurable functions.

Generalization. A generalization of the Kantorovich lifting can be naturally obtained by extending the range of f from $(\mathbb{R}, d_{\mathbb{R}})$ to a generic metric space (V, d_V) , where $V \subseteq \mathbb{R}$ is a convex subset of the reals², and d_V is a metric on V . A function $f : X \rightarrow V$ can be lifted to a function $\hat{f} : \text{Prob}(X) \rightarrow V$ in the same way as before (cfr. (2)); the requirement that V is convex ensures that $\hat{f}(\mu) \in V$.

Then, similarly to the standard case, given a pseudometric space (X, m) we can define a lifted pseudometric $K_V(m)$ on $\text{Prob}(X)$ as:

$$K_V(m)(\mu, \mu') = \sup\{d_V(\hat{f}(\mu), \hat{f}(\mu')) : f \in 1\text{-Lip}[(X, m)(V, d_V)]\} \quad (3)$$

The subscript V in K_V is to emphasize the fact that for each choice of (V, d_V) we may get a different lifting. We should also point out the difference between m , the pseudometric on X being lifted, and d_V , the metric (not pseudo) on V which parametrizes the lifting.

The constructed $K_V(m)$ can be shown to be an extended pseudometric for any choice of (V, d_V) , i.e. it is non-negative, symmetric, identical elements have distance zero, and it satisfies the triangle inequality. However, without extra conditions, it is not guaranteed to be bounded (even if m itself is bounded). For the purposes of this paper this is not an issue. In the report version [11] we show that under the condition that d_V is *ball-convex* (i.e. all its balls are convex sets, which holds for all metrics in this paper), the following bound can be obtained:

$$K_V(m)(\mu, \mu') \leq \text{diam}_m(\text{supp}(\mu) \cup \text{supp}(\mu'))$$

Examples. The standard Kantorovich lifting is obtained by taking $(V, d_V) = (\mathbb{R}, d_{\mathbb{R}})$. When 1-bounded pseudometrics are used, like in the construction of the standard bisimilarity metric, then we can equivalently take $V = [0, 1]$.

Moreover, a multiplicative variant of the Kantorovich lifting can be obtained by taking $(V, d_V) = ([0, 1], d_{\otimes})$ (or equivalently $([0, \infty), d_{\otimes})$) where $d_{\otimes}(x, y) = |\ln x - \ln y|$. The resulting lifting is discussed in detail in Section 5 and its relation to differential privacy is shown in Section 5.1.

4 A General Family of Bisimilarity Pseudometrics

In this section we define a general family of pseudometrics on the states of an LCMC which have the property of extending probabilistic bisimilarity in the usual sense. Following standard lines, we define a transformation on state pseudometrics by first lifting a state pseudometric to a pseudometric on distributions

² V could be further generalized to be a convex subset of a vector space. It is unclear whether such a generalization would be useful, hence it is left as future work.

(over states), using the generalized Kantorovich lifting defined in previous section. Then we apply the standard Hausdorff lifting to obtain a pseudometric on sets of distributions. This last step is to take into account the nondeterminism of the LCMC, i.e., the fact that in general, from a state, we can make transitions to different distributions. The resulting pseudometric naturally corresponds to a state pseudometric, obtained by associating each set of distributions to the states which originate them. Finally, we define the intended bisimilarity pseudometric as the least fixpoint of this transformation wrt the ordering \preceq on the state pseudometrics (or equivalently, as the greatest fixpoint wrt the reverse of \preceq). We recall that $m \preceq m'$ means that $m(s, s') \leq m'(s, s')$ for all $s, s' \in S$.

Let $\mathcal{A} = (S, A, D)$ be a LCMC, let (V, d_V) be a metric space (for some convex $V \subseteq \mathbb{R}$), and let \mathcal{M} be the set of pseudometrics m on S such that $\text{diam}_m(S) \leq \text{diam}_{d_V}(V)$. Recall that $\inf \emptyset = \text{diam}_{d_V}(V)$ and $\sup \emptyset = 0$.

Definition 1. *The transformation $F_V : \mathcal{M} \rightarrow \mathcal{M}$ is defined as follows.*

$$F_V(m)(s, t) = \max\left\{ \sup_{s \xrightarrow{a} \mu} \inf_{t \xrightarrow{a} \nu} K_V(m)(\mu, \nu), \sup_{t \xrightarrow{a} \nu} \inf_{s \xrightarrow{a} \mu} K_V(m)(\nu, \mu) \right\}$$

We can also characterize F_V in terms of the following zigzag formulation:

Proposition 1. *For any $\epsilon \geq 0$, $F_V(m)(s, t) \leq \epsilon$ if and only if:*

- if $s \xrightarrow{a} \mu$, then there exists ν such that $t \xrightarrow{a} \nu$ and $K_V(m)(\mu, \nu) \leq \epsilon$,
- if $t \xrightarrow{a} \nu$, then there exists μ such that $s \xrightarrow{a} \mu$ and $K_V(m)(\nu, \mu) \leq \epsilon$.

The following result states that K_V and F_V are monotonic wrt (\mathcal{M}, \preceq) .

Proposition 2. *Let $m, m' \in \mathcal{M}$. If $m \preceq m'$ then:*

$$\begin{aligned} F_V(m)(s, s') &\leq F_V(m')(s, s') \quad \text{for all states } s, s' \\ K_V(m)(\mu, \mu') &\leq K_V(m')(\mu, \mu') \quad \text{for all distributions } \mu, \mu' \end{aligned}$$

Since (\mathcal{M}, \preceq) is a complete lattice and F_V is monotone on \mathcal{M} , by Tarski's theorem [24] F_V has a least fixpoint, which coincides with the least pre-fixpoint. We define the *bisimilarity pseudometric* bm_V as this least fixpoint:

Definition 2. *The bisimilarity pseudometric bm_V is defined as:*

$$bm_V = \min \{m \in \mathcal{M} \mid F_V(m) = m\} = \min \{m \in \mathcal{M} \mid F_V(m) \preceq m\}$$

In addition, if the states of \mathcal{A} are finite, then the closure ordinal of F_V is ω (cf: [17], Lemma 3.10). Hence we can approximate bm_V by iterating the function F_V from the bottom element:

Proposition 3. *Assume S is finite. Let $m_0 = \perp$ and $m_{i+1} = F_V(m_i)$. Then $bm_V = \sup_i m_i$.*

Next section shows that bm_V is indeed a bisimilarity metric, in the sense that its kernel coincides with probabilistic bisimilarity.

4.1 Bisimilarity as 0-distance

We now show that under certain conditions, the pseudometric constructed by $K_V(m)$ characterizes bisimilarity at its kernel. Recall that the kernel $\ker(m)$ of m is an equivalence relation relating states at distance 0.

Given an equivalence relation R on S , its lifting $\mathcal{L}(R)$ is an equivalence relation on $\text{Disc}(S)$, defined as

$$(\mu, \mu') \in \mathcal{L}(R) \quad \text{iff} \quad \forall s \in S : \mu([s]_R) = \mu'([s]_R)$$

where $[s]_R$ denotes the equivalence class of s wrt R .

To obtain the characterization result we assume that (a) the set of states is finite, and (b) the distance d_V is non-discrete. Under these conditions, the kernel operator and the lifting operator commute (cfr. [15] for the analogous property for the standard Kantorovich lifting).

Lemma 1. *If S is finite and d_V is non-discrete, then $\mathcal{L}(\ker(m)) = \ker(K_V(m))$.*

We recall the notions of probabilistic bisimulation and bisimilarity, following the formulation in terms of post-fixpoints of a transformation on state relations:

Definition 3.

- The transformation $B : S \times S \rightarrow S \times S$ is defined as: $(s, s') \in B(R)$ iff
 - if $s \xrightarrow{a} \mu$, then there exists μ' such that $t \xrightarrow{a} \mu'$ and $(\mu, \mu') \in \mathcal{L}(R)$,
 - if $s' \xrightarrow{a} \mu'$, then there exists μ such that $s \xrightarrow{a} \mu$ and $(\mu', \mu) \in \mathcal{L}(R)$.
- A relation $R \subseteq S \times S$ is called a bisimulation if it is a post-fixpoint of B , i.e. $R \subseteq B(R)$.

It is easy to see that B is monotonic on $(2^{S \times S}, \subseteq)$ and that the latter is a complete lattice, hence by Tarski's theorem there exists the greatest fixpoint of B , and it coincides with the greatest bisimulation:

Definition 4. *The bisimilarity relation $\sim \subseteq S \times S$ is defined as:*

$$\sim = \max\{R \mid R = B(R)\} = \max\{R \mid R \subseteq B(R)\} = \bigcup\{R \mid R \subseteq B(R)\}$$

We are now ready to show the correspondence between pre-fixpoint metrics and bisimulations. Using Lemma 1, we can see that the definition of B corresponds to the characterization of F_V in Proposition 1, for $\epsilon = 0$. Hence we have:

Proposition 4. *For every $m \in \mathcal{M}$, if $F_V(m) \preceq m$ then $\ker(m) \subseteq B(\ker(m))$, i.e., $\ker(m)$ is a bisimulation.*

As a consequence, $\ker(bm_V) \subseteq \sim$. The converse of Proposition 4 does not hold, because the fact that $\ker(m) \subseteq B(\ker(m))$ does not say anything about the effect of F_V on the distance between elements that are not on the kernel. However, in the case of bisimilarity we can make a connection: consider the greatest

metric m_{\sim} whose kernel coincides with bisimilarity, namely, $m_{\sim}(s, s') = 0$ if $s \sim s'$ and $m_{\sim}(s, s') = \text{diam}_{d_V}(V)$ otherwise. We have that $F_V(m_{\sim}) \preceq m_{\sim}$, and therefore $\sim = \ker(m_{\sim}) \subseteq \text{ker}(bm_V)$. Therefore we can conclude that the kernel of the bisimilarity pseudometrics coincides with bisimilarity.

Theorem 1. $\ker(bm_V) = \sim$ for every (V, d_V) ,

4.2 Relation with Trace Distributions

In this section, we show the relation between the bisimilarity metric bm_V and the corresponding metric on traces, in the case of LMCs (labeled Markov chains). Note that we restrict to the fully probabilistic case here, where probabilities on traces can be defined in the way shown in the preliminaries. The full case of LMCs can be treated by using schedulers, but a proper treatment involves imposing scheduler restrictions which complicate the formalism. Since these problems are orthogonal to the goals of this paper, we keep the discussion simple by restricting to the fully probabilistic case.

The distance between trace distributions (i.e. distributions over A^ω) will be measured by the Kantorovich lifting of the *discrete metric*. Given (V, d_V) , let $\delta_V = \text{diam}_{d_V}(V)$. Then let dm_{δ_V} be the δ_V -valued discrete metric on A^ω , defined as $dm_{\delta_V}(\mathbf{t}, \mathbf{t}') = 0$ if $\mathbf{t} = \mathbf{t}'$, and $dm_{\delta_V}(\mathbf{t}, \mathbf{t}') = \delta_V$ otherwise.

Then $K_V(dm_{\delta_V})(\mu, \mu')$ is a pseudometric on $\text{Prob}(A^\omega)$, whose kernel coincides with probabilistic trace equivalence.

Proposition 5. $K_V(dm_{\delta_V})(\mu, \mu') = 0$ iff $\mu(\sigma) = \mu'(\sigma)$ for all measurable $\sigma \subseteq A^\omega$.

The following theorem expresses that our bisimilarity metric bm_V is a bound on the distance on traces, which extends the standard relation between probabilistic bisimilarity and probabilistic trace equivalence.

Theorem 2. Let $\mu = \text{Pr}[s \triangleright \cdot]$ and $\mu' = \text{Pr}[s' \triangleright \cdot]$. Then $K_V(dm_{\delta_V})(\mu, \mu') \leq bm_V(s, s')$

It should be noted that, although the choice of $K_V(dm_{\delta_V})$ as our trace distribution metric might seem arbitrary, this metric is in fact of great interest. In the case of the standard bisimilarity pseudometric, i.e. when $(V, d_V) = ([0, 1], d_{\mathbb{R}})$, this metric is equal to the well-known *total variation distance* (also known as *statistical distance*), defined as $tv(\mu, \mu') = \sup_{\sigma} |\mu(\sigma) - \mu'(\sigma)|$:

$$K(dm_{\delta_V}) = tv \tag{4}$$

Theorem 2 reduces to the result of [12] relating the total variation distance to the bisimilarity pseudometric. Moreover, in the case of the multiplicative pseudometric, discussed in the next section, $K_V(dm_{\delta_V})$ is the same as the multiplicative distance between distributions, discussed in Section 5.1, which plays a central role in differential privacy.

Table 1. The standard Kantorovich metric and its multiplicative variant

| | Standard $K(m)(\mu, \mu')$ | Multiplicative $K_{\otimes}(m)(\mu, \mu')$ |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Primal | $\max_f \hat{f}(\mu) - \hat{f}(\mu') $ <p style="text-align: center;">subject to</p> $\forall s, s'. f(s) - f(s') \leq m(s, s')$ | $\max_f \ln \hat{f}(\mu) - \ln \hat{f}(\mu') $ <p style="text-align: center;">subject to</p> $\forall s, s'. \ln f(s) - \ln f(s') \leq m(s, s')$ |
| Dual | $\min_{\ell} \sum_{i,j} \ell_{ij} m(s_i, s_j)$ <p style="text-align: center;">subject to</p> $\forall i, j. \ell_{ij} \geq 0$ $\forall i. \sum_j \ell_{ij} = \mu(s_i)$ $\forall j. \sum_i \ell_{ij} = \mu'(s_j)$ | $\min \ln z$ <p style="text-align: center;">subject to</p> $\forall i, j. \ell_{ij}, r_i \geq 0$ $\forall i. \sum_j \ell_{ij} - r_i = \mu(s_i)$ $\forall j. \sum_i \ell_{ij} e^{m(s_i, s_j)} - r_j \leq z \cdot \mu'(s_j)$ |

5 The Multiplicative Variant

In this section we investigate the multiplicative variant of the Kantorovich pseudometric, obtained by considering as distance d_V the ratio between two numbers instead than their difference. This is the distance used to define differential privacy. We show that this variant has a dual form, which can be used to compute the metric by using linear programming techniques. In the next section, we will show how to use it to verify differential privacy.

Definition 5. *The multiplicative variant K_{\otimes} of the Kantorovich lifting is defined as the instantiation of K_V with $([0, 1], d_{\otimes})$ where $d_{\otimes}(x, y) = |\ln x - \ln y|$.*

It is well known that the standard Kantorovich metric has a dual form which can be interpreted in terms of *the Transportation Problem*, namely, the lowest total cost of transporting the mass of one distribution μ to the other distribution μ' given the distance m between locations (in our case, states). The dual form is shown in Table 1. Note that both the primal and the dual forms are linear optimization problems. The dual form is particularly suitable for computation, via standard linear programming techniques.

For our multiplicative variant, the objective function of the primal form is not a linear expression, hence the linear programming techniques cannot be applied directly. However, since $\ln \hat{f}(\mu) - \ln \hat{f}(\mu') = \ln \frac{\hat{f}(\mu)}{\hat{f}(\mu')}$ and \ln is a monotonically increasing function, the primal problem is actually a linear-fractional program. It is known that such kind of program can be converted to an equivalent linear programming problem and then to a dual program. The dual form of the multiplicative variant obtained in this way is shown in Table 1. (For the sake of simplicity, the table shows only the dual form of $\ln \hat{f}(\mu) - \ln \hat{f}(\mu')$. The dual form of $\ln \hat{f}(\mu') - \ln \hat{f}(\mu)$ can be obtained by simply switching the roles of μ and

μ' .) Hence, the multiplicative pseudometric can be computed by using linear programming techniques.

5.1 Application to Differential Privacy

Differential privacy [19] is a notion of privacy originating from the area of statistical databases, which however has been recently applied to several other areas. The standard context is that of an analyst who wants to perform a statistical query to a database. Although obtaining statistical information is permitted, privacy issues arise when this information can be linked to that of an individual in the database. In order to hide this link, differentially private mechanisms add noise to the outcome of the query, in a way such that databases differing in a single individual have similar probability of producing the same observation.

More concretely, let \mathcal{X} be the set of all databases; two databases $x, x' \in \mathcal{X}$ are *adjacent*, written $x \sim x'$, if they differ in the value of a single individual. A *mechanism* is a function $M : \mathcal{X} \rightarrow \text{Prob}(\mathcal{Z})$ where \mathcal{Z} is some set of reported values. Intuitively, $M(x)$ gives the outcome of the query when applied to database x , which is a probability distribution since noise is added.

Let tv_{\otimes} be a multiplicative variant of the total variation distance on $\text{Prob}(\mathcal{Z})$ (simply called “multiplicative distance” in [23]), defined as:

$$tv_{\otimes}(\mu, \mu') = \sup_{\mathcal{Z}} \left| \ln \frac{\mu(\mathcal{Z})}{\mu'(\mathcal{Z})} \right|$$

Then differential privacy can be defined as follows.³

Definition 6. *A mechanism $M : \mathcal{X} \rightarrow \text{Prob}(\mathcal{Z})$ is ϵ -differentially private iff*

$$tv_{\otimes}(M(x), M(x')) \leq \epsilon \quad \forall x \sim x'$$

Intuitively, the definition requires that, when run on adjacent databases, the mechanism should produce similar results, since the distance between the corresponding distributions should be bounded by ϵ (a privacy parameter).

In our setting, we assume that the mechanism M is modelled by a LMC, and the result of the mechanism running on x is the trace produced by the execution of the LMC starting from some corresponding state s_x . That is, $\mathcal{Z} = A^{\omega}$ and

$$M(x) = \text{Pr}[s_x \triangleright \cdot] \tag{5}$$

The relation between differential privacy and the multiplicative bisimilarity metric comes from the fact that tv_{\otimes} can be obtained as the K_{\otimes} lifting of the discrete metric on A^{ω} .

Lemma 2. *Let $\delta_V = \text{diam}_{d_{\otimes}}([0, 1]) = +\infty$ and let dm_{δ_V} be the discrete metric on A^{ω} . Then $tv_{\otimes} = K_{\otimes}(dm_{\delta_V})$.*

³ The definition can be generalized to an arbitrary set of secrets \mathcal{X} equipped with a “distinguishability metric” $d_{\mathcal{X}}$ [10]. The results of this section extend to this setting.




Fig. 1. The bisimilarity pseudometric bm does not imply differential privacy

Let bm_{\otimes} be the instantiation of the bisimilarity metric bm_V with K_{\otimes} . The above Lemma, together with Theorem 2, imply the following result, which makes bm_{\otimes} useful to verify differential privacy:

Theorem 3. *Let M be the mechanism defined by (5), and assume that*

$$bm_{\otimes}(s_x, s_{x'}) \leq \epsilon \quad \text{for all } x \sim x'$$

Then M satisfies ϵ -differential privacy.

Note that the use of the multiplicative bm_{\otimes} is crucial in the above result. The following example shows that the standard bisimilarity metric bm (generated by the original Kantorovich lifting) may be very different from the level of differential privacy, which is expected, since bm bounds the additive total variation metric (Theorem 2 and (4)) instead of the multiplicative tv_{\otimes} .

Example 1. Consider the processes s, t shown in Fig. 1 (a). We have that $bm(s, t) = 0.1 - 0.001 = 0.099$ while their level of differential privacy is $\epsilon = \ln \frac{0.1}{0.001} = \ln 100$. Moreover, for the processes s', t' shown in Fig. 1 (b) we have $bm(s', t') = 0.7 - 0.2 = 0.5$ while their level of differential privacy is $\epsilon' = \ln \frac{0.7}{0.2} = \ln 3.5$. Using the original Kantorovich metric, s and t are considered more indistinguishable than s' and t' , in sharp contrast to the corresponding differential privacy levels.

Approximate differential privacy. An approximate, also known as (ϵ, δ) version of differential privacy is also widely used [20], relaxing the definition by an additive factor δ . It requires that:

$$M(x)(Z) \leq e^{\epsilon} M(x')(Z) + \delta \quad \forall x \sim x', Z \subseteq \mathcal{Z}$$

The α -distance on distributions is proposed in [3] to capture (ϵ, δ) -differential privacy. For two real numbers a, b and a skew parameter $\alpha \geq 1$, the α -distance between a and b is $\max\{a - \alpha b, b - \alpha a, 0\}$. An instantiation of the Kantorovich lifting based on the α -distance seems promising for extending Theorem 3 to the approximate case; we leave this extension as future work.

Weak probabilistic anonymity. Weak probabilistic anonymity was proposed in [16] as a measure of the degree of protection of user’s identities. It is defined in a way similar to differential privacy, with the crucial difference (apart from the

lack of an adjacency relation) that it uses the (additive) total variation instead of the multiplicative one. Formally, let \mathcal{X} contain the users' identities, and let $M : \mathcal{X} \rightarrow \text{Prob}(\mathcal{Z})$ be the system in which users operate. We say that M is *ϵ -weakly probabilistically anonymous* iff $tv(M(x), M(x')) \leq \epsilon$ for all $x, x' \in \mathcal{X}$.

For systems modelled by LMCs, by (4) and Theorem 2, we have that if $bm(s_x, s_{x'}) \leq \epsilon$ for all $x, x' \in \mathcal{X}$, then M satisfies ϵ -weak probabilistic anonymity. Hence bm can be used to verify this anonymity property.

6 Process Algebra

Process algebras allow to syntactically describe probabilistic processes in terms of a small set of well-understood operators. The operational semantics of a process term is a LCMC with transitions derived from SOS rules.

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 operators should satisfy in order to facilitate compositional reasoning [18,1]. In this section we will show that the standard non-recursive process algebra operators are non-expansiveness [18] (as most prominent compositionality property) with respect to the bisimilarity metric.

We introduce a simple probabilistic process algebra that comprises the following operators i) constants 0 (stop process) and ϵ (skip process); ii) a family of n -ary prefix operators $a.([p_1]_- \oplus \dots \oplus [p_n]_-)$ with $a \in \text{Act}$, $n \geq 1$, $p_1, \dots, p_n \in (0, 1]$ and $\sum_{i=1}^n p_i = 1$; iii) binary operators $_- ; _-$ (sequential composition), $_- + _-$ (alternative composition), $_- +_p _-$ (probabilistic alternative composition), $_- | _-$ (synchronous parallel composition), $_- || _-$ (asynchronous parallel composition), and $_- ||_p _-$ (probabilistic parallel composition). We assume a set of actions Act with the distinguished action $\surd \in A$ to denote successful termination. The operational semantics of all operators is specified by the rules in Table 2.

We use distribution terms in the target of rules (right hand side of the conclusion of the rules) in order to describe distributions. We briefly recall the semantics of distribution terms of [21,14]. The expression $\delta(x)$ denotes a Dirac distribution on x . The expression $\mu; \delta(y)$ denotes a distribution such that $(\mu; \delta(y))(x; y) = \mu(x)$, the expression $\mu \oplus_p \nu$ denotes a distribution such that $(\mu \oplus_p \nu)(x) = p\mu(x) + (1-p)\nu(x)$, and $(\mu || \nu)(s || t) = \mu(s)\nu(t)$.

The probabilistic prefix operator expresses that the process $a.([p_1]t_1 \oplus \dots \oplus [p_n]t_n)$ 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 $s | t$ describes the simultaneous evolution of processes s and t , while the asynchronous parallel composition $t || t$ describes the interleaving of s and t where both processes can progress by alternating at any rate the execution of their actions. The probabilistic alternative and

Table 2. Probabilistic process algebra operators

| | | | |
|-----------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| $\varepsilon \xrightarrow{\checkmark} \delta(0)$ | $a. \bigoplus_{i=1}^n [p_i]x_i \xrightarrow{a} \bigoplus_{i=1}^n p_i \delta(x_i)$ | | |
| $\frac{x \xrightarrow{a} \mu \quad a \neq \checkmark}{x; y \xrightarrow{a} \mu; \delta(y)}$ | $\frac{x \xrightarrow{\checkmark} \mu \quad y \xrightarrow{a} \nu}{x; y \xrightarrow{a} \nu}$ | $\frac{x \xrightarrow{a} \mu}{x + y \xrightarrow{a} \mu}$ | $\frac{y \xrightarrow{a} \nu}{x + y \xrightarrow{a} \nu}$ |
| $\frac{x \xrightarrow{a} \mu \quad y \xrightarrow{a} \nu}{x \mid y \xrightarrow{a} \mu \mid \nu}$ | $\frac{x \xrightarrow{a} \mu}{x \parallel y \xrightarrow{a} \mu \parallel \delta(y)}$ | $\frac{y \xrightarrow{a} \nu}{x \parallel y \xrightarrow{a} \delta(x) \parallel \nu}$ | |
| $\frac{x \xrightarrow{a} \mu \quad y \not\xrightarrow{a}}{x +_p y \xrightarrow{a} \mu}$ | $\frac{x \not\xrightarrow{a} \quad y \xrightarrow{a} \nu}{x +_p y \xrightarrow{a} \nu}$ | $\frac{x \xrightarrow{a} \mu \quad y \xrightarrow{a} \nu}{x +_p y \xrightarrow{a} \mu \oplus_p \nu}$ | |
| $\frac{x \xrightarrow{a} \mu \quad y \not\xrightarrow{a}}{x \parallel_p y \xrightarrow{a} \mu \parallel_p \delta(y)}$ | $\frac{x \not\xrightarrow{a} \quad y \xrightarrow{a} \nu}{x \parallel_p y \xrightarrow{a} \delta(x) \parallel_p \nu}$ | $\frac{x \xrightarrow{a} \mu \quad y \xrightarrow{a} \nu}{x \parallel_p y \xrightarrow{a} \mu \parallel_p \delta(y) \oplus_p \delta(x) \parallel_p \nu}$ | |

probabilistic parallel composition replaces the nondeterministic choice of their non-probabilistic variants by a probabilistic choice. The probabilistic alternative composition $s +_p t$ evolves to the probabilistic choice between a distribution reached by s (with probability p) and a distribution reached by t (with probability $1 - p$) for actions which can be performed by both processes. For actions that can be performed by either only s or only t , the probabilistic alternative composition $s +_p t$ behaves just like the nondeterministic alternative composition $s + t$. Similarly, the probabilistic parallel composition $s \parallel_p t$ evolves to a probabilistic choice between the nondeterministic choices of asynchronous parallel composition of s and t .

We start by showing an important auxiliary property how the distance between convex combinations of probability distributions relates to the distance between the combined probability distributions.

Proposition 6. *Let $\mu_1, \mu_2, \mu'_1, \mu'_2 \in \text{Disc}(X)$ and $p \in [0, 1]$. Then*

$$K_{\otimes}(bm_{\otimes})(p\mu_1 + (1-p)\mu_2, p\mu'_1 + (1-p)\mu'_2) \leq \max(K_{\otimes}(bm_{\otimes})(\mu_1, \mu_2), K_{\otimes}(bm_{\otimes})(\mu'_1, \mu'_2))$$

Non-expansiveness is the most wildly studied compositionality property stating that the distance between composed processes is at most the sum of the distance between its parts.

Definition 7. *A n -ary operator f is non-expansive wrt a pseudometric m if*

$$m(f(s_1, \dots, s_n), f(t_1, \dots, t_n)) \leq \sum_{i=1}^n m(s_i, t_i)$$

Now we can show that all (non-recursive) operators of the probabilistic process algebra introduced above are non-expansive. In fact, we will provide upper bounds on distance between the composed processes which are in case of the (nondeterministic and probabilistic) alternative composition even stricter than the non-expansiveness condition.

Theorem 4. *Let s, t, s', t' be probabilistic processes. Then*

1. $bm_{\otimes}(s; t, s'; t') \leq bm_{\otimes}(s, s') + bm_{\otimes}(t, t')$
2. $bm_{\otimes}(s + t, s' + t') \leq \max(bm_{\otimes}(s, s'), bm_{\otimes}(t, t'))$
3. $bm_{\otimes}(s +_p t, s' +_p t') \leq \max(bm_{\otimes}(s, s'), bm_{\otimes}(t, t'))$
4. $bm_{\otimes}(s \mid t, s' \parallel t') \leq bm_{\otimes}(s, s') + bm_{\otimes}(t, t')$
5. $bm_{\otimes}(s \parallel t, s' \parallel t') \leq bm_{\otimes}(s, s') + bm_{\otimes}(t, t')$
6. $bm_{\otimes}(s \parallel_p t, s \parallel_p t') \leq bm_{\otimes}(s, s') + bm_{\otimes}(t, t')$

A similar result can be gained for the bisimilarity metric bm based on the standard Kantorovich lifting. This generalizes a similar result of [18] which considered only PTSs without nondeterministic branching and only a small set of process combinators.

For the generalized bisimilarity metric bm_V we can formulate a similar result for the nondeterministic alternative composition.

Theorem 5. *Let s, t, s', t' be probabilistic processes. Then*

$$bm_V(s + t, s' + t') \leq \max(bm_V(s, s'), bm_V(t, t'))$$

7 Conclusion and Future Work

We have proposed a family of Kantorovich pseudometrics depending on the notion of distance used to specify properties over traces. We have developed the theory of this notion, and showed how we can use it to verify the corresponding kind of properties. We have also showed that for the multiplicative variant, which is an interesting case because it corresponds to differential privacy, it is possible to give a dual form that makes the metric computable by standard techniques.

Future work include the investigation of methods to compute other members of this family, and of conditions that make possible a general dual form.

References

1. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: Computing Behavioral Distances, Compositionally. In: Chatterjee, K., Sgall, J. (eds.) MFCS 2013. LNCS, vol. 8087, pp. 74–85. Springer, Heidelberg (2013)
2. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: On-the-fly exact computation of bisimilarity distances. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013. LNCS, vol. 7795, pp. 1–15. Springer, Heidelberg (2013)
3. Barthe, G., Köpf, B., Olmedo, F., Béguelin, S.Z.: Probabilistic relational reasoning for differential privacy. In: Proc. of POPL. ACM (2012)
4. van Breugel, F., Worrell, J.B.: An algorithm for quantitative verification of probabilistic transition systems. In: Larsen, K.G., Nielsen, M. (eds.) CONCUR 2001. LNCS, vol. 2154, pp. 336–350. Springer, Heidelberg (2001)
5. van Breugel, F., Worrell, J.B.: Towards quantitative verification of probabilistic transition systems. In: Orejas, F., Spirakis, P.G., van Leeuwen, J. (eds.) ICALP 2001. LNCS, vol. 2076, pp. 421–432. Springer, Heidelberg (2001)
6. van Breugel, F., Worrell, J.: A behavioural pseudometric for probabilistic transition systems. Theor. Comp. Sci. 331(1), 115–142 (2005)

7. van Breugel, F., Worrell, J.: Approximating and computing behavioural distances in probabilistic transition systems. *Theor. Comp. Sci.* 360(1-3), 373–385 (2006)
8. Cai, X., Gu, Y.: Measuring anonymity. In: Bao, F., Li, H., Wang, G. (eds.) *ISPEC 2009*. LNCS, vol. 5451, pp. 183–194. Springer, Heidelberg (2009)
9. Chatterjee, K., de Alfaro, L., Majumdar, R., Raman, V.: Algorithms for Game Metrics. In: *FSTTCS*, vol. 2, pp. 107–118. Leibniz-Zentrum fuer Informatik (2008)
10. Chatzikokolakis, K., Andrés, M.E., Bordenabe, N.E., Palamidessi, C.: Broadening the scope of Differential Privacy using metrics. In: De Cristofaro, E., Wright, M. (eds.) *PETS 2013*. LNCS, vol. 7981, pp. 82–102. Springer, Heidelberg (2013)
11. Chatzikokolakis, K., Gebler, D., Palamidessi, C., Xu, L.: Generalized bisimulation metrics. Tech. rep., INRIA (2014)
12. Chen, D., van Breugel, F., Worrell, J.: On the complexity of computing probabilistic bisimilarity. In: Birkedal, L. (ed.) *FOSSACS 2012*. LNCS, vol. 7213, pp. 437–451. Springer, Heidelberg (2012)
13. Comanici, G., Precup, D.: Basis function discovery using spectral clustering and bisimulation metrics. In: Vrancx, P., Knudson, M., Grześ, M. (eds.) *ALA 2011*. LNCS, vol. 7113, pp. 85–99. Springer, Heidelberg (2012)
14. D’Argenio, P.R., Gebler, D., Lee, M.D.: Axiomatizing Bisimulation Equivalences and Metrics from Probabilistic SOS Rules. In: Muscholl, A. (ed.) *FOSSACS 2014*. LNCS, vol. 8412, pp. 289–303. Springer, Heidelberg (2014)
15. Deng, Y., Du, W.: The kantorovich metric in computer science: A brief survey. *ENTCS* 253(3), 73–82 (2009)
16. Deng, Y., Palamidessi, C., Pang, J.: Weak probabilistic anonymity. In: *Proc. of SecCo*. *ENTCS*, vol. 180 (1), pp. 55–76. Elsevier (2007)
17. Desharnais, J., Jagadeesan, R., Gupta, V., Panangaden, P.: The metric analogue of weak bisimulation for probabilistic processes. In: *Proc. of LICS*, pp. 413–422. IEEE (2002)
18. Desharnais, J., Jagadeesan, R., Gupta, V., Panangaden, P.: Metrics for labelled Markov processes. *Theor. Comp. Sci.* 318(3), 323–354 (2004)
19. Dwork, C.: Differential privacy. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) *ICALP 2006*. LNCS, vol. 4052, pp. 1–12. Springer, Heidelberg (2006)
20. Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., Naor, M.: Our data, ourselves: Privacy via distributed noise generation. In: Vaudenay, S. (ed.) *EUROCRYPT 2006*. LNCS, vol. 4004, pp. 486–503. Springer, Heidelberg (2006)
21. 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 2012*. *EPTCS*, vol. 89, pp. 115–130 (2012)
22. Reed, J., Pierce, B.C.: Distance makes the types grow stronger: A calculus for differential privacy. In: *Proc. of ICFP*, pp. 157–168. ACM (2010)
23. Smith, A.: Efficient, differentially private point estimators. arXiv preprint arXiv:0809.4794 (2008)
24. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* 5(2), 285–309 (1955)
25. Thorsley, D., Klavins, E.: Approximating stochastic biochemical processes with wasserstein pseudometrics. *Systems Biology, IET* 4(3), 193–211 (2010)
26. Tschantz, M.C., Kaynar, D., Datta, A.: Formal verification of differential privacy for interactive systems (extended abstract). *ENTCS* 276, 61–79 (2011)
27. Xu, L., Chatzikokolakis, K., Lin, H.: Metrics for differential privacy in concurrent systems. In: Ábrahám, E., Palamidessi, C. (eds.) *FORTE 2014*. LNCS, vol. 8461, pp. 199–215. Springer, Heidelberg (2014)