A Logical Characterization of Differential Privacy via Behavioral Metrics

Valentina Castiglioni¹, Konstantinos Chatzikokolakis², and Catuscia Palamidessi³

¹ INRIA Saclay - Île de France, France

² CNRS and LIX Ecole Polytechnique, France

³ INRIA Saclay - Île de France and LIX Ecole Polytechnique, France

Abstract. *Differential privacy* is a formal definition of privacy ensuring that sensitive information relative to individuals cannot be inferred by querying a database. In this paper, we exploit a modeling of this framework via labeled Markov Chains (LMCs) to provide a *logical characterization of differential privacy*: we consider a probabilistic variant of the Hennessy-Milner logic and we define a *syntactical distance* on formulae in it measuring their syntactic disparities. Then, we define a *trace distance* on LMCs in terms of the syntactic distance between the sets of formulae satisfied by them. We prove that such distance *corresponds* to the *level of privacy* of the LMCs. Moreover, we use the distance on formulae to define a real-valued semantics for them, from which we obtain a *logical characterization of weak anonymity*: the level of anonymity is measured in terms of the smallest formula distinguishing the considered LMCs. Then, we focus on *bisimulation semantics* on nondeterministic probabilistic processes and we provide a *logical characterization of generalized bisimulation metrics*, namely those defined via the *generalized Kantorovich lifting*. Our characterization is based on the notion of *mimicking formula of a process* and the *syntactic distance* on formulae, where the former captures the observable behavior of the corresponding process and allows us to characterize bisimilarity. We show that the generalized bisimulation distance on processes is equal to the syntactic distance on their mimicking formulae. Moreover, we use the distance on mimicking formulae to obtain *bounds* on differential privacy.

1 Introduction

With the ever-increasing use of internet-connected devices, such as computers, IoT appliances and GPS-enabled equipment, personal data are collected in larger and larger amounts, and then stored and manipulated for the most diverse purposes. The exposure of personal data raises all kinds of privacy threats, and it has motivated researchers to develop theories and techniques to protect users from these risks.

The state of the art in privacy research is represented by *differential privacy* (DP) [21], a framework originally proposed for protecting the privacy of participants in statistical databases, and now applied to geolocation [34], social

networks [35] and many other domains. DP is based on the idea of obfuscating the link between the answers to queries and the personal data by adding controlled (probabilistic) noise to the answers. One of the main advantages of DP with respect to previous approaches is its compositionality. Namely, if we combine the information that we obtain by querying two differentially-private mechanisms, the resulting mechanism is also differentially-private.

Recently, a distributed variant of DP has emerged, *called local differential privacy* (LDP) [20]. In this variant, users obfuscate their personal data by themselves, before sending them to the data collector. In this way, the data collector can only see, stock and analyze the obfuscated data. LDP, like DP, is compositional, and furthermore it has the further advantages that it does not need to trust the data collector. LDP is having a considerable impact, specially after large companies such as Apple and Google have started to adopt it for collecting the data of their users for statistical purposes [22].

In this paper, we consider $d_{\mathcal{X}}$ -privacy [11], a metric-based generalization of differential privacy that subsumes both DP and LDP by exploiting a metric in the domain of secrets to capture the desired privacy protection semantics. We study $d_{\mathcal{X}}$ -privacy in the context of probabilistic transition systems (PTSs) and labeled Markov chains (LMCs), aiming at importing the rich concepts and techniques that have been developed in the area of Concurrency Theory. In particular, we focus on *behavioral metrics* and on their logical counterparts, exploring their use to specify privacy properties.

Behavioral metrics [1, 8, 12, 13, 17, 25, 31, 32, 39] represent the quantitative analogue of behavioral equivalences and preorders measuring the disparities in the behavior of processes. Here, we consider a probabilistic extension $\mathbb{L}$ of the Hennessy-Milner logic (HML) [28] and we propose a notion of *trace metric* defined via a *syntactic distance* over formulae in $\mathbb{L}$, namely a pseudometric on formulae measuring their syntactic disparities. Informally, we consider formulae expressing probabilistic linear properties and we define the trace metric between two processes as the Hausdorff lifting of the syntactic distance over the sets of formulae satisfied by them. Such trace metric will allow us to obtain the first *logical characterization of $d_{\mathcal{X}}$ -privacy*.

Although $d_{\mathcal{X}}$ -privacy is defined in terms of a *multiplicative* variant of the *total variation* distance, one could also define privacy properties based on the standard total variation, as in the case of *weak anonymity* from [16]. We exploit the distance on formulae to define a real-valued semantics for them, from which we obtain a *logical characterization of weak anonymity*.

Then we switch from trace to bisimulation semantics and we provide a *logical bound* on $d_{\mathcal{X}}$ -privacy. We consider the generalized notion of Kantorovich lifting [12], which allows to define distances suitable to deal with privacy and security properties. We provide a *logical characterization* of the *generalized bisimulation metrics* from [12], using the syntactic distance over formulae in the probabilistic extension $\mathcal{L}$ of HML [14], and the notion of *mimicking formula of a process* [9]. The latter is a special formula in $\mathcal{L}$ that captures the observable behavior of a process and allows us to characterize bisimilarity. We show

that the generalized bisimulation distance between two processes is equal to the (generalized) distance between their mimicking formulae, called *logical distance*. Moreover, we show that we can exploit the logical distance to obtain bounds on $d_{\mathcal{X}}$ -privacy. Notice that dealing with bisimulation semantics instead of traces would allow us to develop efficient algorithms for the evaluation of the logical distance (following, e.g., [4]), and thus of approximations on $d_{\mathcal{X}}$ -privacy. Furthermore, we could exploit the *non-expansiveness* results obtained in [12] to favor compositional reasoning over $d_{\mathcal{X}}$ -privacy.

Related work. As already mentioned, this paper builds on the work of [9, 12]. The main novelty is that we develop a technique for characterizing privacy properties, and that we deal with $d_{\mathcal{X}}$ -privacy rather than DP.

Verification of differential privacy has been itself an active area of research. Prominent approaches based on formal methods are those based on type systems [23, 36] and logical formulations [5]. Earlier papers [40, 41] defined bisimulation metrics suitable for proving DP, however they suffered from the fact that the respective kernel relations do not fully characterize probabilistic bisimilarity.

Contribution. Summarizing, the main contributions of this paper are:

1. We define a trace metric over LMCs in terms of a syntactic distance on formulae in $\mathbb{L}$, a probabilistic refinement of HML.
2. We show that such trace metric allows us to obtain a logical characterization of $d_{\mathcal{X}}$ -privacy.
3. We exploit the syntactic distance on formulae to define a real-valued semantics for them, from which we get a logical characterization of weak anonymity.
4. We provide a logical characterization of the generalized bisimilarity metric by using the syntactic distance over $\mathcal{L}$ and the notion of mimicking formulae of processes in a PTS.
5. We exploit the characterization of the bisimilarity metric to obtain bounds on $d_{\mathcal{X}}$ -privacy.

2 Background

The PTS model. PTSs [37] combine LTSs [30] and discrete time Markov chains [27], to model reactive behavior, nondeterminism and probability. The state space is a set $\mathcal{S}$ of *processes*, ranged over by $s, t, \dots$ and transition steps take processes to *probability distributions* over $\mathcal{S}$, namely mappings $\pi: \mathcal{S} \rightarrow [0, 1]$ with $\sum_{s \in \mathcal{S}} \pi(s) = 1$. By $\Delta(\mathcal{S})$ we denote the set of all distributions over $\mathcal{S}$, ranged over by $\pi, \pi', \dots$. The support of $\pi \in \Delta(\mathcal{S})$ is the set $\text{supp}(\pi) = \{s \in \mathcal{S} \mid \pi(s) > 0\}$. We consider only distributions with *finite* support. For $s \in \mathcal{S}$ we denote by δ_s the *Dirac distribution* defined by $\delta_s(s) = 1$ and $\delta_s(t) = 0$ for $s \neq t$.

Definition 1 (PTS, [37]). A nondeterministic probabilistic labeled transition system (PTS) is a triple $(\mathcal{S}, \mathcal{A}, \rightarrow)$, where: $\mathcal{S}$ is a countable set of processes, $\mathcal{A}$ is a countable set of actions, and $\rightarrow \subseteq \mathcal{S} \times \mathcal{A} \times \Delta(\mathcal{S})$ is a transition relation.

We write $s \xrightarrow{a} \pi$ for $(s, a, \pi) \in \rightarrow$, $s \xrightarrow{a}$ if there is a distribution $\pi \in \Delta(\mathcal{S})$ with $s \xrightarrow{a} \pi$, and $s \not\xrightarrow{a}$ otherwise. Let $\text{init}(s) = \{a \in \mathcal{A} \mid s \xrightarrow{a}\}$ denote the set of the actions that can be performed by s . Let $\text{der}(s, a) = \{\pi \in \Delta(\mathcal{S}) \mid s \xrightarrow{a} \pi\}$ denote the set of the distributions reachable from s through action a . We say that a process $s \in \mathcal{S}$ is *image-finite* if for all actions $a \in \text{init}(s)$ the set $\text{der}(s, a)$ is finite [29]. We consider only image-finite processes.

Labeled Markov Chains. We call *trace* any finite sequence of action labels in $\mathcal{A}^*$, ranged over by $\alpha, \alpha', \dots$, and we use ϵ to denote the empty trace.

A *labeled Markov chain* (LMC) is a fully probabilistic PTS, namely a PTS in which for each process we have at most one available transition. In a LMC, a process s induces a probability measure over traces $\text{Pr}(s, \cdot)$, defined for each trace α recursively as follows:

$$\text{Pr}(s, \alpha) = \begin{cases} 1 & \text{if } \alpha = \epsilon \\ 0 & \text{if } \alpha = a\alpha' \text{ and } s \not\xrightarrow{a} \\ \sum_{s' \in \text{supp}(\pi)} \pi(s') \text{Pr}(s', \alpha') & \text{if } \alpha = a\alpha' \text{ and } s \xrightarrow{a} \pi. \end{cases}$$

We can express the *observable behavior* of processes in a LMC in terms of the *linear properties* that they satisfy, or equivalently in terms of the traces that they can perform. Hence, it is natural to compare process behavior in LMCs by means of *trace semantics* (see for instance [4]).

Definition 2 (Trace equivalence on LMCs). Assume a LMC $(\mathcal{S}, \mathcal{A}, \rightarrow)$. Processes $s, t \in \mathcal{S}$ are trace equivalent, written $s \sim_{\text{Tr}} t$, if for all traces $\alpha \in \mathcal{A}^*$ it holds that $\text{Pr}(s, \alpha) = \text{Pr}(t, \alpha)$.

Pseudometric spaces. For a countable set X , a non-negative function $d: X \times X \rightarrow \mathbb{R}^+$ is a *metric* on X whenever it satisfies: (i) $d(x, y) = 0$ iff $x = y$, for all $x, y \in X$; (ii) $d(x, y) = d(y, x)$, for all $x, y \in X$; (iii) $d(x, y) \leq d(x, z) + d(z, y)$, for all $x, y, z \in X$. By relaxing the first axiom to (i)' $d(x, x) = 0$ for all $x \in X$, we obtain the notion of *pseudometric*. We say that d is an *extended* (pseudo)metric if we allow its value to be $+\infty$, notation $d: X \times X \rightarrow [0, +\infty]$. Given a (pseudo)metric d on X , the pair (X, d) is called *(pseudo)metric space*. The *kernel* of a (pseudo)metric d on X is the set $\ker(d) = \{(x, y) \in X \times X \mid d(x, y) = 0\}$. Given two (pseudo)metric spaces $(X, d_X), (Y, d_Y)$, the function $f: X \rightarrow Y$ is *1-Lipschitz* w.r.t. 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 such functions. Given any (pseudo)metric space (X, d) , the *diameter* of X w.r.t. d , denoted by $\mathcal{O}_d(X)$, is the maximal distance of two elements in X , namely $\mathcal{O}_d(X) = \sup_{x, y \in X} d(x, y)$.

The Hausdorff lifting allows us to lift a (pseudo)metric over elements in a set X to a (pseudo)metric over the power set of X , denoted by $\mathcal{P}(X)$.

Definition 3 (Hausdorff metric). Let $d: X \times X \rightarrow [0, +\infty]$ be a pseudometric. The Hausdorff lifting of d is the pseudometric $\mathbf{H}(d): \mathcal{P}(X) \times \mathcal{P}(X) \rightarrow$

$[0, +\infty]$ defined for all sets $X_1, X_2 \subseteq X$ by

$$\mathbf{H}(d)(X_1, X_2) = \max \left\{ \sup_{x_1 \in X_1} \inf_{x_2 \in X_2} d(x_1, x_2), \sup_{x_2 \in X_2} \inf_{x_1 \in X_1} d(x_2, x_1) \right\}$$

with, by convention $\sup_{\emptyset} = 0$ and $\inf_{\emptyset} = \mathcal{O}_d(X)$.

3 Logical characterization of differential privacy: a trace metric approach

In this section we present the first logical characterization for $d_{\mathcal{X}}$ -privacy.

We recall briefly the definitions. The interested reader can find more details in [11]. Let $\mathcal{X}$ be an arbitrary set of *secrets* provided with distance $d_{\mathcal{X}}$. Let $\mathcal{Z}$ be a set of *observables*, and let M be a randomized mechanism from $\mathcal{X}$ to $\mathcal{Z}$, namely a function that assigns to every element of $\mathcal{X}$ a probability distribution on $\mathcal{Z}$. We say that M is $\varepsilon \cdot d_{\mathcal{X}}$ -private if for any two secrets $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ and any measurable subset Z of $\mathcal{Z}$, we have $M(\mathbf{x})(Z)/M(\mathbf{x}')(Z) \leq e^{\varepsilon \cdot d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}')}$. The idea is that $d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}')$ represents a *distinguishability level* between $\mathbf{x}$ and $\mathbf{x}'$: the more we want to confuse them, the more similar the probabilities of producing the same answers in the randomization process should be. Notice that $d_{\mathcal{X}}$ -privacy subsumes standard DP, by setting $\mathcal{X}$ to be the set of databases, and $d_{\mathcal{X}}$ the Hamming distance between databases: $d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}')$ is the number of records in which $\mathbf{x}$ and $\mathbf{x}'$ differ. The resulting property is, by transitivity, equivalent to say that for all $\mathbf{x}$ and $\mathbf{x}'$ which are adjacent (i.e., $d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}') = 1$), $M(\mathbf{x})(Z)/M(\mathbf{x}')(Z) \leq e^{\varepsilon}$. Note that we consider here an equivalent definition of DP in which the adjacency relation is defined as differing in the value of one record. The standard definition, in which $\mathbf{x}$ and $\mathbf{x}'$ are adjacent if $\mathbf{x}'$ is obtained from $\mathbf{x}$ by adding or removing one record, can be specified by using an extra value to indicate absence of the record.

Furthermore, $d_{\mathcal{X}}$ -privacy subsumes LDP as well, by setting $d_{\mathcal{X}}$ to be the discrete distance, i.e., $d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}') = 0$ if $\mathbf{x} = \mathbf{x}'$ and $d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}') = 1$ otherwise.

To formalize $d_{\mathcal{X}}$ -privacy, we will exploit the *multiplicative variant of the total variation distance* on probability distributions.

Definition 4 (Multiplicative total variation distance, [38]). Let X be a set. The multiplicative variant of the total variation distance on $\Delta(X)$ is the function $tv_{\otimes}: \Delta(X) \times \Delta(X) \rightarrow [0, +\infty]$ defined, for all $\pi, \pi' \in \Delta(X)$, as $tv_{\otimes}(\pi, \pi') = \sup_{x \in X} |\ln(\pi(x)) - \ln(\pi'(x))|$.

For $\mathcal{X}$ set of secrets and $\mathcal{Z}$ set of observables, $d_{\mathcal{X}}$ -privacy is defined as follows.

Definition 5 ($d_{\mathcal{X}}$ -privacy, [11]). Let $\varepsilon > 0$ and $d_{\mathcal{X}}$ be any distance on $\mathcal{X}$. A randomized mechanism $M: \mathcal{X} \rightarrow \Delta(\mathcal{Z})$ is $\varepsilon \cdot d_{\mathcal{X}}$ -private iff

$$tv_{\otimes}(M(\mathbf{x}), M(\mathbf{x}')) \leq \varepsilon \cdot d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}') \quad \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}.$$

Interestingly, each randomized mechanisms can be modeled as a LMC. Each secret $\mathbf{x}$ is mapped to a state $s_{\mathbf{x}}$ in the LMC and the observable result of the

Fig. 1. The mechanism ‘Randomized responses’ as a LMC. For simplicity, an arrow $u \xrightarrow{a}$ with no target models the evolution of process u to the Dirac distribution δ_{nil} , with nil process that can execute no action, via the execution of a .

mechanism applied to $\mathbf{x}$ is modeled by the traces executable by $s_{\mathbf{x}}$ in the LMC. The randomized mechanism M on $\mathbf{x}$ is then modeled as the trace distribution induced by $s_{\mathbf{x}}$. More formally, we consider $\mathcal{Z} = \mathcal{A}^*$ and we define $M(\mathbf{x})(\alpha) = \Pr(s_{\mathbf{x}}, \alpha)$ for each $\alpha \in \mathcal{A}^*$.

We give an example based on local differential privacy. The mechanism is called “Randomized responses” and is a simplified instance of the system RAP-POR used by Google to protect the privacy of their users [22].

Example 1 (Randomized responses). Suppose that we want to collect the answers to some embarrassing question (for instance “Have you ever cheated on your partner?”) for some statistic purpose. To persuade people to answer truly, we allow them to report the true answer with probability $3/4$, and the opposite answer with probability $1/4$. In this way, the privacy of the user will be protected in the sense that the answers collector will not know for sure whether the person has cheated or not. In fact, the system is $\log 3$ -locally differentially private. At the same time, if the population is large enough, the collector will be able to obtain a good statistical approximation of the real percentage of cheaters.

To implement the system, we can use a (fair) coin: the person tosses the coin twice, and if the first result is head, he answers truly, otherwise he answers “yes” or “no” depending on whether the second result is, resp., head or tail. The results of the coin tossings, of course, has to be invisible to the data collector, and thus we represent it as an internal action τ .

The LMCs s_y and s_n in Fig. 1 represent the mechanism applied to two individuals: s_y that has cheated and s_n has not. s_y will toss the coin and make a transition τ . Then, depending on the result, will go in a state s_h or s_t with even probability. From s_h it will toss a coin again, and then make a transition yes to a final state. From s_t it will toss the coin and go in states s_{th} and s_{tt} with even probability. From s_{th} and s_{tt} it will then make transitions yes and no , resp., and then terminate. The system s_n is analogous, with yes and no inverted.

To obtain the logical characterization of $d_{\mathcal{X}}$ -privacy we can investigate the semantics of the so obtained LMCs. In particular, we will exploit a notion of *trace*

metric evaluated on *modal formulae* expressing linear properties of processes in LMCs. Informally, we will consider a simple probabilistic variant of the modal logic capturing the trace semantics in the fully nondeterministic case to define a *probabilistic trace semantics* for processes. Then, we will define a metric for such a semantics in terms of a syntactic distance over the formulae in the considered logic and we will use such a distance to characterize $d_{\mathcal{X}}$ -privacy. Interestingly, although the considered trace semantics is based on a quite limited observation power, it will allow us to obtain the first *logical characterization* of $d_{\mathcal{X}}$ -privacy (Thm. 2): we will show that the trace metrics so defined on LMCs coincides with the multiplicative variant of the total variation distance (Prop. 2).

3.1 Trace metrics on LMCs

Probabilistic trace semantics compares the behavior of processes w.r.t. the probabilities that they assign to the same linear properties, namely to the same traces. In the literature we can find several notions of probabilistic trace equivalence, of which $\sim_{\text{Tr}}$ given in Def. 2 is an example, and we refer the interested reader to [6] for a survey. Such a wealth of notions derives from the interplay of non-determinism and probability that we can witness in quantitative systems and the different interpretations that researchers have given to it. We can also find several proposals of behavioral distances measuring the disparities of processes w.r.t. the same linear properties, that is their differences in the probabilities of executing the same traces (see, e.g., [1, 4, 39]).

As the focus of this paper is on $d_{\mathcal{X}}$ -privacy, we adopt a different approach, w.r.t. to those referenced, to the definition of a trace metric on LMCs. In fact, we hark back to the seminal work [17] on bisimulation metrics and: (i) We provide a logical characterization of $\sim_{\text{Tr}}$ by means of a simple modal logic $\mathbb{L}$ that allows us to express traces and their probability of being executed, so that s and t are trace equivalent if they satisfy the same formulae in $\mathbb{L}$. (ii) We quantify the trace metric on processes in terms to the formulae distinguishing them. Informally, in [17] this is obtained by transforming formulae into functional expressions and by interpreting the satisfaction relation as integration. Then, the distance on processes is defined on the so obtained *real-valued* logic by considering the maximal disparity between the images of processes through all functional expressions. Here, we propose a much simpler approach based on the *boolean-valued* logic $\mathbb{L}$: we introduce a (family of generalized) *syntactic distance* on formulae in $\mathbb{L}$ and we define the *trace metric* on processes as the Hausdorff lifting of the syntactic distance to the sets of formulae satisfied by processes.

The logic $\mathbb{L}$ extends the one used in the nondeterministic case to express trace semantics [7] (and corresponding to the subclass of *linear formulae*) with a probabilistic modality expressing the execution probabilities of traces.

Definition 6 (Modal logic $\mathbb{L}$). *The logic $\mathbb{L} = \mathbb{L}^1 \cup \mathbb{L}^P$ is given by the classes of linear formulae $\mathbb{L}^1$ and of probabilistic formulae $\mathbb{L}^P$ over $\mathcal{A}$, defined by:*

$$\mathbb{L}^1: \quad \Phi ::= \quad \top \mid \langle a \rangle \Phi \qquad \mathbb{L}^P: \quad \Psi ::= \quad r\Phi$$

where: (i) Φ ranges over $\mathbb{L}^1$, (ii) Ψ ranges over $\mathbb{L}^p$, (iii) $a \in \mathcal{A}$; (iv) $r \in [0, 1]$.

We say that a trace α is compatible with the linear formula Φ , notation $\alpha = \text{Tr}(\Phi)$, if the sequence of action labels in α is exactly the same sequence of labels of the diamond modalities in Φ , i.e., $\alpha = \text{Tr}(\langle a_1 \rangle \dots \langle a_n \rangle \top)$ iff $\alpha = a_1 \dots a_n$.

Definition 7 (Semantics of $\mathbb{L}$). For any $s \in \mathcal{S}$, the satisfaction relation $\models \subseteq \mathcal{S} \times \mathbb{L}^1 \cup \mathbb{L}^p$ is defined by structural induction over formulae in $\mathbb{L}^1 \cup \mathbb{L}^p$ by

- $s \models \top$ always;
- $s \models \langle a \rangle \Phi$ iff $s \xrightarrow{a} \pi$ for some π such that $s' \models \Phi$ for some $s' \in \text{supp}(\pi)$;
- $s \models r\Phi$ iff $s \models \Phi$ and $\text{Pr}(s, \text{Tr}(\Phi)) = r$.

For each process $s \in \mathcal{S}$, we let $\mathbb{L}(s) = \{\Psi \in \mathbb{L}^p \mid s \models \Psi\}$.

Example 2 (Randomized responses II). Consider processes s_y, s_n in Fig. 1. One can easily check that

$$\begin{aligned} \mathbb{L}(s_y) &= \{1\langle\tau\rangle\top, 1\langle\tau\rangle\langle\tau\rangle\top, 3/4\langle\tau\rangle\langle\tau\rangle\langle\text{yes}\rangle\top, 1/4\langle\tau\rangle\langle\tau\rangle\langle\text{no}\rangle\top\} \\ \mathbb{L}(s_n) &= \{1\langle\tau\rangle\top, 1\langle\tau\rangle\langle\tau\rangle\top, 1/4\langle\tau\rangle\langle\tau\rangle\langle\text{yes}\rangle\top, 3/4\langle\tau\rangle\langle\tau\rangle\langle\text{no}\rangle\top\} \end{aligned}$$

By means of $\mathbb{L}$ we can provide a logical characterization of $\sim_{\text{Tr}}$: two processes are trace equivalent if and only if they satisfy the same formulae in $\mathbb{L}$.

Theorem 1. Assume an LMC $(\mathcal{S}, \mathcal{A}, \rightarrow)$. Then for all processes $s, t \in \mathcal{S}$ we have that $s \sim_{\text{Tr}} t$ iff $\mathbb{L}(s) = \mathbb{L}(t)$.

We can now proceed to the definition of the *trace metric*. The definition of the *syntactic distance* on formulae in $\mathbb{L}$ is parametric w.r.t. a generic metric $\mathcal{D}$ on $[0, 1]$ that plays the role of a ground distance on the weights of probabilistic formulae, to which a syntactic distance could not be applied. For this reason we shall sometimes speak of *generalized syntactic distance* and *trace metric*.

Definition 8 (Distance on $\mathbb{L}$). Let $([0, 1], \mathcal{D})$ be a metric space. The function $\mathbf{dm}_{\mathcal{D}}: \mathbb{L}^1 \times \mathbb{L}^1 \rightarrow \{0, \mathcal{D}([0, 1])\}$ is defined as the discrete metric over $\mathbb{L}^1$, namely $\mathbf{dm}_{\mathcal{D}}(\Phi_1, \Phi_2) = 0$ if $\Phi_1 = \Phi_2$ and $\mathbf{dm}_{\mathcal{D}}(\Phi_1, \Phi_2) = \mathcal{D}([0, 1])$ otherwise. The function $\mathbf{d}_{\mathcal{D}}^p: \mathbb{L}^p \times \mathbb{L}^p \rightarrow [0, \mathcal{D}([0, 1])]$ is defined over $\mathbb{L}^p$ as follows:

$$\mathbf{d}_{\mathcal{D}}^p(r_1\Phi_1, r_2\Phi_2) = \begin{cases} \mathcal{D}(r_1, r_2) & \text{if } \mathbf{dm}_{\mathcal{D}}(\Phi_1, \Phi_2) = 0 \\ \mathcal{D}([0, 1]) & \text{otherwise.} \end{cases}$$

Definition 9 (Trace metric). Let $([0, 1], \mathcal{D})$ be a metric space. The trace metric over processes $\mathbf{d}_{\mathcal{D}}^T: \mathcal{S} \times \mathcal{S} \rightarrow [0, \mathcal{D}([0, 1])]$ is defined for all $s, t \in \mathcal{S}$ by

$$\mathbf{d}_{\mathcal{D}}^T(s, t) = \mathbf{H}(\mathbf{d}_{\mathcal{D}}^p)(\mathbb{L}(s), \mathbb{L}(t)).$$

The kernel of each generalized trace metric corresponds to $\sim_{\text{Tr}}$.

Proposition 1. For all possible choices of the metric $\mathcal{D}$, trace equivalence is the kernel of the trace metric, namely $\sim_{\text{Tr}} = \ker(\mathbf{d}_{\mathcal{D}}^T)$.

3.2 Logical characterization of $d_{\mathcal{X}}$ -privacy

We can now present the logical characterization result for $d_{\mathcal{X}}$ -privacy. As the $d_{\mathcal{X}}$ -privacy property is basically a measure of the level of privacy of a system, a logical characterization for it should be interpreted as a logical characterization of a behavioral metric, in the sense of [9, 10, 17], rather than in the sense of behavioral equivalences. Roughly speaking, we evaluate the $d_{\mathcal{X}}$ -privacy property by exploiting the linear properties of the mechanism as expressed by our trace metric, and thus by the logic $\mathbb{L}$. More formally, we let $\mathbf{d}_{\otimes}^T$ denote the multiplicative variant of our trace metric, i.e., the one with $\mathcal{D}(r_1, r_2) = |\ln(r_1) - \ln(r_2)|$. Then, we prove that $\mathbf{d}_{\otimes}^T$ coincides with the multiplicative total variation distance on the trace distributions induced by processes.

Proposition 2. *For any $s \in \mathcal{S}$ let $\mu_s = \Pr(s, \cdot)$. Then $\mathbf{d}_{\otimes}^T(s, t) = tv_{\otimes}(\mu_s, \mu_t)$.*

We can then formalize our logical characterization of $d_{\mathcal{X}}$ -privacy.

Theorem 2 (Logical characterization of $d_{\mathcal{X}}$ -privacy). *Let M be a randomized mechanism defined by $M(\mathbf{x})(\alpha) = \Pr(s_{\mathbf{x}}, \alpha)$ for all $\mathbf{x} \in \mathcal{X}$, $\alpha \in \mathcal{A}^*$. Then, given $\varepsilon > 0$, M is $\varepsilon \cdot d_{\mathcal{X}}$ -private if $\mathbf{d}_{\otimes}^T(s_{\mathbf{x}}, s_{\mathbf{x}'}) \leq \varepsilon \cdot d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}') \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$.*

Example 3 (Randomized responses, III). We can show that the mechanism ‘Randomized responses’ described in Ex. 1 is log 3-locally differentially private by evaluating the trace distance between processes s_y and s_n in Fig. 1. By comparing the sets of formulae $\mathbb{L}(s_y)$ and $\mathbb{L}(s_n)$ given in Ex. 2, we can infer that

$$\begin{aligned} \mathbf{d}_{\otimes}^T(s_y, s_n) &= \max \left\{ \begin{array}{l} \mathfrak{d}_{\otimes}^P(3/4 \langle \tau \rangle \langle \tau \rangle \langle yes \rangle \top, 1/4 \langle \tau \rangle \langle \tau \rangle \langle yes \rangle \top) \\ \mathfrak{d}_{\otimes}^P(1/4 \langle \tau \rangle \langle \tau \rangle \langle no \rangle \top, 3/4 \langle \tau \rangle \langle \tau \rangle \langle no \rangle \top) \end{array} \right\} \\ &= |\ln(3/4) - \ln(1/4)| = \ln(3). \end{aligned}$$

3.3 Logical characterization of weak anonymity: from boolean to real semantics

So far, we have seen how we can express the $d_{\mathcal{X}}$ -privacy property as a syntactic distance over modal formulae capturing trace semantics. However, in the literature, when behavioral metrics are considered, logics equipped with a real-valued semantics are usually used for the characterization, which is then expressed as

$$d(s, t) = \sup_{\phi \in L} |\llbracket \phi \rrbracket(s) - \llbracket \phi \rrbracket(t)| \quad (1)$$

where d is the behavioral metric of interest, L is the considered logic and $\llbracket \phi \rrbracket(s)$ denotes the value of the formula ϕ in process s accordingly to the real-valued semantics (see eg. [1, 3, 17–19]). In this Section, we exploit the syntactic distance on $\mathbb{L}$ to provide a real valued semantics for formulae and thus a characterization of *weak probabilistic anonymity* expressed accordingly to classic schema in (1).

Weak probabilistic anonymity [16] uses the *additive* total variation distance tv to measure the degree of protection of the identity of a user while performing

a particular task. Hence, the set of secrets $\mathcal{X}$ is now the set of users' identities and a randomized mechanism $M: \mathcal{X} \rightarrow \Delta(\mathcal{Z})$ has to introduce some noise so that from the 'performed tasks' in $\mathcal{Z}$ an adversary cannot discover the identity of the user that actually performed them. Finally, we recall that the total variation distance is defined by $tv(\mu, \mu') = \sup_{Z \in \mathcal{Z}} |\mu(Z) - \mu'(Z)|$ for all $\mu, \mu' \in \Delta(\mathcal{Z})$.

Definition 10 (Weak probabilistic anonymity [16]). *Let M be a randomized mechanism defined by $M(\mathbf{x})(\alpha) = \Pr(s_{\mathbf{x}}, \alpha)$ for all $\mathbf{x} \in \mathcal{X}$, $\alpha \in \mathcal{A}^*$. Then, given $\varepsilon > 0$, M satisfies ε -weak anonymity if $tv(M(\mathbf{x}), M(\mathbf{x}')) \leq \varepsilon \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$.*

So, we consider all metric spaces $([0, 1], \mathcal{D})$ with $\mathcal{O}_{\mathcal{D}}([0, 1]) < \infty$ and: 1. We use the syntactic distance over formulae in $\mathbb{L}$ to define a (generalized) real valued semantics for those modal formulae. 2. We show that the total variation distance satisfies the general schema in (1) w.r.t. such real semantics. 3. We express the ε -weak anonymity property as an upper bound to the total variation distance on the values of formulae in the processes of the LMCs.

Equipping modal formulae with a real-valued semantics means assigning to each formula ϕ a real number in $[0, 1]$ expressing *how much* a given process s satisfies ϕ ; value 1 stands for $s \models \phi$. We exploit our distance over formulae to define such a semantics. Informally, let L be the class of formulae of interest, let $D_{\mathcal{D}}$ be any generalized syntactic distance defined on L (like, eg., the distance $\mathfrak{d}_{\mathcal{D}}^p$ for the logic $\mathbb{L}$) and for each process s let $L(s)$ denote the set of formulae in L satisfied by s . To quantify how much the formula $\phi \in L$ is satisfied by process s , we evaluate first how far ϕ is from being satisfied by s . This corresponds to the minimal distance between ϕ and a formula satisfied by s , namely to $\inf_{\phi' \in L(s)} D_{\mathcal{D}}(\phi, \phi')$. Then we simply notice that being $D_{\mathcal{D}}(\phi, \phi')$ far from s is equivalent to be $\mathcal{O}_{\mathcal{D}}([0, 1]) - D_{\mathcal{D}}(\phi, \phi')$ close to it (notice that $\mathcal{O}_{\mathcal{D}}([0, 1])$ has to be finite in order to obtain a meaningful value). Thus, we assign to ϕ the value $\mathcal{O}_{\mathcal{D}}([0, 1]) - \inf_{\phi' \in L(s)} D_{\mathcal{D}}(\phi, \phi') / \mathcal{O}_{\mathcal{D}}([0, 1])$ in s , where the normalization w.r.t. $\mathcal{O}_{\mathcal{D}}([0, 1])$ ensures that this value is in $[0, 1]$.

Definition 11 (Real valued semantics). *Let $([0, 1], \mathcal{D})$ be a metric space with $\mathcal{O}_{\mathcal{D}}([0, 1]) < \infty$. Assume any class of formulae L , let $D_{\mathcal{D}}$ be any generalized syntactic distance over L . We define the value of $\phi \in L$ in process $s \in \mathcal{S}$ as*

$$\llbracket \phi \rrbracket_{\mathcal{D}}(s) = 1 - \frac{\inf_{\phi' \in L(s)} D_{\mathcal{D}}(\phi, \phi')}{\mathcal{O}_{\mathcal{D}}([0, 1])}$$

Example 4. Consider s_y in Fig. 1 and $\mathcal{D}(r_1, r_2) = |r_1 - r_2|$ for all $r_1, r_2 \in [0, 1]$. Notice that $\mathcal{O}_{\mathcal{D}}([0, 1]) = 1$. For any $r \in [0, 1]$, consider the formula $\varphi_r = r\langle\tau\rangle\langle\tau\rangle\langle yes \rangle \top$. Then $\inf_{\varphi \in L(s_y)} \mathfrak{d}_{\mathcal{D}}^p(\varphi_r, \varphi) = \mathfrak{d}_{\mathcal{D}}^p(\varphi_r, {}^{3/4}\langle\tau\rangle\langle\tau\rangle\langle yes \rangle \top) = |r - 3/4|$. Hence, the value of φ_r in s_y is given by $\llbracket \varphi_r \rrbracket_{\mathcal{D}}(s_y) = 1 - |r - 3/4|$.

Before proceeding to the characterization, notice that for each class of formulae L equipped with a generalized syntactical distance $D_{\mathcal{D}}$ we can provide an equivalent reformulation of the Hausdorff metric as in the following Proposition.

Proposition 3. *Let $([0, 1], \mathcal{D})$ be a metric space. Assume a class of formulae L and let $D_{\mathcal{D}}$ be any generalized syntactic distance over L . For any $L_1, L_2 \subseteq L$ we have that $\mathbf{H}(D_{\mathcal{D}})(L_1, L_2) = \sup_{\phi \in L} |\inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2)|$.*

If we focus on the class of formulae $\mathbb{L}$, from Prop. 3 we can immediately derive the characterization of trace metrics in terms of real-valued formulae.

Lemma 1. *Let $([0, 1], \mathcal{D})$ be a metric space with $\mathcal{O}_{\mathcal{D}}([0, 1]) < \infty$. For all processes $s, t \in \mathcal{S}$ it holds that $\mathbf{d}_{\mathcal{D}}^T(s, t) = \sup_{\Psi \in \mathbb{L}^1} |\llbracket \Psi \rrbracket_{\mathcal{D}}(s) - \llbracket \Psi \rrbracket_{\mathcal{D}}(t)|$.*

By abuse of notation, for any linear formula $\Phi \in \mathbb{L}^1$, we write $\llbracket \Phi \rrbracket_{\mathcal{D}}(s)$ in place of $\llbracket 1\Phi \rrbracket_{\mathcal{D}}(s)$. Moreover, we write the ‘generalized’ metrics defined on the metric space $([0, 1], \mathcal{D})$, with $\mathcal{D}(x, y) = |x - y|$, with no $\mathcal{D}$ subscripts. Then, the following characterization of the total variation distance holds.

Proposition 4. *Let $([0, 1], \mathcal{D})$ be a metric space with $\mathcal{O}_{\mathcal{D}}([0, 1]) < \infty$. For any $s \in \mathcal{S}$ define $\mu_s = \Pr(s, \cdot)$. Then, $tv_{\mathcal{D}}(\mu_s, \mu_t) = \sup_{\Phi \in \mathbb{L}^1} |\llbracket \Phi \rrbracket_{\mathcal{D}}(s) - \llbracket \Phi \rrbracket_{\mathcal{D}}(t)|$. In particular, we have $tv(\mu_s, \mu_t) = \sup_{\Phi \in \mathbb{L}^1} |\llbracket \Phi \rrbracket(s) - \llbracket \Phi \rrbracket(t)|$.*

Finally, we can express ε -weak anonymity property as an upper bound to the total variation distance on the values of formulae in the processes of the LMCs, accordingly to the general schema in (1).

Theorem 3 (Logical characterization of weak anonymity). *Let M be a randomized mechanism defined by $M(\mathbf{x})(\alpha) = \Pr(s_{\mathbf{x}}, \alpha)$ for all $\mathbf{x} \in \mathcal{X}$, $\alpha \in \mathcal{A}^*$. Then, given $\varepsilon > 0$, M satisfies ε -weak anonymity if $\sup_{\Phi \in \mathbb{L}^1} |\llbracket \Phi \rrbracket(s_{\mathbf{x}}) - \llbracket \Phi \rrbracket(s_{\mathbf{x}'})| \leq \varepsilon \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$.*

4 A logical bound on $d_{\mathcal{X}}$ -privacy: from traces to bisimulations

So far we have shown how it is possible to obtain a characterization of $d_{\mathcal{X}}$ -privacy by exploiting trace semantics and a notion of syntactic distance on modal formulae. However, one could argue that there are no efficient algorithms to evaluate the trace metric, and therefore the $d_{\mathcal{X}}$ -privacy property, especially if the state space of the LMC is infinite. In [4] it is proved that we can obtain upper bounds on the evaluation of trace metrics by exploiting bisimulation-like distances, for which polynomial-time algorithms can be provided. Here, we follow a similar reasoning: we switch from LMCs to the more general semantic model of *PTSSs*, we consider the *generalized bisimulation metrics* introduced in [12] and we provide a *logical characterization* for them. This is based on the notion of syntactic distance on formulae and the notion of *mimicking formula* of a process from [9, 10]. As in previous Sect. 3.1, the former is a pseudometric on a probabilistic version of HML $\mathcal{L}$ that extends $\mathbb{L}$ with modalities allowing us to express the interplay of nondeterminism and probability typical of PTSSs (Sect. 4.2). The latter is a special formula in $\mathcal{L}$ that alone expresses the observable behavior w.r.t. bisimulation semantics of the process to which it is related and

allows us to characterize bisimilarity (Sect. 4.3). Then we show that the syntactic distance between the mimicking formulae of processes equals their bisimulation distance (Sect. 4.4) and that, when we focus on LMCs, it gives an upper bound on $d_{\mathcal{X}}$ -privacy properties of mechanisms (Sect. 4.5).

As a final remark, note that using bisimulation metrics and their characterization would allow us to apply the compositional results obtained for them in [12] also to $d_{\mathcal{X}}$ -privacy properties. Due to space limitations, we leave their formal development as future work. Now, we proceed to recall some base notions on bisimulation semantics and generalized bisimulation metrics.

4.1 Generalized bisimilarity metric

Probabilistic (bi)simulations. A probabilistic bisimulation is an equivalence relation over $\mathcal{S}$ that equates processes $s, t \in \mathcal{S}$ if they can mimic each other's transitions and evolve to distributions that are in turn related by the same bisimulation. To formalize this, we need to lift relations over processes to relations over distributions. Informally, given a relation $\mathcal{R}$ on processes we say that two distributions $\pi, \pi' \in \Delta(\mathcal{S})$ are related by the lifting of $\mathcal{R}$, denoted by $\mathcal{R}^\dagger$, iff they assign the same probabilistic weights to the same equivalence classes in $\mathcal{R}$.

Definition 12 (Relation lifting, [15]). *Let X be a set. The lifting of a relation $\mathcal{R} \subseteq X \times X$ is the relation $\mathcal{R}^\dagger \subseteq \Delta(X) \times \Delta(X)$ with $\pi \mathcal{R}^\dagger \pi'$ whenever there is a set of indexes I s.t.*

$$(i) \pi = \sum_{i \in I} p_i \delta_{x_i}, \quad (ii) \pi' = \sum_{i \in I} p_i \delta_{y_i}, \quad \text{and} \quad (iii) x_i \mathcal{R} y_i \text{ for all } i \in I.$$

Definition 13 (Probabilistic bisimulation, [33]). *Assume a PTS. A binary relation $\mathcal{R} \subseteq \mathcal{S} \times \mathcal{S}$ is a probabilistic bisimulation if whenever $s \mathcal{R} t$*

- *if $s \xrightarrow{a} \pi_s$ then there is a transition $t \xrightarrow{a} \pi_t$ such that $\pi_s \mathcal{R}^\dagger \pi_t$;*
- *if $t \xrightarrow{a} \pi_t$ then there is a transition $s \xrightarrow{a} \pi_s$ such that $\pi_t \mathcal{R}^\dagger \pi_s$;*

The union of all probabilistic bisimulations is the greatest probabilistic bisimulation, denoted by $\sim$ and called bisimilarity, and is an equivalence.

Generalized bisimulation metrics. For our purposes, we need to lift a pseudometric over processes to a pseudometric over distributions over processes. We follow the approach of [12], that considers the *generalized Kantorovich lifting*. Take a generic metric space (V, d_V) , with $V \subseteq \mathbb{R}$ is a convex subset of the reals. A function $f: X \rightarrow V$ can be lifted to a function $\hat{f}: \Delta(X) \rightarrow V$ by taking its expected value, i.e., $\hat{f}(\pi) = \sum_{x \in X} \pi(x) f(x)$ (requiring V to be convex ensures that $\hat{f}(\pi) \in V$). Then, for each V , we define the lifting of a pseudometric d_X over X to a pseudometric over $\Delta(X)$ via the *generalized Kantorovich metric* $\mathbf{K}_V$.

Definition 14 (Generalized Kantorovich lifting, [12]). *For a pseudometric space (X, d_X) and a metric space (V, d_V) with $V \subseteq \mathbb{R}$ convex, the generalized Kantorovich lifting of d_X w.r.t. (V, d_V) is the pseudometric $\mathbf{K}_V: \Delta(X) \times \Delta(X) \rightarrow [0, +\infty]$ defined, for all $\pi, \pi' \in \Delta(X)$ by*

$$\mathbf{K}_V(d_X)(\pi, \pi') = \sup \left\{ d_V(\hat{f}(\pi), \hat{f}(\pi')) \mid f \in 1\text{-Lip}[(X, d_X), (V, d_V)] \right\}.$$

Bisimulations answer the question of whether two processes behave precisely the same way or not. Bisimulation metrics answer the more general question of how far the behavior of two processes is. They are defined as the least fixed points of a suitable functional on the following structure. Let (V, d_V) be a metric space and let $\mathbf{D}$ be the set of pseudometrics d on $\mathcal{S}$ such that $\mathcal{O}_d(\mathcal{S}) \leq \mathcal{O}_{d_V}(V)$. Then $(\mathbf{D}, \preceq)$ with $d_1 \preceq d_2$ iff $d_1(s, t) \leq d_2(s, t)$ for all processes $s, t \in \mathcal{S}$, is a complete lattice. In detail, for each set $D \subseteq \mathbf{D}$ the supremum and infimum are $\sup(D)(s, t) = \sup_{d \in D} d(s, t)$ and $\inf(D)(s, t) = \inf_{d \in D} d(s, t)$ for all $s, t \in \mathcal{S}$. The bottom element is function $\mathbf{0}$ with $\mathbf{0}(s, t) = 0$ for all $s, t \in \mathcal{S}$, and the top element is function $\mathbf{1}$ with $\mathbf{1}(s, t) = \mathcal{O}_{d_V}(V)$ if $s \neq t$, and $\mathbf{1}(s, t) = 0$ otherwise.

The quantitative analogue of bisimulation is defined by means of a functional $\mathbf{B}_V$ over the lattice $(\mathbf{D}, \preceq)$. By means of a *discount factor* $\lambda \in (0, 1]$, $\mathbf{B}_V$ allows us to specify how much the behavioral distance of future transitions is taken into account to determine the distance between two processes [2, 17]. $\lambda = 1$ expresses no discount, so that the differences in the behavior of s and t are considered irrespective of after how many steps can be observed.

Definition 15 (Generalized bisimulation metric functional, [12]). *Let (V, d_V) be a metric space, with $V \subseteq \mathbb{R}$ convex. Let $\mathbf{B}_V: \mathbf{D} \rightarrow \mathbf{D}$ be the function defined for all $d \in \mathbf{D}$ and $s, t \in \mathcal{S}$ by*

$$\mathbf{B}_V(d)(s, t) = \sup_{a \in \mathcal{A}} \mathbf{H}(\lambda \cdot \mathbf{K}_V(d))(\text{der}(s, a), \text{der}(t, a)).$$

Remark 1. It is easy to show that for any pseudometric d the lifting $\mathbf{K}_V(d)$ is an extended pseudometric for any choice of (V, d_V) . However, in general the lifting does not preserve the boundedness properties of d . To guarantee $\mathbf{K}_V(d)$ to be bounded we need to assume that the metric d_V is *ball-convex*, namely the open balls in the generated topology are convex sets. This is not an issue for this paper, since all the considered metrics satisfy the ball-convex property. Thus, henceforth, whenever we consider a metric space (V, d_V) with $V \subseteq \mathbb{R}$ convex, we subsume also the ball-convex property for the metric d_V .

We can show that $\mathbf{B}_V$ is monotone [12]. Then, as $(\mathbf{D}, \preceq)$ is a complete lattice, by the Tarski theorem $\mathbf{B}_V$ has the least fixed point. Bisimulation metrics are the pseudometrics being prefixed points of $\mathbf{B}_V$ and the *bisimilarity metric* $\mathbf{d}_{\lambda, V}$ is the least fixed point of $\mathbf{B}_V$ and its kernel is probabilistic bisimilarity [12].

Definition 16 (Generalized bisimulation metric, [12]). *A pseudometric $d: \mathcal{S} \times \mathcal{S} \rightarrow [0, +\infty]$ is a bisimulation metric iff $\mathbf{B}_V(d) \preceq d$. The least fixed point of $\mathbf{B}_V$ is denoted by $\mathbf{d}_{\lambda, V}$ and called the bisimilarity metric.*

4.2 The modal logic $\mathcal{L}$

We introduce the *modal logic* $\mathcal{L}$ of [14], which extends HML [28] with a probabilistic choice modality that allows us to express the behavior of probability distributions over processes.

Fig. 2. The classic bisimilarity distance between s, t is $\mathbf{d}_\lambda(s, t) = 1/2 \cdot \lambda$.

Definition 17 (Modal logic $\mathcal{L}$, [14]). The logic $\mathcal{L} = \mathcal{L}^s \cup \mathcal{L}^d$ is given by the classes of state formulae $\mathcal{L}^s$ and distribution formulae $\mathcal{L}^d$ over $\mathcal{A}$ defined by:

$$\mathcal{L}^s: \varphi ::= \top \mid \neg\varphi \mid \bigwedge_{j \in J} \varphi_j \mid \langle a \rangle \psi \quad \mathcal{L}^d: \psi ::= \bigoplus_{i \in I} r_i \varphi_i$$

where: (i) φ ranges over $\mathcal{L}^s$, (ii) ψ ranges over $\mathcal{L}^d$, (iii) $a \in \mathcal{A}$, (iv) $J \neq \emptyset$ is a countable set of indexes, (v) $I \neq \emptyset$ is a finite set of indexes and (vi) $r_i \in (0, 1]$ for all $i \in I$ and $\sum_{i \in I} r_i = 1$.

We shall write $\varphi_1 \wedge \varphi_2$ for $\bigwedge_{j \in J} \varphi_j$ with $J = \{1, 2\}$, and $\langle a \rangle \varphi$ for $\langle a \rangle \bigoplus_{i \in I} r_i \varphi_i$ with $I = \{i\}$, $r_i = 1$ and $\varphi_i = \varphi$. We use $\top$ instead of $\bigwedge_\emptyset$ to improve readability.

Formulae are interpreted over a PTS. A distribution π satisfies the formula $\bigoplus_{i \in I} r_i \varphi_i$ if, for each $i \in I$, π assigns probability (at least) r_i to processes satisfying the formula φ_i . This is formalized by requiring that π can be rewritten as a convex combination of distributions π_i , using the r_i as weights, such that all the processes in $\text{supp}(\pi_i)$ satisfy the formula φ_i .

Definition 18 (Semantics of $\mathcal{L}$, [14]). The satisfaction relation $\models \subseteq (\mathcal{S} \times \mathcal{L}^s) \cup (\Delta(\mathcal{S}) \times \mathcal{L}^d)$ is defined by structural induction on formulae in $\mathcal{L}$ by

- $s \models \top$ always;
- $s \models \neg\varphi$ iff $s \models \varphi$ does not hold;
- $s \models \bigwedge_{j \in J} \varphi_j$ iff $s \models \varphi_j$ for all $j \in J$;
- $s \models \langle a \rangle \psi$ iff $s \xrightarrow{a} \pi$ for a distribution $\pi \in \Delta(\mathcal{S})$ with $\pi \models \psi$,
- $\pi \models \bigoplus_{i \in I} r_i \varphi_i$ iff $\pi = \sum_{i \in I} r_i \pi_i$ for some distributions $\pi_i \in \Delta(\mathcal{S})$ such that for all $i \in I$ we have $s \models \varphi_i$ for all states $s \in \text{supp}(\pi_i)$.

We introduce the relation of $\mathcal{L}$ -equivalence over formulae in $\mathcal{L}$, which identifies formulae that are indistinguishable by their syntactic structure. Such an equivalence is obtained as the greatest fixed point of a proper transformation E of relations on state formulae.

Definition 19 ([9]). Let $\mathcal{R} \subseteq \mathcal{L}^s \times \mathcal{L}^s$ be any equivalence relation on $\mathcal{L}^s$. The transformation $E: \mathcal{L}^s \times \mathcal{L}^s \rightarrow \mathcal{L}^s \times \mathcal{L}^s$ is defined as: $(\varphi, \varphi') \in E(\mathcal{R})$ iff

- $\varphi = \varphi'$;

- $\varphi = \bigwedge_{j \in J} \varphi_j, \varphi' = \bigwedge_{j \in (J \setminus \{h\}) \cup I} \varphi_j$ and $(\varphi_h, \bigwedge_{i \in I} \varphi_i) \in \mathcal{R}$, for $I \cap J = \emptyset$;
- $\varphi = \bigwedge_{j \in J} \varphi_j, \varphi' = \bigwedge_{j \in (J \setminus \{i\})} \varphi_j$ and $(\varphi_i, \bigwedge_{j \in I} \varphi_j) \in \mathcal{R}$, for $I \subseteq J \setminus \{i\}$;
- $\varphi = \bigwedge_{j \in J} \varphi_j, \varphi' = \bigwedge_{i \in I} \varphi_i$ and there is a bijection $f: J \rightarrow I$ with $(\varphi_j, \varphi_{f(j)}) \in \mathcal{R}, \forall j \in J$;
- $\varphi = \neg \varphi_1, \varphi' = \neg \varphi_2$ and $(\varphi_1, \varphi_2) \in \mathcal{R}$;
- $\varphi = \langle a \rangle \psi, \varphi' = \langle a \rangle \psi'$ and $(\psi, \psi') \in \mathcal{R}^\dagger$.

It is easy to check that the transformation E is monotone on the complete lattice $(\mathcal{L}^s \times \mathcal{L}^s, \subseteq)$ and hence, by Tarski's theorem E has a greatest fixed point. We define the $\mathcal{L}$ -equivalence of formulae as such a greatest fixed point.

Definition 20 ($\mathcal{L}$ -equivalence). *The $\mathcal{L}$ -equivalence of formulae $\equiv_{\mathcal{L}} \subseteq \mathcal{L}^s \times \mathcal{L}^s$ is defined as $\equiv_{\mathcal{L}} = \max\{\mathcal{R} \subseteq \mathcal{L}^s \times \mathcal{L}^s \mid \mathcal{R} \subseteq E(\mathcal{R})\}$.*

4.3 The mimicking formulae

In [14] it was proved that the logic $\mathcal{L}$ is *adequate* for bisimilarity, i.e., two processes are bisimilar iff they satisfy the same formulae in $\mathcal{L}$. The drawback of this valuable result is in that to verify the equivalence we would need to test all the formulae definable in the logic, that is infinitely many formulae. As an alternative, in [15] a characterization of bisimilarity was given in terms of *characteristic formulae* [26] of processes, i.e., particular formulae that alone capture the entire equivalence class of the related process: if ϕ_s is the characteristic formula of process s for bisimilarity, then $s \sim t$ iff $t \models \phi_s$. This is the so called *expressive* characterization of an equivalence and allows us to establish process equivalence by testing a single formula. Unfortunately, also in this case there is a little drawback: to guarantee the possibility of constructing the characteristic formulae we need a very rich logic. For instance, [15] uses the probabilistic μ -calculus which, differently from $\mathcal{L}$, allows for arbitrary formulae to occur after the diamond modality and includes fixpoint operators.

Recently, [9,10] proposed a different technique for the characterization. When we compare the behavior of two processes, we compare those properties that are observable for them w.r.t. the considered semantics. The idea is to introduce a special formula, called *mimicking formula*, for each process expressing all and only its observable properties. In a broader sense, the mimicking formula of a process can be regarded as its specification. [9,10] showed that semantic equivalence of processes holds iff their mimicking formulae are syntactically equivalent (Thm. 4 below). Hence, to establish process equivalence we need only two formulae. Moreover, the logic on which the mimicking formulae are constructed is always *minimal* w.r.t. the chosen semantics, i.e., it only includes the operators necessary to express the observable properties w.r.t. that semantics.

Here, we recall the definition of mimicking formula and the *weak expressive* characterization of bisimilarity from [9]. Mimicking formulae are defined inductively over the depth of formulae as *up-to- k mimicking formulae*. Intuitively,

the *up-to- k mimicking formula* of process s , denoted by φ_s^k , characterizes the branching structure of the first k -steps of s by specifying which transitions are enabled for s as well as all the actions that it cannot perform.

Definition 21 (Mimicking formula, [9]). For a process $s \in \mathcal{S}$ and $k \in \mathbb{N}$, the up-to- k mimicking formula of s , notation φ_s^k , is defined inductively by

$$\begin{aligned}\varphi_s^0 &= \top \\ \varphi_s^k &= \bigwedge_{(s,a,\pi) \in \rightarrow} \langle a \rangle \bigoplus_{t \in \text{supp}(\pi)} \pi(t) \varphi_t^{k-1} \wedge \bigwedge_{b \notin \text{init}(s)} \neg \langle b \rangle \top\end{aligned}$$

Then, the mimicking formula of s , notation φ_s , is defined as $\varphi_s = \lim_{k \rightarrow \infty} \varphi_s^k$.

Example 5. Consider s in Fig. 2 and assume that $\mathcal{A} = \{a, b, c\}$. We aim to construct the mimicking formula of s . We have

$$\begin{aligned}\varphi_{\text{nil}} &= \neg \langle a \rangle \top \wedge \neg \langle b \rangle \top \wedge \neg \langle c \rangle \top & \varphi_{s_1} &= \langle b \rangle \varphi_{\text{nil}} \wedge \neg \langle a \rangle \top \wedge \neg \langle c \rangle \top \\ \varphi_{s_2} &= \langle b \rangle \varphi_{\text{nil}} \wedge \langle c \rangle \varphi_{\text{nil}} \wedge \neg \langle a \rangle \top & \varphi_{s_3} &= \langle c \rangle \varphi_{\text{nil}} \wedge \neg \langle a \rangle \top \wedge \neg \langle b \rangle \top \\ \varphi_{s_3} &= \varphi_{s_4} \\ \varphi_s &= \langle a \rangle \varphi_{s_1} \wedge \langle a \rangle (1/2 \varphi_{s_2} \oplus 1/2 \varphi_{s_3}) \wedge \langle a \rangle \varphi_{s_4} \wedge \neg \langle b \rangle \top \wedge \neg \langle c \rangle \top.\end{aligned}$$

Mimicking formulae allow us to characterize probabilistic bisimilarity.

Theorem 4 ([9]). Given any $s, t \in \mathcal{S}$ we have that $\varphi_s \equiv_{\mathcal{L}} \varphi_t$ iff $s \sim t$.

4.4 $\mathcal{L}$ -characterization of a family of bisimilarity metrics

In this Section we exploit the relation between the semantic properties of a process and the syntactic structure of its mimicking formula to provide a logical characterization of the family of bisimilarity metrics introduced in Sect. 4.1. The idea follows that of [9, 10]: 1. Firstly we transform the logic $\mathcal{L}$ into a family of metric spaces by defining a suitable *syntactic distance* over formulae. Intuitively, since distribution formulae are defined as probability distributions over state formulae, we can exploit the generalized Kantorovich metric to lift the distance over state formulae to a distance over distribution formulae. 2. Then we lift these syntactic distances to a family of pseudometrics over processes, called *logical distances*. Briefly, the logical distance $\ell_{\lambda, V}$ between two processes is defined as the syntactic distance between their mimicking formulae. 3. We show that the logical distance $\ell_{\lambda, V}$ coincides with the bisimilarity metric $\mathbf{d}_{\lambda, V}$ (Thm. 5 below).

The family of syntactic distances over formulae is defined inductively over the depth of formulae and their structure.

Definition 22 (Up-to- k distance on $\mathcal{L}$). Let $\lambda \in (0, 1]$ and let (V, d_V) be a metric space with $V \subseteq \mathbb{R}$ convex. For $k \in \mathbb{N}$, the up-to- k distance on state formulae is the mapping $\mathfrak{d}_{\lambda, V}^k: \mathcal{L}^s \times \mathcal{L}^s \rightarrow [0, +\infty]$ defined by:

$$\mathfrak{d}_{\lambda, V}^0(\varphi_1, \varphi_2) = 0 \text{ for all } \varphi_1, \varphi_2 \in \mathcal{L}^s$$

$$\mathfrak{d}_{\lambda,V}^k(\varphi_1, \varphi_2) = \begin{cases} 0 & \text{if } \varphi_1 = \top, \varphi_2 = \top \\ \mathfrak{d}_{\lambda,V}^k(\varphi'_1, \varphi'_2) & \text{if } \varphi_1 = \neg\varphi'_1, \varphi_2 = \neg\varphi'_2 \\ \lambda \cdot \mathbf{K}_V(\mathfrak{d}_{\lambda,V}^{k-1})(\psi_1, \psi_2) & \text{if } \varphi_1 = \langle a \rangle \psi_1, \varphi_2 = \langle a \rangle \psi_2 \\ \mathbf{H}(\mathfrak{d}_{\lambda,V}^k)(\{\varphi_j\}_{j \in J}, \{\varphi_i\}_{i \in I}) & \text{if } \varphi_1 = \bigwedge_{j \in J} \varphi_j, \varphi_2 = \bigwedge_{i \in I} \varphi_i \\ \mathcal{O}_{d_V}(V) & \text{otherwise.} \end{cases}$$

Clearly, the mapping $\mathfrak{d}_{\lambda,V}^k$ is a pseudometric and it is bounded whenever $\mathbf{K}_V$ is bounded. The discount factor $\lambda \in (0, 1]$ allows us to specify how much the distance between state formulae at the same depth is taken into account. For this reason, the discount factor λ is introduced in the evaluation of the distance between equally labeled diamond modalities.

We define the family of *syntactic distances over formulae*, denoted by $\mathfrak{d}_{\lambda,V}$, as the limit of their up-to- k distances. Since we consider only the metric spaces (V, d_V) for which $\mathbf{K}_V$ is bounded (cf. Remark 1), the existence of such a limit is ensured by the following two results.

Lemma 2. *For each $k \in \mathbb{N}$ and for all $\varphi, \varphi' \in \mathcal{L}^s$, $\mathfrak{d}_{\lambda,V}^{k+1}(\varphi, \varphi') \geq \mathfrak{d}_{\lambda,V}^k(\varphi, \varphi')$.*

Proposition 5. *The mapping $\mathfrak{d}_{\lambda,V}: \mathcal{L}^s \times \mathcal{L}^s \rightarrow [0, +\infty]$ defined, for all $\varphi, \varphi' \in \mathcal{L}^s$, by $\mathfrak{d}_{\lambda,V}(\varphi, \varphi') = \lim_{k \rightarrow \infty} \mathfrak{d}_{\lambda,V}^k(\varphi, \varphi')$ is well-defined.*

We are now ready to lift the metric on $\mathcal{L}$ to a metric on $\mathcal{S}$. To this aim, we exploit the close relation between processes and their own mimicking formulae.

Definition 23 (Logical distance). *For any $k \in \mathbb{N}$, the up-to- k logical distance $\ell_{\lambda,V}^k: \mathcal{S} \times \mathcal{S} \rightarrow [0, +\infty]$ over processes is defined for all $s, t \in \mathcal{S}$ by $\ell_{\lambda,V}^k(s, t) = \mathfrak{d}_{\lambda,V}^k(\varphi_s^k, \varphi_t^k)$. Then, the logical distance $\ell_{\lambda}: \mathcal{S} \times \mathcal{S} \rightarrow [0, +\infty]$ over processes is defined, for all $s, t \in \mathcal{S}$ by $\ell_{\lambda,V}(s, t) = \mathfrak{d}_{\lambda,V}(\varphi_s, \varphi_t)$.*

The next Theorem gives us the logical characterization of the generalized bisimilarity metrics in terms of the logical distances over processes.

Theorem 5. *Let $\lambda \in (0, 1]$. For any $s, t \in \mathcal{S}$ we have $\ell_{\lambda,V}(s, t) = \mathbf{d}_{\lambda,V}(s, t)$.*

4.5 A logical bound on $d_{\mathcal{X}}$ -privacy: the logical distance

We exploit the *multiplicative variant* of the logical distance over processes to obtain a *logical bound* on $d_{\mathcal{X}}$ -privacy. In detail, we model randomized mechanisms as LMCs and then: 1. We show that the multiplicative variant of the logical distance on the states of the LMC is an upper bound to the multiplicative total variation distance on the trace distributions induced by them. 2. We rephrase the $d_{\mathcal{X}}$ -privacy property as an upper bound on the logical distance between states corresponding to the considered secrets.

We remark that since we will use traces as a mere representation of the information on secrets, the actual length of the trace should play no role in the

evaluation of the distances. More precisely, the depth of the mimicking formula of the process that induces those traces in the LMC should not interfere in the evaluation of the distance as we are not interested in keeping track of the number of computation steps performed by a process, but, rather, in the possibility of executing them and the related execution probability. Hence, in the remaining of this Section we assume the discount factor $\lambda = 1$ and we omit it.

As shown in [12], we can express the multiplicative total variation distance in terms of the multiplicative variant of the Kantorovich lifting $\mathbf{K}_\otimes$ of the discrete metric over traces. More precisely, we let $\mathbf{dm}_{\otimes_V}$ be the $\otimes_{d_V}(V)$ -valued discrete metric over $\mathcal{A}^*$ which is defined as $\mathbf{dm}_{\otimes_V}(\alpha, \alpha') = 0$ if $\alpha = \alpha'$ and $\mathbf{dm}_{\otimes_V}(\alpha, \alpha') = \otimes_{d_V}(V)$ otherwise. To define $\mathbf{K}_\otimes$, we need to consider $V = [0, 1]$ and $d_\otimes(x, y) = |\ln(x) - \ln(y)|$. In [12] it has been proved that for $\otimes_{d_\otimes}([0, 1]) = +\infty$ it holds $tv_\otimes = \mathbf{K}_\otimes(\mathbf{dm}_{\otimes_\otimes})$. Hence, from $\mathbf{d}_{\lambda, \otimes} \geq \mathbf{K}_\otimes(\mathbf{dm}_{\otimes_\otimes})$ (cf. [12]) and Thm. 5 we obtain the following result.

Proposition 6. *Assume a LMC and let s, t be two processes in it. Let $\pi_s = \text{Tr}(s, \cdot)$ and $\pi_t = \text{Pr}(t, \cdot)$. Then $tv_\otimes(\pi_s, \pi_t) \leq \ell_\otimes(s, t)$.*

We remark that Prop. 2, Thm. 5 and Prop. 6 imply that $\mathbf{d}_\otimes^T \preceq \mathbf{d}_\otimes$.

We can then restate Def. 5 in terms of an upper bound on the multiplicative logical distance, thus obtaining the logical bound on $d_{\mathcal{X}}$ -privacy.

Theorem 6 (Logical bound on $d_{\mathcal{X}}$ -privacy). *Let M be a randomized mechanism defined by $M(\mathbf{x})(\alpha) = \text{Pr}(s_{\mathbf{x}}, \alpha)$ for all $\mathbf{x} \in \mathcal{X}$, $\alpha \in \mathcal{A}^*$. Then, given $\varepsilon > 0$, M is $\varepsilon \cdot d_{\mathcal{X}}$ -private if $\ell_\otimes(s_{\mathbf{x}}, s_{\mathbf{x}'}) \leq \varepsilon \cdot d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}') \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$.*

The following example illustrates a case of standard differential privacy.

Example 6. Consider two medical databases x and x' , both of size n ,⁴ and assume that they are adjacent, i.e. that they differ only for one individual record. Assume that we ask a counting query of the form $a = \text{"How many people in the database have the disease } d_a \text{"}$. Assume that, to sanitize the answer, we use a geometric mechanism [24], namely a probabilistic function that reports as answer the integer j with a probability distribution of the form $p_a(j) = c e^{-|i-j|\varepsilon}$, where i is the true answer, ε is the desired privacy level, and c is a normalization factor. In order to obtain a finite support, we can truncate the mechanism in the interval $[0, n]$, namely accumulate on 0 all the probability mass of the interval $(-\infty, 0]$, and on n all the probability mass of the interval $[n, +\infty)$. It is well known that the resulting mechanism is ε -differentially private. Consider now a new counting query of the form $b = \text{"How many people in the database have the disease } d_b \text{"}$, and again, assume that the answer is sanitized by a truncated geometric mechanism of the same form, with probability distribution p_b .

From the differential privacy literature we know that the combination of both mechanisms, in which the second query is asked after having obtained the answer from the first one, is 2ε -differentially private. However, we can obtain a better

⁴ We recall that we are using a notion of privacy in which all databases have the same number of records n , and the absence of a record is represented by a special value.

bound by looking at the various situations. To this purpose, let us consider the systems s and s' corresponding to the two databases x and x' respectively, and let p_a, p_b, p'_a and p'_b the probability distributions for the queries a and b in x and x' respectively. We can completely describe them by the mimicking formulae (which in this case are also characteristic formulae) φ and φ' defined as (for simplicity we omit the negative parts and the probabilities when they are 1):

$$\begin{aligned}\varphi_s &= \langle a \rangle \bigoplus_{j \in [0, n]} p_a(j) \langle j \rangle \langle b \rangle \bigoplus_{m \in [0, n]} p_b(m) \langle m \rangle \top \\ \varphi_{s'} &= \langle a \rangle \bigoplus_{j \in [0, n]} p'_a(j) \langle j \rangle \langle b \rangle \bigoplus_{m \in [0, n]} p'_b(m) \langle m \rangle \top\end{aligned}$$

Consider now the four scenarios obtained by combining the various cases that the individual corresponding to the new record in x' has or does not have the diseases d_a and d_b .

- If he does not have either of them, then p_a coincides with p'_a and p_b coincides with p'_b , which means that the distance between φ_s and $\varphi_{s'}$ is 0: the two systems are indistinguishable (0-differentially private).
- If he has d_a but not d_b , or vice versa, then either p_a coincides p'_a and the ratio between p_b and p'_b is bound by ε , or vice versa. The distance between φ_s and $\varphi_{s'}$ is ε : the two systems are ε -differentially private.
- If he has both d_a and d_b , then the ratio between p_a and p'_a , and that between p_b and p'_b , are bound by ε . The distance between φ_s and $\varphi_{s'}$ is 2ε : the two systems are 2ε -differentially private.

5 Conclusions

We have provided a logical characterization of generalized bisimulation metrics, based on the notions of mimicking formulae, i.e., formulae capturing the observable behavior of a particular process, and distance on formulae, i.e., a pseudometric on formulae measuring their syntactic disparities. Moreover, we have used the distance on formulae to obtain logical bounds on differential privacy properties. Then we have applied the same method to a simpler class of formulae expressing the trace semantics, thus obtaining a logical characterization of differential privacy and a classic logical characterization of weak anonymity.

As future work, we will further investigate the relation between the distance on formulae and real valued semantics on richer classes of formulae, by providing a thorough comparison with the real-valued semantics proposed in [17, 18] for the characterization of bisimulation semantics. Moreover, we aim at using the metrics and logical properties explored in this paper to reason about privacy in concurrent systems. This will require to deal with nondeterminism, which is already considered in the present paper, but probably we will need to reason explicitly about the scheduler and to restrict its capabilities, in order to avoid the problem of the “omniscient scheduler”, which could break any privacy defense. Finally, we aim at developing quantitative analysis techniques and tools for proving privacy properties.

References

1. de Alfaro, L., Faella, M., Stoelinga, M.: Linear and branching system metrics. *IEEE Trans. Software Eng.* 35(2), 258–273 (2009)
2. de Alfaro, L., Henzinger, T.A., Majumdar, R.: Discounting the Future in Systems Theory. In: *Proceedings of ICALP’03, Lecture Notes in Computer Science*, vol. 2719, pp. 1022–1037 (2003)
3. de Alfaro, L., Majumdar, R., Raman, V., Stoelinga, M.: Game refinement relations and metrics. *Logical Methods in Computer Science* 4(3) (2008)
4. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: Converging from branching to linear metrics on Markov chains. In: *Proceedings of ICTAC 2015. Lecture Notes in Computer Science*, vol. 9399, pp. 349–367 (2015)
5. Barthe, G., Köpf, B., Olmedo, F., Béguelin, S.Z.: Probabilistic relational reasoning for differential privacy. In: *Proc. of POPL. ACM* (2012)
6. Bernardo, M., De Nicola, R., Loreti, M.: Revisiting trace and testing equivalences for nondeterministic and probabilistic processes. *Logical Methods in Computer Science* 10(1) (2014)
7. Bloom, B., Fokkink, W.J., van Glabbeek, R.J.: Precongruence formats for decorated trace semantics. *ACM Trans. Comput. Log.* 5(1), 26–78 (2004)
8. van Breugel, F., Worrell, J.: A behavioural pseudometric for probabilistic transition systems. *Theor. Comput. Sci.* 331(1), 115–142 (2005)
9. Castiglioni, V., Gebler, D., Tini, S.: Logical characterization of bisimulation metrics. In: *Proceedings of QAPL’16. EPTCS*, vol. 227, pp. 44–62 (2016)
10. Castiglioni, V., Tini, S.: Logical characterization of trace metrics. In: *Proceedings of QAPL@ETAPS 2017. EPTCS*, vol. 250, pp. 39–74 (2017)
11. Chatzikokolakis, K., Andrés, M.E., Bordenabe, N.E., Palamidessi, C.: Broadening the scope of differential privacy using metrics. In: *Proceedings of PETS 2013. LNCS*, vol. 7981, pp. 82–102 (2013)
12. Chatzikokolakis, K., Gebler, D., Palamidessi, C., Xu, L.: Generalized bisimulation metrics. In: *Proceedings of CONCUR 2014. Lecture Notes in Computer Science*, vol. 8704, pp. 32–46 (2014)
13. Deng, Y., Chothia, T., Palamidessi, C., Pang, J.: Metrics for action-labelled quantitative transition systems. *Electr. Notes Theor. Comput. Sci.* 153(2), 79–96 (2006)
14. Deng, Y., Du, W.: Logical, metric, and algorithmic characterisations of probabilistic bisimulation. *CoRR* abs/1103.4577 (2011), <http://arxiv.org/abs/1103.4577>
15. Deng, Y., van Glabbeek, R.J.: Characterising probabilistic processes logically - (extended abstract). In: *Proceedings of LPAR-17*. pp. 278–293 (2010)
16. Deng, Y., Palamidessi, C., Pang, J.: Weak probabilistic anonymity. *ENTCS* 180(1), 55–76 (2007)
17. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled Markov processes. *Theor. Comput. Sci.* 318(3), 323–354 (2004)
18. Desharnais, J., Jagadeesan, R., Gupta, V., Panangaden, P.: The metric analogue of weak bisimulation for probabilistic processes. In: *Proceedings of LICS 2002*. pp. 413–422 (2002)
19. Du, W., Deng, Y., Gebler, D.: Behavioural pseudometrics for nondeterministic probabilistic systems. In: *Proceedings of SETTA 2016. Lecture Notes in Computer Science*, vol. 9984, pp. 67–84 (2016)
20. Duchi, J.C., Jordan, M.I., Wainwright, M.J.: Local privacy and statistical minimax rates. In: *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. pp. 429–438. *IEEE Computer Society* (2013)

21. Dwork, C.: Differential privacy. In: Proceedings of ICALP 2006. LNCS, vol. 4052, pp. 1–12 (2006)
22. Erlingsson, Ú., Pihur, V., Korolova, A.: RAPPOR: randomized aggregatable privacy-preserving ordinal response. In: Ahn, G., Yung, M., Li, N. (eds.) Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS). pp. 1054–1067. ACM (2014)
23. Gaboardi, M., Haeberlen, A., Hsu, J., Narayan, A., Pierce, B.C.: Linear dependent types for differential privacy. In: POPL. pp. 357–370 (2013)
24. Ghosh, A., Roughgarden, T., Sundararajan, M.: Universally utility-maximizing privacy mechanisms. In: Proceedings of the 41st annual ACM Symposium on Theory of Computing (STOC). pp. 351–360. ACM (2009)
25. Giacalone, A., Jou, C.C., Smolka, S.A.: Algebraic reasoning for probabilistic concurrent systems. In: Proceedings of IFIP Work. Conf. on Programming, Concepts and Methods. pp. 443–458 (1990)
26. Graf, S., Sifakis, J.: A modal characterization of observational congruence on finite terms of CCS. *Information and Control* 68(1-3), 125–145 (1986)
27. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *FAC* 6(5), 512–535 (1994)
28. Hennessy, M., Milner, R.: Algebraic laws for nondeterminism and concurrency. *J. Assoc. Comput. Mach.* 32, 137–161 (1985)
29. Hermanns, H., Parma, A., Segala, R., Wachter, B., Zhang, L.: Probabilistic logical characterization. *Inf. Comput.* 209(2), 154–172 (2011)
30. Keller, R.M.: Formal verification of parallel programs. *Commun. ACM* 19(7), 371–384 (1976)
31. Kwiatkowska, M.Z., Norman, G.: Probabilistic metric semantics for a simple language with recursion. In: Proceedings of MFCS’96. Lecture Notes in Computer Science, vol. 1113, pp. 419–430 (1996)
32. Larsen, K.G., Mardare, R., Panangaden, P.: Taking it to the limit: Approximate reasoning for Markov processes. In: Proceedings of MFCS 2012. Lecture Notes in Computer Science, vol. 7464, pp. 681–692 (2012)
33. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. *Inf. Comput.* 94(1), 1–28 (1991)
34. Machanavajjhala, A., Kifer, D., Abowd, J.M., Gehrke, J., Vilhuber, L.: Privacy: Theory meets practice on the map. In: Proceedings of ICDE 2008. pp. 277–286 (2008)
35. Narayanan, A., Shmatikov, V.: De-anonymizing social networks. In: Proceedings of S&P 2009. pp. 173–187 (2009)
36. 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)
37. Segala, R.: Modeling and Verification of Randomized Distributed Real-Time Systems. Ph.D. thesis, MIT (1995)
38. Smith, A.D.: Efficient, differentially private point estimators. CoRR abs/0809.4794 (2008), <http://arxiv.org/abs/0809.4794>
39. Song, L., Deng, Y., Cai, X.: Towards automatic measurement of probabilistic processes. In: Proceedings of QSIC 2007. pp. 50–59 (2007)
40. Tschantz, M.C., Kaynar, D., Datta, A.: Formal verification of differential privacy for interactive systems (extended abstract). *ENTCS* 276, 61–79 (sep 2011)
41. Xu, L., Chatzikokolakis, K., Lin, H.: Metrics for differential privacy in concurrent systems. In: Proc. of FORTE. LNCS, vol. 8461, pp. 199–215. Springer (2014)