A Logical Characterization of Differential Privacy

Valentina Castiglioni^{a,b,*}, Konstantinos Chatzikokolakis^c, Catuscia Palamidessi^d

^aINRIA Saclay - Île de France, France

^bReykjavik University, Iceland

^cCNRS and LIX Ecole Polytechnique, France

^dINRIA 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 *syntactic 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 formulae 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.

Keywords: Differential privacy, Metric semantics, Logical characterization, Nondeterministic probabilistic processes, Labeled Markov chains

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) [32], a framework originally proposed for protecting the privacy of participants in statistical databases, and now applied to geolocation [47], social networks [49] 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 by means of a *randomized mechanism*. 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.

In the literature we can find a wealth of proposals of notions of differential-privacy, based on variations on how the outputs of the randomized mechanism are compared. To avoid confusion, henceforth when we refer to DP, or standard DP, we intend the so called *pure* DP, also known as $(\varepsilon, 0)$ -DP.

*Corresponding author. Full address: Reykjavik University, Menntavegur 1, 101, Reykjavik, Iceland.
Email address: vale.castiglioni@gmail.com; valentinac@ru.is (Valentina Castiglioni)

Recently, a distributed variant of DP has emerged, called *local differential privacy* (LDP) [31]. 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, especially after large companies such as Apple and Google have started to adopt it for collecting the data of their users for statistical purposes [33].

1.1. Our goal

In this paper, we consider $d_{\mathcal{X}}$ -privacy [14], 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, and *weak probabilistic anonymity* [24], which measures the information leakage on user's identities. We study them in the context of *nondeterministic probabilistic transition systems* (PTSs) [52] and *labeled Markov chains* (LMCs) [46], 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. More formally, we provide a *logical characterization* of $d_{\mathcal{X}}$ -privacy and weak anonymity. To the best of our knowledge, ours is the first attempt in this direction.

1.2. Our contribution

Our work starts from a simple observation: each application of a mechanism, that we use to guarantee privacy protection, to a secret can be encoded as a LMC. Considering that the privacy protection guarantees are obtained from the comparison of the results of such applications, it is reasonable to exploit the studies on LMCs to obtain information on privacy properties.

A natural approach is that of *behavioral metrics* [19, 22, 27, 36, 44, 53, 56]. They were introduced to overcome the high sensitivity of behavioral equivalences and preorders with respect to tiny variations in the values of probabilities. Instead of stating whether the behavior of two processes is exactly the same or not, behavioral metrics *measure* the disparities in their behavior. Since, moreover, for verification purposes, the desired properties (and observable behavior) of processes are usually expressed in terms of modal formulae, *logical characterizations* of behavioral metrics have been thoroughly investigated [2, 11, 18, 25, 27]. As $d_{\mathcal{X}}$ -privacy and weak anonymity are measures over privacy protection guarantees of mechanisms, we aim at providing logical characterizations of them by exploiting the characterizations of the behavioral metrics on the LMCs induced by those mechanisms.

To this end, we consider the novel characterization technique recently proposed in [11, 12]. The idea is as follows:

- (i) We consider a boolean modal logic powerful enough to express the desired semantics.
- (ii) We define a *syntactic distance* over the formulae in the chosen logic, namely a pseudometric on formulae measuring their syntactic disparities.
- (iii) We express the differences in the behavior of processes, and thus the behavioral metric, in terms of such distance.

In detail, to obtain the *logical characterization* of $d_{\mathcal{X}}$ -privacy, we reason in terms of *trace semantics* over LMCs. We consider a probabilistic refinement $\mathbb{L}$ of the Hennessy-Milner logic (HML) [40] and we propose a novel notion of *trace metric* defined via the syntactic distance over formulae in $\mathbb{L}$. 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. We show that the value of the trace distance between two LMCs equals the privacy guarantee of the mechanism that induced those LMCs. Interestingly, we also show how it is possible to define a *real-valued semantics* for formulae in $\mathbb{L}$ starting from their syntactic distance. From this, we obtain a *logical characterization* of *weak anonymity* in the classic sense of [2, 27] (see Section 5 for a detailed description).

Then we switch from trace to bisimulation semantics and we provide a *logical bound* on $d_{\mathcal{X}}$ -privacy. We consider processes in the PTS model, which enriches LMCs with nondeterministic choices, and we study *generalized bisimulation metrics* [15] on them. Informally, [15] defines bisimulation metrics via a generalized notion of *Kantorovich lifting*, which allows to define distances suitable to deal with privacy and security properties. Then we consider the modal logic $\mathcal{L}$ from [23], which extends HML with a probabilistic choice operator that allows us to properly express the probabilistic behavior of processes with respect to the bisimulation semantics. By means of $\mathcal{L}$ we provide a logical characterization of generalized bisimulation metrics via the syntactic distance over formulae in $\mathcal{L}$ and the notion of *mimicking formula* of a process [11]. 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., [2]), and thus of approximations on $d_{\mathcal{X}}$ -privacy. Furthermore, we could exploit the *non-expansiveness* results obtained in [15] to favor compositional reasoning over $d_{\mathcal{X}}$ -privacy.

1.3. Summary of results

Our contribution can then be summarized as follows:

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 $\mathcal{L}$, a probabilistic extension of HML, 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.

1.4. Organization of contents

We start by reviewing the background in Section 2 and we recall some basic notions on $d_{\mathcal{X}}$ -privacy and weak anonymity in Section 3. Section 4 comes with our first contribution, namely the definition of a trace metric on processes in terms of a syntactic distance on modal formulae and the logical characterization of $d_{\mathcal{X}}$ -privacy obtained from it. Our second contribution, the definition of a real-valued semantics for formulae via the syntactic distance and the logical characterization of weak anonymity built on it, is presented in Section 5. In Section 6 we present the generalized bisimilarity metrics and in Section 7 we introduce the modal logic $\mathcal{L}$ and the mimicking formulae of processes. We present, in Section 8, our third contribution, namely how we can combine the mimicking formulae and the syntactic distance to obtain a logical characterization of generalized bisimilarity metrics. Then, in Section 9, we show how to obtain bounds on $d_{\mathcal{X}}$ -privacy from such a characterization. In Section 10 we discuss related work and some possible extensions of our work. Finally, we draw some conclusions in Section 11.

1.5. What's new

A preliminary version of this paper appeared as [10]. Besides providing the full proofs of our results and new examples, we have enriched our previous contribution as follows:

- a. We study the *Dining Cryptographers Protocol* and show how the well known anonymity results for it can be obtained via our technique (Section 5.2).

- b. We show that the generalized bisimilarity metrics can be obtained as the limit of the *up-to-k generalized bisimilarity metrics*, namely distances taking into account the behavioral differences observable in the first k computation steps (Section 6.3).
- c. We discuss the expressive power of the modal logic $\mathcal{L}$ with respect to other logics used in the literature to provide characterization results for probabilistic relations, showing that $\mathcal{L}$ is more expressive (Section 10.2).
- d. We discuss the extension of our characterization technique to processes with recursion. We argue that by means of the *equational μ -calculus* framework [1, 45, 51] our results can also be obtained in the case of recursion (Section 10.3).

2. Background

In this section we review the preliminary notions on probabilistic processes and metric spaces that are necessary for our dissertation.

The PTS model. Nondeterministic probabilistic labeled transition systems (PTSs) [52] combine LTSs [42] and discrete time Markov chains [39], to model reactive behavior, nondeterminism and probability. In a PTS, 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 $\sum_{s \in \mathcal{S}} \pi(s) = 1$. The *support* of π is the set $\text{supp}(\pi) = \{x \in X \mid \pi(x) > 0\}$. By $\Delta(X)$ we denote the set of all *finitely supported* distributions over X , ranged over by $\pi, \pi', \dots$. 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, [52]). A *nondeterministic probabilistic labeled transition system (PTS)* is a triple $(\mathcal{S}, \mathcal{A}, \rightarrow)$, where: (i) $\mathcal{S}$ is a countable set of processes, (ii) $\mathcal{A}$ is a countable set of actions, and (iii) $\rightarrow \subseteq \mathcal{S} \times \mathcal{A} \times \Delta(\mathcal{S})$ 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 . Finally, a PTS is *image-finite* [41] if $\text{der}(s, a)$ is finite for each $s \in \mathcal{S}$ and $a \in \mathcal{A}$. We consider only image-finite PTSs. Moreover, for sake of readability and to limit the amount of purely technical content of the paper, we consider processes *without recursion*. We refer the interested reader to Section 10.3 for a discussion on how our results can be obtained when also recursion is taken into account.

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 $\alpha \in \mathcal{A}^*$ recursively as follows:

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

For a process $s \in \mathcal{S}$ and a trace $\alpha \in \mathcal{A}^*$, we will sometimes refer to $\text{Pr}(s, \alpha)$ as to the *execution probability* of α by s .

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 [2]).

Figure 1: s is an arbitrary process in the PTS model, whereas t, u and v are three LMCs. For simplicity, an arrow $s \xrightarrow{a}$ with no target models the evolution of process s , via the execution of action a , to the Dirac distribution δ_{nil} , with nil process that can execute no action.

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 $\Pr(s, \alpha) = \Pr(t, \alpha)$.

Example 1. Consider processes t, u, v in Figure 1. We have

$\Pr(t, a) = 1$	$\Pr(u, a) = 1$	$\Pr(v, a) = 1$
$\Pr(t, ab) = 1$	$\Pr(u, ab) = 1$	$\Pr(v, ab) = 1$
$\Pr(t, abc) = 0.3$	$\Pr(u, abc) = 0.3$	$\Pr(v, abc) = 0.35$
$\Pr(t, abd) = 0.3$	$\Pr(u, abd) = 0.3$	$\Pr(v, abd) = 0.35$

thus giving that t and u are trace equivalent, whereas v is not trace equivalent to them. $\square$

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 $d: X \times X \rightarrow [0, +\infty]$. Henceforth, we shall use the term pseudometric to refer to both, pseudometrics and extended pseudometrics, since the meaning will always be clear from the co-domain. 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* with respect to 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 $\text{1-Lip}[(X, d_X), (Y, d_Y)]$ the set of such functions.

Given any (pseudo)metric space (X, d) , the *diameter* of X with respect to d , denoted by $\odot_d(X)$, is the maximal distance of two elements in X , namely $\odot_d(X) = \sup_{x, y \in X} d(x, y)$.

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 non-empty sets $X_1, X_2 \subseteq X$ by

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

3. Differential privacy

In this section we briefly recall the definitions of the privacy notions that are of central interest in this paper, namely *d_X -privacy* and *weak probabilistic anonymity*. We refer the interested reader to [14] for more details on the differential privacy framework and to [24] for weak probabilistic anonymity.

3.1. $d_{\mathcal{X}}$ -privacy

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, namely $d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}')$ 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 the 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). 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, [14]). 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 if and only if

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

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 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 LDP. The mechanism is called “Randomized responses” and is a simplified instance of the system RAPPOR used by Google to protect the privacy of their users [33].

Example 2 (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 Figure 2 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, it 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. $\square$

Figure 2: The mechanism ‘Randomized responses’ as a LMC.

Figure 3: LMCs representing the behavior of an arbitrary paying cryptographer c_i and of the master m .

3.2. Weak anonymity

Weak probabilistic anonymity [24] 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. Informally, a randomized mechanism M is ε -*weak anonymous* if after its interaction with the system an adversary is more likely to identify $\mathbf{x}$ as the user who performed the tasks than user $\mathbf{x}'$ by an additive factor ε . 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 6 (Weak probabilistic anonymity [24]). 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 and only if

$$tv(M(\mathbf{x}), M(\mathbf{x}')) \leq \varepsilon \quad \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}.$$

Example 3 (Dining cryptographers). One of the most known anonymity protocols, from the literature, is that of *Dining Cryptographers* [16]. Briefly, an arbitrary number of cryptographers are dining together with their master. At the end of the dinner, the master will choose whether to pay the bill or not. In the latter case, he will secretly say to each cryptographer if he has to pay or not. An external observer will be able to discover whether the payer is the master or one of the cryptographers. However, in the latter case, the identity of the payer should not be disclosed.

The original solution in [16], to maintain the paying cryptographer anonymous, is to associate a (fair) coin to each cryptographer making it visible to himself and to his right neighbor. These coins are then tossed, and each cryptographer computes the binary sum of the adjacent coins (by letting, for instance, *head*= 0 and *tail*= 1), adds 1 if he is the payer, and outputs the result. It is then proved that the payer is one of the cryptographers if and only if the binary sum of all the so obtained outputs is 1.

The strong anonymity of the protocol is based on the fact that under fair coins all output configurations have the same probability and, moreover, for each output obtained in a configuration where the payer is cryptographer i , there is a different coin configuration producing the same output when the payer is cryptographer j .

	c_1	c_2	c_3
p_1^i	$p_1 \cdot (1 - p_2)$	$(1 - p_1) \cdot (1 - p_2)$	$(1 - p_1) \cdot p_2$
p_2^i	$p_1 \cdot p_2$	$(1 - p_1) \cdot p_2$	$(1 - p_1) \cdot (1 - p_2)$
p_3^i	$(1 - p_1) \cdot p_2$	$p_1 \cdot p_2$	$p_1 \cdot (1 - p_2)$
p_4^i	$(1 - p_1) \cdot (1 - p_2)$	$p_1 \cdot (1 - p_2)$	$p_1 \cdot p_2$

Table 1: The probabilities p_j^i evaluated with respect to each paying cryptographer c_i .

We are therefore interested in analyzing the case of weak anonymity, namely the one in which we consider *biased* coins, and thus some information on the payer's identity may be leaked due to the *biased factor* of the coins. For simplicity of presentation, we will consider the case of a master with three cryptographers. The LMCs in Figure 3 represent the behavior of any cryptographer c_i who has been selected as a payer, and the behavior of the master m when he is willing to pay the bill. The action of the randomized mechanism, in this case, can be subsumed in the choice of the biased factor of the coins. The eight actions $o_1, \dots, o_8$ represent the possible outputs of the computation by the three cryptographers. In detail, we have

$$\begin{aligned} o_1 &= 111 & o_2 &= 100 & o_3 &= 010 & o_4 &= 001 \\ o_5 &= 110 & o_6 &= 101 & o_7 &= 011 & o_8 &= 000 \end{aligned}$$

in which the i -th bit represents the output of the i -th cryptographer. The most significant ones are the first four $o_1, \dots, o_4$, in that they are obtained only when the payer is one of the cryptographers and the different probability of observing an output with respect to another could reveal information about the identity of the payer. Having three cryptographers, we only need two coins, which we assume to be biased: let p_1 be the probability of obtaining 0 on the first coin, and p_2 the corresponding probability for the second coin. Then, the probability p_j^i of the observable o_j being obtained when cryptographer c_i is the payer is evaluated in Table 1. In [24] it is then proved that weak anonymity parameter ε of the protocol depends on the biased factor of the coins as follows

$$\varepsilon = \begin{cases} |1 - (p_1 + p_2)| & \text{if } p_1, p_2 \leq 1/2 \text{ or } p_1, p_2 \geq 1/2 \\ |p_1 - p_2| & \text{otherwise.} \end{cases}$$

□

4. Logical characterization of $d_{\mathcal{X}}$ -privacy: a trace metric approach

In this section we present the first proposal of a logical characterization of $d_{\mathcal{X}}$ -privacy. To obtain it, we investigate the semantics of the LMCs induced by the randomized mechanisms. In particular, we exploit a notion of *trace metric* evaluated on *modal formulae* expressing linear properties of LMCs. Informally, we 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 define a metric for such a semantics in terms of a syntactic distance over the formulae in the considered logic and we 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 allows us to obtain the first *logical characterization* of $d_{\mathcal{X}}$ -privacy (Theorem 2): we show that the trace metrics so defined on LMCs coincides with the multiplicative variant of the total variation distance (Proposition 2).

4.1. Trace metrics on LMCs

Probabilistic trace semantics compares the behavior of processes with respect to 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 Definition 2 is an example, and we refer the interested reader to [7] for a survey. Such a wealth of notions derives from the interplay of nondeterminism 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 with respect to the same linear properties, that is their differences in the probabilities of executing the same traces (see, e.g., [9, 19, 53]).

As the focus of this paper is on $d_{\mathcal{X}}$ -privacy, we adopt a different approach, with respect to those referenced, to the definition of a trace metric on LMCs. In fact, we hark back to the seminal work [27] 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 of the formulae distinguishing them.

Informally, in [27] 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 modal logic $\mathbb{L}$ extends the class of formulae used in the fully nondeterministic case to express trace semantics [8] (and corresponding to the subclass of *linear formulae*) with a probabilistic modality allowing us to express the execution probabilities of traces.

Definition 7 (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:

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

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 $\text{Tr}(\Phi) = \alpha$, if the sequence of action labels in α is exactly the same sequence of labels of the diamond modalities in Φ , i.e., $\text{Tr}(\langle a_1 \rangle \dots \langle a_n \rangle \top) = \alpha$ iff $\alpha = a_1 \dots a_n$. In particular we have that $\text{Tr}(\top) = \epsilon$.

Definition 8 (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
 - for $r > 0$, $s \models \Phi$ and $\text{Pr}(s, \text{Tr}(\Phi)) = r$;
 - for $r = 0$, $\text{Pr}(s, \text{Tr}(\Phi)) = 0$.

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

Example 4 (Randomized responses II). Consider processes s_y, s_n in Figure 2. One can easily check that, by omitting the occurrences of probabilistic formulae of the form 0Φ , for some linear formula Φ ,

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

□

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 $\mathbb{L}(s) = \mathbb{L}(t)$.

Proof. The proof can be found in Appendix A. □

We can now proceed to the definition of the *trace metric*. The definition of the *syntactic distance* on formulae in $\mathbb{L}$ is parametric with respect to 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 9 (Distance on $\mathbb{L}$). Let $([0, 1], \mathcal{D})$ be a metric space. The function $\mathbf{dm}_{\mathcal{D}}: \mathbb{L}^1 \times \mathbb{L}^1 \rightarrow \{0, \mathcal{O}_{\mathcal{D}}([0, 1])\}$ 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 $\mathbf{dm}_{\mathcal{D}}(\Phi_1, \Phi_2) = \mathcal{O}_{\mathcal{D}}([0, 1])$ otherwise. The function $\mathfrak{d}_{\mathcal{D}}^p: \mathbb{L}^p \times \mathbb{L}^p \rightarrow [0, \mathcal{O}_{\mathcal{D}}([0, 1])]$ is defined over $\mathbb{L}^p$ as follows:

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

The trace metric on processes is then defined as the Hausdorff lifting of the syntactic distance on $\mathbb{L}$ to the sets of formulae satisfied by the processes.

Definition 10 (Trace metric). Let $([0, 1], \mathcal{D})$ be a metric space. The *trace metric* over processes $\mathbf{d}_{\mathcal{D}}^T: \mathcal{S} \times \mathcal{S} \rightarrow [0, \mathcal{O}_{\mathcal{D}}([0, 1])]$ is defined for all $s, t \in \mathcal{S}$ by

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

Example 5. Consider processes u, v in Figure 1. We have

$$\begin{aligned}\mathbb{L}(u) &= \{1\langle a \rangle\top, 1\langle a \rangle\langle b \rangle\top, 0.3\langle a \rangle\langle b \rangle\langle c \rangle\top, 0.3\langle a \rangle\langle b \rangle\langle d \rangle\top\} \\ \mathbb{L}(v) &= \{1\langle a \rangle\top, 1\langle a \rangle\langle b \rangle\top, 0.35\langle a \rangle\langle b \rangle\langle c \rangle\top, 0.35\langle a \rangle\langle b \rangle\langle d \rangle\top\}.\end{aligned}$$

Consider, for instance, as ground distance $\mathcal{D}$ the Euclidean distance, namely $\mathcal{D}(r_1, r_2) = |r_1 - r_2|$ for all $r_1, r_2 \in [0, 1]$. Then, we have

$$\begin{aligned}\mathbf{d}_{\mathcal{D}}^T(u, v) &= \mathbf{H}(\mathfrak{d}_{\mathcal{D}}^p)(\mathbb{L}(u), \mathbb{L}(v)) \\ &= \max \left\{ \begin{array}{l} \mathfrak{d}_{\mathcal{D}}^p(1\langle a \rangle\top, 1\langle a \rangle\top), \\ \mathfrak{d}_{\mathcal{D}}^p(1\langle a \rangle\langle b \rangle\top, 1\langle a \rangle\langle b \rangle\top), \\ \mathfrak{d}_{\mathcal{D}}^p(0.3\langle a \rangle\langle b \rangle\langle c \rangle\top, 0.35\langle a \rangle\langle b \rangle\langle c \rangle\top), \\ \mathfrak{d}_{\mathcal{D}}^p(0.3\langle a \rangle\langle b \rangle\langle d \rangle\top, 0.35\langle a \rangle\langle b \rangle\langle d \rangle\top) \end{array} \right\} \\ &= \mathfrak{d}_{\mathcal{D}}^p(0.3\langle a \rangle\langle b \rangle\langle c \rangle\top, 0.35\langle a \rangle\langle b \rangle\langle c \rangle\top) \\ &= 0.05.\end{aligned}$$

By considering a different distance on the probability weights, we can obtain the *multiplicative variant* of our trace metric, which we will denote by $\mathbf{d}_{\otimes}^T$ and that will play a fundamental role in our characterization theorem. Let $\mathcal{D}_{\otimes}(r_1, r_2) = |\ln(r_1) - \ln(r_2)|$, namely $\mathcal{D}$ measures the absolute value of the difference between

the natural logarithms of the probability weights and let $\mathbf{d}_{\otimes}^T$ be the trace metric obtained using $\mathcal{D}_{\otimes}$ as ground distance. Then, for the processes u, v of this example we obtain

$$\mathbf{d}_{\otimes}^T(u, v) = \mathfrak{d}_{\otimes}^p(0.3\langle a \rangle \langle b \rangle \langle c \rangle \top, 0.35\langle a \rangle \langle b \rangle \langle c \rangle \top) = |\ln(0.3) - \ln(0.35)| \approx 0.15.$$

□

We can prove that 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 $\sim_{\text{Tr}} = \ker(\mathbf{d}_{\mathcal{D}}^T)$.*

Proof. The proof can be found in Appendix A. □

4.2. 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 [11, 12, 27], 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 consider the multiplicative variant $\mathbf{d}_{\otimes}^T$ of our trace metric introduced in Example 5. 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 $\mathbf{d}_{\otimes}^T(s, t) = tv_{\otimes}(\mu_s, \mu_t)$.*

Proof. The proof can be found in Appendix A. □

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

Theorem 2 (Logical characterization of $d_{\mathcal{X}}$ -privacy). *Consider the randomized mechanism M 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 and only if*

$$\mathbf{d}_{\otimes}^T(s_{\mathbf{x}}, s_{\mathbf{x}'}) \leq \varepsilon \cdot d_{\mathcal{X}}(\mathbf{x}, \mathbf{x}') \quad \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}.$$

Proof. Immediate from Proposition 2 and Definition 5. □

Example 6 (Randomized responses, III). We can show that the mechanism ‘Randomized responses’ described in Example 2 is log 3-locally differentially private by evaluating the trace distance between processes s_y and s_n in Figure 2. By comparing the sets of formulae $\mathbb{L}(s_y)$ and $\mathbb{L}(s_n)$ given in Example 4, we obtain

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

□

5. 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 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, e.g., [19, 21, 27, 28, 30]). In this section, we exploit the syntactic distance on $\mathbb{L}$ to provide a real-valued semantics for formulae in $\mathbb{L}$ and thus a characterization of *weak probabilistic anonymity* expressed according to the classic schema in (1) (Theorem 3).

In detail, we consider all metric spaces $([0, 1], \mathcal{D})$ with $0 < \mathcal{O}_{\mathcal{D}}([0, 1]) < \infty$ and:

- (i) We use the syntactic distance over formulae in $\mathbb{L}$ to define a (generalized) real-valued semantics for those modal formulae.
- (ii) We show that the total variation distance satisfies the general schema in (1) with respect to such real semantics.
- (iii) 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.

5.1. Real valued semantics

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 $s \models \phi$. Our aim is to exploit our distance over formulae to define such a semantics.

Example 7. Consider process s in Figure 1. We aim at evaluating *how much* s satisfies the formulae $\Psi_{1/2} = 1/2\langle a \rangle \langle c \rangle \top$ and $\Psi^q = q\langle a \rangle \langle d \rangle \top$, for $q \in [0, 1]$.

We start with $\Psi_{1/2}$. First of all, notice that process s can perform the trace ac and, in particular, we have that $s \models 3/5\langle a \rangle \langle c \rangle \top$ and $s \models 1\langle a \rangle \langle c \rangle \top$. By comparing the formula $\Psi_{1/2}$ with the latter, we shall say that there is a discrepancy of $1/2$ between the required execution probability (namely $1/2$) and the actual one (namely 1). Clearly, by comparing $\Psi_{1/2}$ with $3/5\langle a \rangle \langle c \rangle \top$, such discrepancy decreases to $1/10$, and hence s shows a better behavior with respect to the required one in this second case. Roughly speaking, we should say that process s satisfies the formula $\Psi_{1/2}$ except for $1/10$. It is then reasonable to set the value of $\Psi_{1/2}$ in s to $9/10$.

Consider now the family of formulae Ψ^q . Clearly, process s does not perform trace ad , namely $s \models 0\langle a \rangle \langle d \rangle \top$. This could suggest to set the value of Ψ^q in s to 0, for each $q \neq 0$. However, such solution would flatten the differences in the quantitative requirements imposed by those formulae on varying of $q \in [0, 1]$. In fact the discrepancy between not performing a trace and performing it with a tiny probability should be negligible with respect to the one between not performing a trace and performing it with high probability. It is then reasonable to set the value of Ψ^q in s to $1 - q$: the higher q , the lower the value of Ψ^q . $\square$

The intuitions discussed in Example 7 can be generalized and formalized as follows. Let L be the class of formulae of interest, let $D_{\mathcal{D}}$ be any generalized syntactic distance defined on L , like, e.g., the distance $\mathfrak{D}_{\mathcal{D}}^p$ for the logic $\mathbb{L}$. 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 :

1. 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 $\inf_{\phi' \in L(s)} D_{\mathcal{D}}(\phi, \phi')$.
2. Then we simply notice that being $D_{\mathcal{D}}(\phi, \phi')$ far from s is equivalent to be $\mathcal{O}_{\mathcal{D}}([0, 1]) - \inf_{\phi' \in L(s)} D_{\mathcal{D}}(\phi, \phi')$ close to it. We remark that $\mathcal{O}_{\mathcal{D}}([0, 1])$ has to be finite in order to obtain a meaningful value.
3. Finally, we assign to ϕ the value

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

in s , where the normalization with respect to $\mathcal{O}_{\mathcal{D}}([0, 1])$ ensures that this value is in $[0, 1]$.

Definition 11 (Real valued semantics). Let $([0, 1], \mathcal{D})$ be a metric space with $0 < \mathcal{O}_{\mathcal{D}}([0, 1]) < \infty$. 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 8. Consider process s in Figure 1 and the family of formulae $\Psi_p = p\langle a \rangle \langle c \rangle \top$, for $p \in [0, 1]$. In Example 7 we have discussed the value of the formula $\Psi_{1/2}$, from such a family, in s . We now formalize the intuition provided there and study the value of formulae Ψ_p in s on varying of $p \in [0, 1]$. By Definition 11, for $\mathcal{D}$ the Euclidean distance (and thus omitting the $\mathcal{D}$ subscript), we have that for each $p \in [0, 1]$

$$\llbracket \Psi_p \rrbracket(s) = 1 - \frac{\inf_{\Psi \in \mathbb{L}(s)} \mathfrak{d}^P(\Psi_p, \Psi)}{\mathcal{O}([0, 1])} = 1 - \inf_{\Psi \in \mathbb{L}(s)} \mathfrak{d}^P(\Psi_p, \Psi)$$

since $\mathcal{O}([0, 1]) = 1$. We recall that process s performs the trace ac so that the infimum on $\mathbb{L}(s)$ is to be evaluated over the formulae satisfied by s that subsume the execution of the trace ac . In particular, we have that $s \models \Psi_{3/5}$ and $s \models \Psi_1$, so that

$$\llbracket \Psi_p \rrbracket(s) = 1 - \min\{|p - 3/5|, |p - 1|\}.$$

□

5.2. Logical characterization of weak anonymity

Before proceeding to the characterization of weak anonymity, 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 non-empty $L_1, L_2 \subseteq L$ we have that

$$\mathbf{H}(D_{\mathcal{D}})(L_1, L_2) = \sup_{\phi \in L} \left| \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \right|.$$

Proof. The proof can be found in Appendix B. □

If we focus on the class of formulae $\mathbb{L}$, from Proposition 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 $0 < \mathcal{O}_{\mathcal{D}}([0, 1]) < \infty$. For all processes $s, t \in \mathcal{S}$ it holds that

$$\mathbf{d}_{\mathcal{D}}^T(s, t) = \sup_{\Psi \in \mathbb{L}^P} |\llbracket \Psi \rrbracket_{\mathcal{D}}(s) - \llbracket \Psi \rrbracket_{\mathcal{D}}(t)|.$$

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 $\mathcal{D}(x, y) = |x - y|$ Euclidean distance, 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 $0 < \mathcal{O}_{\mathcal{D}}([0, 1]) < \infty$. For any $s \in \mathcal{S}$ define $\mu_s = \text{Pr}(s, \cdot)$. 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

$$tv(\mu_s, \mu_t) = \sup_{\Phi \in \mathbb{L}^1} |\llbracket \Phi \rrbracket(s) - \llbracket \Phi \rrbracket(t)|.$$

Proof. The proof can be found in Appendix B. □

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).

	c_1	c_2	c_3
$\llbracket \Phi_1 \rrbracket(c_i)$	$1 - p_1 \cdot (1 - p_2)$	$1 - (1 - p_1) \cdot (1 - p_2)$	$1 - (1 - p_1) \cdot p_2$
$\llbracket \Phi_2 \rrbracket(c_i)$	$1 - p_1 \cdot p_2$	$1 - (1 - p_1) \cdot p_2$	$1 - (1 - p_1) \cdot (1 - p_2)$
$\llbracket \Phi_3 \rrbracket(c_i)$	$1 - (1 - p_1) \cdot p_2$	$1 - p_1 \cdot p_2$	$1 - p_1 \cdot (1 - p_2)$
$\llbracket \Phi_4 \rrbracket(c_i)$	$1 - (1 - p_1) \cdot (1 - p_2)$	$1 - p_1 \cdot (1 - p_2)$	$1 - p_1 \cdot p_2$

Table 2: The real-valued semantics of formulae Φ_j evaluated with respect to each paying cryptographer c_i .

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, for $\varepsilon > 0$, M satisfies ε -weak anonymity if and only if*

$$\sup_{\Phi \in \mathbb{L}^1} |\llbracket \Phi \rrbracket(s_{\mathbf{x}}) - \llbracket \Phi \rrbracket(s_{\mathbf{x}'})| \leq \varepsilon \quad \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}.$$

Proof. Immediate from Proposition 4 and Definition 6. $\square$

Example 9 (Dining cryptographers II). We show that we can regain the ε -weak anonymity result discussed in Example 3 by applying Theorem 3. For $j \in \{1, \dots, 4\}$, consider the linear formulae $\Phi_j = \langle \tau \rangle \langle o_j \rangle \top$ that express the fact that the output o_j is observable. Then, accordingly to Definition 11, where we omit the $\mathcal{D}$ subscript since we consider the Euclidean distance, and the probabilities reported in Table 1, for the formulae Φ_j we obtain the real-valued semantics showed in Table 2. Hence, we have that

$$\varepsilon = \max_{i, h \in \{1, 2, 3\}, j \in \{1, 2, 3, 4\}} |\llbracket \Phi_j \rrbracket(c_i) - \llbracket \Phi_j \rrbracket(c_h)|.$$

Firstly, we notice that

$$\begin{aligned} \max_{i, h \in \{1, 2, 3\}} |\llbracket \Phi_1 \rrbracket(c_i) - \llbracket \Phi_1 \rrbracket(c_h)| &= \max \{ |(2p_1 - 1) \cdot (1 - p_2)|, |(1 - p_1) \cdot (1 - 2p_2)|, |p_1 - p_2| \} \\ \max_{i, h \in \{1, 2, 3\}} |\llbracket \Phi_2 \rrbracket(c_i) - \llbracket \Phi_2 \rrbracket(c_h)| &= \max \{ |(2p_1 - 1) \cdot p_2|, |(1 - p_1) \cdot (2p_2 - 1)|, |1 - (p_1 + p_2)| \} \\ \max_{i, h \in \{1, 2, 3\}} |\llbracket \Phi_3 \rrbracket(c_i) - \llbracket \Phi_3 \rrbracket(c_h)| &= \max \{ |(1 - 2p_1) \cdot p_2|, |p_1 \cdot (2p_2 - 1)|, |p_1 - p_2| \} \\ \max_{i, h \in \{1, 2, 3\}} |\llbracket \Phi_4 \rrbracket(c_i) - \llbracket \Phi_4 \rrbracket(c_h)| &= \max \{ |(1 - 2p_1) \cdot (1 - p_2)|, |p_1 \cdot (1 - 2p_2)|, |1 - (p_1 + p_2)| \} \end{aligned}$$

From these, through case analysis and simple algebra, we regain

$$\varepsilon = \begin{cases} |1 - (p_1 + p_2)| & \text{if } p_1, p_2 \leq 1/2 \text{ or } p_1, p_2 \geq 1/2 \\ |p_1 - p_2| & \text{otherwise.} \end{cases}$$

$\square$

6. 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 [2] 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 *PTSs*, we consider the *generalized bisimulation metrics* introduced in [15] 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 [11, 12]. As in Section 4.1, the former is a pseudometric on a probabilistic version of HML, $\mathcal{L}$ from [23], that extends $\mathbb{L}$ with modalities allowing us to express the interplay of nondeterminism and probability typical of PTSs (Section 7). The latter is a special formula in $\mathcal{L}$ that alone expresses the observable behavior with respect to bisimulation semantics of the process to which it is related and allows us to characterize bisimilarity (Section 7.1). Then we show that the syntactic distance between the mimicking formulae of processes equals their bisimulation distance (Section 8) and that, when we focus on LMCs, it gives an upper bound on $d_{\mathcal{X}}$ -privacy properties of mechanisms (Section 9).

As a final remark, note that using bisimulation metrics and their characterization would allow us to apply the compositional results obtained for them in [15] also to $d_{\mathcal{X}}$ -privacy properties (see Section 11 for more insights into this research line). Now, we proceed to recall some base notions on bisimulation semantics and generalized bisimulation metrics.

6.1. Probabilistic bisimulations

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 equivalence. 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 $\mathcal{R}^\dagger$, if and only if they assign the same probabilistic weights to processes related by $\mathcal{R}$.

Definition 12 (Relation lifting, [25]). 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 finite set of indexes I s.t.

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

Notice that the x_i , for $i \in I$, are not required to be distinct. Similarly for the y_i , for $i \in I$. A different formulation to Definition 12 has been provided elsewhere in the literature in terms of couplings [52]. In [25] it was proved that the two formulations are equivalent. We decided to present the version of [25] since it will simplify the technical reasoning in some of the upcoming proofs.

Definition 13 (Probabilistic bisimulation, [46]). 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.

6.2. Generalized bisimulation metrics

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: two processes can be at some given distance $\varepsilon < 1$ only if they can mimic each other's transitions and evolve to distributions that are, in turn, at a distance $\leq \varepsilon$. Hence, for our purposes, we need to lift a pseudometric over processes to a pseudometric over distributions. We follow the approach of [15], that considers the *generalized Kantorovich lifting*. Take a generic metric space (V, d_V) , with $V \subseteq \mathbb{R}$ a convex subset. 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, [15]). 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 with respect to (V, d_V) is the pseudometric $\mathbf{K}_V(d_X): \Delta(X) \times \Delta(X) \rightarrow [0, +\infty]$ defined, for all $\pi, \pi' \in \Delta(X)$ by

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

Example 10 (The multiplicative variant of the Kantorovich lifting). The generalization of the Kantorovich lifting proposed in Definition 14 allows us to define the so called *multiplicative variant* of the Kantorovich lifting [15], which will play a fundamental role in the definition of the logical bound on d_X -privacy. To define the multiplicative variant, denoted by $\mathbf{K}_\otimes$, we consider $V = [0, 1]$ equipped with the metric $d_\otimes: [0, 1]^2 \rightarrow [0, 1]$ defined for all $x, y \in [0, 1]$ by

$$d_\otimes(x, y) = |\ln(x) - \ln(y)|$$

namely, $d_\otimes$ measures the euclidean distance between the natural logarithms of the reals in $[0, 1]$. $\square$

Generalized bisimulation metrics are then 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 $d_1 \preceq d_2$ if and only if $d_1(s, t) \leq d_2(s, t)$ 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 $s, t \in \mathcal{S}$.

The quantitative analogous to 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 [20, 27]. Intuitively, any difference that can be observed only after a long sequence of computation steps does not have the same impact of the differences that can be witnessed at the beginning of the computation. $\lambda = 1$ expresses no discount, so that the differences in the behavior of s and t are considered irrespective of after how many steps they can be observed.

Definition 15 (Generalized bisimulation metric functional, [15]). 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

$$\mathbf{B}_V(d)(s, t) = \begin{cases} \mathcal{O}_{d_V}(V) & \text{if } \text{init}(s) \neq \text{init}(t) \\ \sup_{a \in \mathcal{A}} \mathbf{H}(\lambda \cdot \mathbf{K}_V(d))(\text{der}(s, a), \text{der}(t, a)) & \text{otherwise.} \end{cases}$$

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 (cf. [15]).

Clearly, if the pseudometric d is not bounded, its Kantorovich lifting cannot be bounded.

We can show that $\mathbf{B}_V$ is monotone [15]. 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 [15].

Definition 16 (Generalized bisimulation metric, [15]). 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*.

Example 11. The classic bisimilarity metric $\mathbf{d}_\lambda$, is defined in terms of the Kantorovich metric. Since we are considering finitely supported distributions, one can prove that the Kantorovich-Rubinstein theorem can be applied to the metric space $(\mathcal{S}, \mathbf{d}_\lambda)$, and thus we can express the Kantorovich metric in its dual formulation, namely the infimum over the matchings for the distributions. Such formulation of the Kantorovich metric

Figure 4: The classic bisimilarity distance between s, t is $\mathbf{d}_\lambda(s, t) = 1/2 \cdot \lambda$.

is known as the *Wasserstein metric* and stems from optimal transport analysis (reason for which it is also known as the *Earth mover's distance*). Given two distributions π, π' a *matching* (or *coupling*, or *weight function*) for $\pi, \pi' \in \Delta(\mathcal{S})$ is a distribution over the product space $\mathfrak{w} \in \Delta(\mathcal{S} \times \mathcal{S})$ with π and π' as left and right marginal, namely: (i) $\sum_{t \in \mathcal{S}} \mathfrak{w}(s, t) = \pi(s)$, for all $s \in \mathcal{S}$, and (ii) $\sum_{s \in \mathcal{S}} \mathfrak{w}(s, t) = \pi'(t)$, for all $t \in \mathcal{S}$. We let $\mathfrak{W}(\pi, \pi')$ denote the set of all matchings for π and π' . Then, given any pseudometric d on $\mathcal{S}$, the Kantorovich lifting of d to $\Delta(\mathcal{S})$ is defined as

$$\mathbf{K}(d)(\pi, \pi') = \min_{\mathfrak{w} \in \mathfrak{W}(\pi, \pi')} \sum_{s, t \in \mathcal{S}} \mathfrak{w}(s, t) \cdot d(s, t)$$

for all $\pi, \pi' \in \Delta(\mathcal{S})$. Consider processes s, t in Figure 4. We aim at evaluating $\mathbf{d}_\lambda(s, t)$. Notice that $\mathbf{d}_\lambda(s_1, t_1) = \mathbf{d}_\lambda(s_3, t_3) = \mathbf{d}_\lambda(s_4, t_4) = 0$ whereas, since s_2 can perform both b and c , $\mathbf{d}_\lambda(s_2, t_2) = \mathbf{d}_\lambda(s_2, t_3) = 1$. Let $\pi_s = 1/2\delta_{s_2} + 1/2\delta_{s_3}$ and $\pi_t = 1/2\delta_{t_2} + 1/2\delta_{t_3}$, namely the distributions reached, respectively, by s and t via the execution of the central a action. Then we have

$$\begin{aligned} \mathbf{d}_\lambda(s, t) &= \lambda \cdot \max \begin{cases} \mathbf{K}(\mathbf{d}_\lambda)(\delta_{s_1}, \delta_{t_1}) \\ \mathbf{K}(\mathbf{d}_\lambda)(\pi_s, \pi_t) \\ \mathbf{K}(\mathbf{d}_\lambda)(\delta_{s_4}, \delta_{t_4}) \end{cases} \\ &= \lambda \cdot \min_{\mathfrak{w} \in \mathfrak{W}(\pi_s, \pi_t)} \sum_{s', t' \in \mathcal{S}} \mathfrak{w}(s', t') \cdot \mathbf{d}_\lambda(s', t') \\ &= \lambda \cdot (1/2 \cdot \mathbf{d}_\lambda(s_2, t_2) + 1/2 \cdot \mathbf{d}_\lambda(s_3, t_3)) \\ &= \lambda \cdot 1/2. \end{aligned}$$

Consider now the multiplicative variant of the Kantorovich metric introduced in Example 10 and let us evaluate the multiplicative bisimilarity distance $\mathbf{d}_{\lambda, \otimes}$ built on it, on the same processes s, t . Clearly, since $\mathbf{d}_{\lambda, \otimes}(s_2, t_2) = \mathbf{d}_{\lambda, \otimes}(s_2, t_3) = +\infty$, we obtain $\mathbf{K}_{\otimes}(\pi_s, \pi_t) = +\infty$ and thus $\mathbf{d}_{\lambda, \otimes}(s, t) = +\infty$. $\square$

6.3. Up-to k reasoning

We recall that on an image finite PTS, the bisimulation equivalence can be approximated by relations that consider only the first k transition steps [3, 41].

Definition 17 (Up-to- k bisimulation). Assume an image finite PTS. The family of the *up-to- k bisimulations* $\sim_k$, for $k \in \mathbb{N}$, is inductively defined as follows:

1. $\sim_0 = \mathcal{S} \times \mathcal{S}$;
2. $s \sim_{k+1} t$ if
 - whenever $s \xrightarrow{a} \pi_s$ there is a transition $t \xrightarrow{a} \pi_t$ such that $\pi_s \sim_k^\dagger \pi_t$;
 - whenever $t \xrightarrow{a} \pi_t$ there is a transition $s \xrightarrow{a} \pi_s$ such that $\pi_t \sim_k^\dagger \pi_s$.

Finally, we define $\sim_\omega = \bigcap_{k \geq 0} \sim_k$.

Proposition 5 ([41]). On image-finite PTSs, $\sim_\omega$ coincides with $\sim$.

Similarly, the bisimulation functional $\mathbf{B}_V$ allows us to define a notion of distance between processes that considers only the first k transition steps.

Definition 18 (Up-to- k bisimilarity metric). We define the *up-to- k bisimilarity metric* $\mathbf{d}_{\lambda,V}^k$ for $k \in \mathbb{N}$ by $\mathbf{d}_{\lambda,V}^k = \mathbf{B}_V^k(\mathbf{0})$.

Due to the continuity of the lifting functional $\mathbf{K}_V$ we can infer that also the functional $\mathbf{B}_V$ is continuous, besides monotone, thus ensuring that the closure ordinal of $\mathbf{B}_V$ is ω [55]. Hence, the up-to- k bisimilarity metrics converge to the bisimilarity metric when $k \rightarrow \infty$.

Proposition 6. Assume an image-finite PTS such that for each transition $s \xrightarrow{a} \pi$ we have that the probability distribution π has finite support. Then $\mathbf{d}_{\lambda,V} = \lim_{k \rightarrow \infty} \mathbf{d}_{\lambda,V}^k$.

Proof. The proof can be found in Appendix C. $\square$

7. The modal logic $\mathcal{L}$

We introduce the *modal logic* $\mathcal{L}$ of [23], which extends HML [40] with a probabilistic choice modality that allows us to express the behavior of probability distributions over processes.

Definition 19 (Modal logic $\mathcal{L}$, [23]). 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 20 (Semantics of $\mathcal{L}$, [23]). 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 21 ([11]). We define $E: \mathcal{P}(\mathcal{L}^s \times \mathcal{L}^s) \rightarrow \mathcal{P}(\mathcal{L}^s \times \mathcal{L}^s)$ as the functional such that for all relations $\mathcal{R} \in \mathcal{P}(\mathcal{L}^s \times \mathcal{L}^s)$ we have that $E(\mathcal{R})$ is the greatest relation satisfying:

1. $(\varphi, \varphi) \in E(\mathcal{R})$;
2. $(\varphi', \varphi) \in E(\mathcal{R})$ iff $(\varphi, \varphi') \in E(\mathcal{R})$;
3. $(\neg\varphi_1, \neg\varphi_2) \in E(\mathcal{R})$ iff $(\varphi_1, \varphi_2) \in E(\mathcal{R})$;
4. $(\bigwedge_{j \in \mathcal{J}} \varphi_j, \bigwedge_{j \in (\mathcal{J} \setminus \mathcal{I})} \varphi'_j) \in E(\mathcal{R})$ for some $\mathcal{I} \neq \emptyset, \mathcal{I} \subset \mathcal{J}$ iff
 - for each $j \in \mathcal{J} \setminus \mathcal{I}$ we have $(\varphi_j, \varphi'_j) \in E(\mathcal{R})$,
 - for each $i \in \mathcal{I}$ we have $(\varphi_i, \varphi'_{j_i}) \in E(\mathcal{R})$ for some $j_i \in \mathcal{J} \setminus \mathcal{I}$;
5. $(\bigwedge_{j \in \mathcal{J}} \varphi_j, \bigwedge_{i \in \mathcal{I}} \varphi_i) \in E(\mathcal{R})$ iff there is a bijection $f: \mathcal{J} \rightarrow \mathcal{I}$ with $(\varphi_j, \varphi_{f(j)}) \in E(\mathcal{R})$ for all $j \in \mathcal{J}$;
6. $(\langle a \rangle \psi, \langle a \rangle \psi') \in E(\mathcal{R})$ iff $\psi \mathcal{R}^\dagger \psi'$.

We briefly explain the six conditions in Definition 21: 1. Items 1 and 2 ensure, respectively, that $E(\mathcal{R})$ is a reflexive and symmetric relation. 2. Item 3 extends $E(\mathcal{R})$ to negation by stating that whenever φ_1 and φ_2 are related by $E(\mathcal{R})$, then also their negations are related by $E(\mathcal{R})$. 3. Item 4 establishes the relation among two formulae defined as conjunctions on the same set of formulae: it states that if we delete multiple copies of the same formula or copies of formulae that are related to other formulae already occurring in the conjunction, then we obtain a formula which is related by $E(\mathcal{R})$ to the original formula. For instance, item 4 allows us to infer that $(\varphi \wedge \varphi, \varphi) \in E(\mathcal{R})$. 4. Item 5 establishes the relation among two formulae defined as conjunctions on two different sets of formulae: $\bigwedge_{j \in \mathcal{J}} \varphi_j$ and $\bigwedge_{i \in \mathcal{I}} \varphi_i$ are in $E(\mathcal{R})$ if and only if the formulae φ_j and φ_i are in turn related by $E(\mathcal{R})$ two by two. 5. Item 6 states that whenever two distribution formulae ψ, ψ' are related by the (lifted) ground relation $\mathcal{R}$ then $E(\mathcal{R})$ relates any pair of equally labeled diamond modalities having, respectively, ψ and ψ' in the scope.

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

Definition 22 ($\mathcal{L}$ -equivalence). The $\mathcal{L}$ -equivalence of formulae $\equiv_{\mathcal{L}} \subseteq \mathcal{L}^s \times \mathcal{L}^s$ is defined as

$$\equiv_{\mathcal{L}} = \max\{\mathcal{R} \subseteq \mathcal{L}^s \times \mathcal{L}^s \mid \mathcal{R} \subseteq E(\mathcal{R})\}.$$

7.1. The mimicking formulae

In [23] it was proved that the logic $\mathcal{L}$ is *adequate* for bisimilarity, i.e., two processes are bisimilar if and only if they satisfy the same formulae in $\mathcal{L}$. The drawback of this valuable result is 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 [25] a characterization of bisimilarity was given in terms of *characteristic formulae* of processes, i.e., particular formulae that alone capture the entire equivalence class of the related process [38]: if ϕ_s is the characteristic formula of process s for bisimilarity, then $s \sim t$ if and only if $t \models \phi_s$. 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, [25] 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, [11–13] proposed a different technique for the characterization. When we compare the behavior of two processes, we compare those properties that are observable for them with respect to 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. [11–13] showed that semantic equivalence of processes holds if and only if their mimicking formulae are syntactically equivalent (Theorem 4 below). Hence, to establish process equivalence we need to compare only two formulae. Moreover, the logic on which the mimicking formulae are constructed is always *minimal* with respect to the chosen semantics, i.e., it only includes the operators necessary to express the observable properties with respect to that semantics.

Here, we recall the definition of mimicking formula and the *weak expressive* characterization of bisimilarity from [11]. 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 23 (Mimicking formula, [11]). 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 12. Consider process s in Figure 4 and assume that $\mathcal{A} = \{a, b, c\}$. We aim at constructing the mimicking formula of s . We have

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

□

Mimicking formulae allow us to characterize probabilistic bisimilarity.

Theorem 4 ([11]). *Given any $s, t \in \mathcal{S}$ we have that $\varphi_s \equiv_{\mathcal{L}} \varphi_t$ iff $s \sim t$.*

8. Logical characterization of generalized bisimulation 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 Section 6.2. The idea follows that of [11, 12]:

- (i) 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.
- (ii) 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.
- (iii) We show that the logical distance $\ell_{\lambda, V}$ coincides with the bisimilarity metric $\mathbf{d}_{\lambda, V}$ (Theorem 6).

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

Definition 24 (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, whose existence is guaranteed by the following two results.

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

Proof. The proof can be found in Appendix D. $\square$

Proposition 7. 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.

Proof. The proof can be found in Appendix D. $\square$

For a complete presentation of our generalized syntactic distances over $\mathcal{L}$, we remark that each of them is a pseudometric whose kernel is $\mathcal{L}$ -equivalence. However, the proof of the latter property requires an additional restriction on the metric space (V, d_V) on which the distance is built: we need that (V, d_V) has a geodesic, which can be informally seen as the shortest path between two points in a curved space. Notice that such requirement is the same imposed in [15] to obtain that the kernel of any generalized bisimilarity metric is $\sim$, and thus it does not compromise the relevance of our results.

Proposition 8. If (V, d_V) has a geodesic, then $\equiv_{\mathcal{L}} = \ker(\mathfrak{d}_{\lambda, V})$.

Proof. The proof can be found in Appendix D. $\square$

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

Definition 25 (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 up-to- k generalized bisimilarity metrics in terms of the up-to- k logical distances over processes.

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

Proof. The proof is by induction on $k \in \mathbb{N}$ and can be found in Appendix D. $\square$

From the characterization of the up-to- k metrics, we can derive the logical characterization of generalized bisimilarity metric.

Theorem 6. *Let $\lambda \in (0, 1]$. For any $s, t \in \mathcal{S}$ we have $\ell_{\lambda, V}(s, t) = \mathbf{d}_{\lambda, V}(s, t)$.*

Proof. Let us prove now that

$$\ell_{\lambda, V}(s, t) = \mathbf{d}_{\lambda, V}(s, t). \quad (2)$$

We have

$$\begin{aligned} & \ell_{\lambda, V}(s, t) \\ &= \mathfrak{d}_{\lambda, V}(\varphi_s, \varphi_t) && \text{(by definition of } \ell_{\lambda, V} \text{ (Definition 25))} \\ &= \lim_{k \rightarrow \infty} \mathfrak{d}_{\lambda, V}^k(\varphi_s^k, \varphi_t^k) && \text{(by def. of } \mathfrak{d}_{\lambda, V} \text{ (Definition 7) and def. of } \varphi_s \text{ (Definition 23))} \\ &= \lim_{k \rightarrow \infty} \mathbf{d}_{\lambda, V}^k(s, t) && \text{(by Theorem 5)} \\ &= \mathbf{d}_{\lambda, V}(s, t) && \text{(by Proposition 6).} \end{aligned}$$

which gives Equation 2 and concludes the proof. $\square$

9. 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:

- (i) 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 probability measures over traces induced by them.
- (ii) 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 [15], 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 $\otimes_V(V)$ -valued discrete metric over $\mathcal{A}^*$ which is defined as $\mathbf{dm}_{\otimes_V}(\alpha, \alpha') = 0$ if $\alpha = \alpha'$ and $\mathbf{dm}_{\otimes_V}(\alpha, \alpha') = \otimes_V(V)$ otherwise. Then let $\mathbf{K}_{\otimes}$ be the multiplicative variant of the Kantorovich lifting introduced in Example 10. In [15] it has been proved that for $\otimes_{d_{\otimes}}([0, 1]) = +\infty$ it holds $tv_{\otimes} = \mathbf{K}_{\otimes}(\mathbf{dm}_{\otimes})$. Hence, from $\mathbf{d}_{\lambda, \otimes} \geq \mathbf{K}_{\otimes}(\mathbf{dm}_{\otimes})$ (cf. [15]) and Theorem 6 we obtain the following result.

Proposition 9. *Assume a LMC and let s, t be two processes in it. Let $\pi_s = \text{Pr}(s, \cdot)$ and $\pi_t = \text{Pr}(t, \cdot)$. Then $tv_{\otimes}(\pi_s, \pi_t) \leq \ell_{\otimes}(s, t)$.*

Proof. The proof can be found in Appendix E. $\square$

We remark that Proposition 2, Theorem 6 and Proposition 9 imply that $\mathbf{d}_{\otimes}^T \preceq \mathbf{d}_{\otimes}$.

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

Theorem 7 (Logical bound on $d_{\mathcal{X}}$ -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*

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

Proof. In [15, Theorem 3] it was proved that M is ε -differentially private if $\mathbf{d}_{\lambda, \otimes}(s_{\mathbf{x}}, s_{\mathbf{x}'}) \leq \varepsilon$ for all adjacent q, q' . Thus, the proof follows by this result and our characterization Theorem 6. $\square$

The following example illustrates a case of standard differential privacy.

Example 13. We recall that we are using an alternative notion of privacy in which all the databases have the same number of records n , and where the absence of a record is represented by a special value. Consider two medical databases x and x' , both of size n , 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 [35], 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 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.

$\square$

10. Related work and extensions

In this section we discuss related work and possible extensions of our results. In particular, we compare the expressive power of the modal logic $\mathcal{L}$ to that of *PHML* [46], a probabilistic variant of HML usually used in the literature to provide the characterization of probabilistic bisimilarity. We show that $\mathcal{L}$ is more expressive thus assessing its choice in our characterization technique. Moreover, we consider the problem of adding recursion to our framework. We argue that this extension would be made possible by exploiting the *equational μ -calculus* approach [1, 45, 51].

10.1. Related work

First of all we remark that, as already mentioned, this paper builds on the work of [11, 15]. The main novelty is in that we develop a technique for characterizing privacy properties, and that we deal with $d_{\mathcal{X}}$ -privacy rather than DP. We stress that ours are the first proposal of a logical characterization of $d_{\mathcal{X}}$ -privacy and weak anonymity.

We can, however, find logical formulations of the verification tools for DP properties [4–6] mainly based on Hoare like logics and the tools and theorem provers constructed for them. In particular, such formulations are based on *apRHL* (approximated probabilistic relational Hoare logic) a relational logic designed to support differentially private computations. By means of apRHL the authors were able to verify and provide accurate bounds on the (ε, δ) -differential privacy of some randomized mechanisms. Interestingly, in [5] apRHL is combined with probabilistic coupling techniques to obtain novel techniques for privacy verification.

Moreover, prominent approaches to the verification of DP that are based on type systems have been proposed [34, 50].

To obtain our characterization we have exploited a reformulation of $d_{\mathcal{X}}$ -privacy in terms of the bisimulation metrics defined via the multiplicative Kantorovich lifting.

In the recent paper [17] the problem of verifying differential privacy in concurrent systems was considered. More precisely, the authors considered (ε, δ) -DP, and, assuming ε known, they established a bound on δ by computing a bisimilarity metric based on the (standard) Kantorovich distance, and provided an analysis of the complexity.

Earlier papers [54, 57] also defined bisimulation metrics suitable for proving DP. Briefly, in [54] the authors consider the model of probabilistic I/O automata on which they use a notion of *differential noninterference* to sanitize data. The level of privacy is measured via the notions of *δ -approximated lifting* and *ε -unwinding relation*. The former allows for lifting a relation over the states of the automaton to probability distributions over them. The latter extends the concept of unwinding relation usually used for noninterference [37] to a sound relation for DP.

In [57], by using notions of approximated lifting inspired by that in [54], two behavioral metrics are presented and shown to preserve differential privacy properties on probabilistic automata: the *accumulative pseudometric* and the *amortized pseudometric*. Although the kernels of both metrics induce a sort of ε -bisimulation equivalence on the states of automata, they both suffer from the fact that such relations do not fully characterize probabilistic bisimilarity.

10.2. The expressive power of $\mathcal{L}$

In the seminal work [46], probabilistic bisimilarity was defined as the equivalence resulting from testing processes on modal formulae in PHML, namely HML equipped with a quantitative version of the diamond operator $\langle \cdot \rangle_r$. Briefly, a process s satisfies the formula $\langle a \rangle_r \phi$ if and only if there a distribution π such that $s \xrightarrow{a} \pi$ and $\pi(\{s' \mid s' \models \phi\}) \geq r$. Basically, the formula $\langle a \rangle_r \phi$ imposes a lower bound r on the total probability of a process to evolve into a process satisfying ϕ via the execution of action a . The same logic was then equipped with a real-valued semantics and used in [27] to define the bisimilarity metric.

The reason why we decided to use the modal logic $\mathcal{L}$ in place of PHML is to be found in its expressive power. We can easily prove that $\mathcal{L}$ is more expressive than PHML, in the sense that all formulae in PHML can be expressed by formulae in $\mathcal{L}$, but the converse is not true. As the difference in the two classes of formulae is in the diamond operator, we only outline this case.

Figure 5: Processes s, t show that $\mathcal{L}$ is more expressive than PHML.

Given any formula PHML ϕ , we let $\phi^{\mathcal{L}}$ denote the rewriting of ϕ as a formula in $\mathcal{L}$. Symmetrically, for any formula $\varphi \in \mathcal{L}$, we let φ^{PHML} denote the rewriting of ϕ as a formula in PHML. Then, one can easily prove that, given any process $s \in \mathcal{S}$, any formula ϕ in PHML and any state formula $\varphi \in \mathcal{L}^s$:

1. $s \models \langle a \rangle_r \phi$ if and only if $s \models \langle a \rangle [r\phi^{\mathcal{L}} \oplus (1-r)\top]$;
2. $s \models \langle a \rangle \bigoplus_{i \in I} r_i \varphi_i$ implies $s \models \bigwedge_{i \in I} \langle a \rangle_{r_i} \varphi_i^{\text{PHML}}$;
3. $s \models \bigwedge_{i \in I} \langle a \rangle_{r_i} \varphi_i^{\text{PHML}}$ does not imply $s \models \langle a \rangle \bigoplus_{i \in I} r_i \varphi_i$.

The first two items are immediate. To give more insights into the third item, consider process t in Figure 5. Clearly, we have that $t \models \langle a \rangle_{1/2} \langle b \rangle_0 \top \wedge \langle a \rangle_{1/2} \langle c \rangle_0 \top$ since process t_2 satisfies both $\langle b \rangle_0 \top$ and $\langle c \rangle_0 \top$ and can thus contribute to reach the required probability bound in both formulae in the conjunction. However, we have $t \not\models \langle a \rangle (1/2 \langle b \rangle \top \oplus 1/2 \langle c \rangle \top)$. In fact, process t_2 still satisfies both formulae $\langle b \rangle \top$ and $\langle c \rangle \top$ but its the probability weight, which is $1/4$, can either be used to reach the required probability on only one of the two formulae in the $\oplus$, or it can be split between them. In the former case only half of the probabilistic choice is satisfied; whereas in the latter no part of the formula is satisfied.

Interestingly, this difference leads to a more consistent disparity in the expressiveness of the two classes of formulae. In fact, we have that $\mathcal{L}$ without negation characterizes *probabilistic simulation*, namely the asymmetric version of probabilistic bisimulation, whereas PHML without negation is not powerful enough to do so and we must add the disjunction modality to obtain such a characterization [26].

To see this, consider processes s, t in Figure 5. We have that neither s is simulated by t nor vice versa. Accordingly, the two processes are distinguished by the $\mathcal{L}$ -formula $\langle a \rangle (1/2 \langle b \rangle \top \oplus 1/2 \langle c \rangle \top)$ which is satisfied by s but not by t , as discussed above. However, no formula in PHML can distinguish s from t . In particular, we notice that both processes satisfy the PHML-formula $\langle a \rangle_{1/2} \langle b \rangle_0 \top \wedge \langle a \rangle_{1/2} \langle c \rangle_0 \top$.

10.3. How to treat recursion

We would like to stress that our characterization technique can be applied alto to processes with recursion. We could in fact apply the *equational μ -calculus* approach of [45] (later generalized in [1, 51]) which allows us to define the fixed point semantics of formulae without introducing the fix-point operators in the syntax of formulae. As one can expect, this is a great advantage from the point of view of our characterization technique based on a *syntactic* distance on formulae. Since the development of the equational μ -calculus in our setting would be technically involved, we decided not to present it in this paper, thus favoring a simplified presentation of the characterization technique and of its application to privacy. Still, for sake of the interested reader, we dedicate this paragraph to a brief description of how recursion could be dealt with. We refer the interested reader to [13] for a detailed presentation of the following technique.

First of all we notice that even in the case of processes with recursion we would still require 1. the image-finiteness hypothesis, since without it defining probabilistic bisimilarity as the limit of its approximations would not be possible [41] and 2. the use of finitely supported probability distributions, in order to guarantee the continuity of the bisimulation metric functional (cf. Section 6.3). Then, we can focus on bisimulation semantics and the equational μ -calculus approach. In detail, we extend the logic $\mathcal{L}$ to a *modal $\mathcal{S}$ -indexed logic* by adding the $\mathcal{S}$ -indexed family of variables $\{X_s \mid s \in \mathcal{S}\}$. Intuitively, these variables allow for a

recursive specification of modal properties. Then, an appropriate *interpretation* (called *model* in [45]) to each variable is provided as the solution of a system of equations obtained by means of an *endodeclaration*, which is a function $\mathcal{E}$ mapping variables to formulae of the logic. More specifically, $\mathcal{E}$ allows us to implicitly define a system of equations

$$\gamma(X) = \llbracket \mathcal{E}(X) \rrbracket_\gamma \quad (3)$$

whose solution will correspond to the *proper* variable interpretation for the formula: an interpretation γ is a solution for the system (3) if the semantics of X under γ corresponds to the interpretation of the formula $\mathcal{E}(X)$ assigned to X by $\mathcal{E}$. As *solution* of the system we consider the variable interpretation corresponding to the *greatest fixed point* of the system. We remark that also the notion of equivalence of formulae, which is a *syntactic* equivalence, would depend on the endodeclaration $\mathcal{E}$, and thus be defined as the fixed-point of a proper Scott-continuous functional $\mathcal{F}_\mathcal{E}$. Finally, we assign to each process s the related mimicking formula φ_s defined as $\mathcal{E}(X_s)$. Once we have obtained the mimicking formulae, we can easily regain all the results of this paper.

11. Conclusions

We have provided a logical characterization of privacy properties based on behavioral metrics. The underlying idea is quite simple: (i) We consider a boolean modal logic powerful enough to express the desired semantics; (ii) We define a syntactic distance over the formulae in the chosen logic; (iii) We express the differences in the behavior of processes in terms of such distance. We have shown that this technique applied in the case of trace semantics on LMCs results in a logical characterization of $d_\mathcal{X}$ -privacy. Moreover, we have also shown how it is possible to define a real-valued semantics for formulae starting from the syntactic distance on them and we have exploited this result to obtain a logical characterization of weak probabilistic anonymity. Then we have switched to bisimulation semantics and we have provided a logical characterization of generalized bisimulation metrics. Beside the syntactic distance on formulae, this characterization is based on the notion of mimicking formula of a process, namely a special formula that captures the observable behavior of that process. We have shown that from the characterization of generalized bisimulation metrics we can obtain bounds on $d_\mathcal{X}$ -privacy properties.

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 [27, 28] 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.

It would also be interesting to apply our characterization method to notions of DP which are not related to the semantics of LMCs. For instance, we could consider the cases of Rényi DP (RDP) [48] and Gaussian DP (GDP) [29]. These notions of DP are based on a comparison of the outputs of randomized mechanisms obtained, respectively, via a *Kullback-Leibler-like distance*, for RDP, and a *trade-off function*, for GDP. Despite there is no relation between these and the semantics of LMCs, we strongly believe that our technique could also be applied to them. The idea behind their characterization would be the same: find a class of modal formulae which allows for expressing the desired properties, and then define a suitable metric over that class.

Finally, we aim at developing quantitative analysis techniques and tools for proving privacy properties. In particular, it could be interesting to write an algorithm for the evaluation of (approximated) weak anonymity based on our characterization of it in terms of a total variation distance on the values of formulae and the results in [43] on the complexity of evaluating (approximated) total variation distances on Markov processes.

Acknowledgements We thank the reviewers for their careful reading of the paper and their constructive feedback.

This paper has been partially funded by the ANR project REPAS. The work of Catuscia Palamidessi has been funded by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme, Grant agreement nr. 835294.

The work of Valentina Castiglioni has been partially supported by the project ‘Open Problems in the Equational Logic of Processes’ (OPEL) of the Icelandic Research Fund (grant nr. 196050-051).

- [1] Aceto, L., Ingólfssdóttir, A., Levy, P. B., Sack, J., 2012. Characteristic formulae for fixed-point semantics: a general framework. *Mathematical Structures in Computer Science* 22 (2), 125–173.
- [2] Bacci, G., Bacci, G., Larsen, K. G., Mardare, R., 2015. Converging from branching to linear metrics on Markov chains. In: *Proceedings of ICTAC 2015*. Vol. 9399 of *Lecture Notes in Computer Science*. pp. 349–367.
- [3] Baier, C., 1998. On algorithmic verification methods for probabilistic systems. Ph.D. thesis, Fakultät für Mathematik und Informatik Universität Mannheim.
- [4] Barthe, G., Gaboardi, M., Arias, E. J. G., Hsu, J., Kunz, C., Strub, P., 2014. Proving differential privacy in hoare logic. In: *Proc. of CSF 2014*. pp. 411–424.
URL <https://doi.org/10.1109/CSF.2014.36>
- [5] Barthe, G., Gaboardi, M., Grégoire, B., Hsu, J., Strub, P., 2016. Proving differential privacy via probabilistic couplings. In: *Proc. of LICS ’16*. pp. 749–758.
URL <https://doi.org/10.1145/2933575.2934554>
- [6] Barthe, G., Köpf, B., Olmedo, F., Béguelin, S. Z., 2012. Probabilistic relational reasoning for differential privacy. In: *Proc. of POPL*. ACM.
- [7] Bernardo, M., De Nicola, R., Loret, M., 2014. Revisiting trace and testing equivalences for nondeterministic and probabilistic processes. *Logical Methods in Computer Science* 10 (1).
- [8] Bloom, B., Fokink, W. J., van Glabbeek, R. J., 2004. Precongruence formats for decorated trace semantics. *ACM Trans. Comput. Log.* 5 (1), 26–78.
- [9] Castiglioni, V., 2018. Trace and testing metrics on nondeterministic probabilistic processes. In: *Proceedings of EXPRESS/SOS 2018*. Vol. 276 of *EPTCS*. pp. 19–36.
URL <https://doi.org/10.4204/EPTCS.276.4>
- [10] Castiglioni, V., Chatzikokolakis, K., Palamidessi, C., 2018. A logical characterization of differential privacy via behavioral metrics. In: *Proceedings of FACS 2018*. Vol. 11222 of *LNCS*. pp. 75–96.
URL https://doi.org/10.1007/978-3-030-02146-7_4
- [11] Castiglioni, V., Gebler, D., Tini, S., 2016. Logical characterization of bisimulation metrics. In: *Proceedings of QAPL’16*. Vol. 227 of *EPTCS*. pp. 44–62.
- [12] Castiglioni, V., Tini, S., 2017. Logical characterization of trace metrics. In: *Proceedings of QAPL@ETAPS 2017*. Vol. 250 of *EPTCS*. pp. 39–74.
- [13] Castiglioni, V., Tini, S., 2019. Logical characterization of branching metrics for nondeterministic probabilistic transition systems. *Inf. Comput.* 268.
URL <https://doi.org/10.1016/j.ic.2019.06.001>
- [14] Chatzikokolakis, K., Andrés, M. E., Bordenabe, N. E., Palamidessi, C., 2013. Broadening the scope of differential privacy using metrics. In: *Proceedings of PETS 2013*. Vol. 7981 of *LNCS*. pp. 82–102.
- [15] Chatzikokolakis, K., Gebler, D., Palamidessi, C., Xu, L., 2014. Generalized bisimulation metrics. In: *Proceedings of CONCUR 2014*. Vol. 8704 of *Lecture Notes in Computer Science*. pp. 32–46.
- [16] Chaum, D., 1988. The dining cryptographers problem: Unconditional sender and recipient untraceability. *J. Cryptology* 1 (1), 65–75.
URL <https://doi.org/10.1007/BF00206326>
- [17] Chistikov, D., Murawski, A. S., Purser, D., 2019. Asymmetric distances for approximate differential privacy. In: *Proceedings of CONCUR 2019*. Vol. 140 of *LIPIcs*. pp. 10:1–10:17.
URL <https://doi.org/10.4230/LIPIcs.CONCUR.2019.10>
- [18] Daga, P., Henzinger, T. A., Křetínský, J., Petrov, T., 2016. Linear distances between Markov chains. In: *Proceedings of CONCUR 2016*. Vol. 59 of *LIPIcs*. pp. 20:1–20:15.
- [19] de Alfaro, L., Faella, M., Stoelinga, M., 2009. Linear and branching system metrics. *IEEE Trans. Software Eng.* 35 (2), 258–273.
- [20] de Alfaro, L., Henzinger, T. A., Majumdar, R., 2003. Discounting the Future in Systems Theory. In: *Proceedings of ICALP’03*. Vol. 2719 of *Lecture Notes in Computer Science*. pp. 1022–1037.
- [21] de Alfaro, L., Majumdar, R., Raman, V., Stoelinga, M., 2008. Game refinement relations and metrics. *Logical Methods in Computer Science* 4 (3).
- [22] Deng, Y., Chothia, T., Palamidessi, C., Pang, J., 2006. Metrics for action-labelled quantitative transition systems. *Electr. Notes Theor. Comput. Sci.* 153 (2), 79–96.
- [23] Deng, Y., Du, W., 2011. Logical, metric, and algorithmic characterisations of probabilistic bisimulation. *CoRR* abs/1103.4577.
URL <http://arxiv.org/abs/1103.4577>
- [24] Deng, Y., Palamidessi, C., Pang, J., 2007. Weak probabilistic anonymity. *ENTCS* 180 (1), 55–76.

- [25] Deng, Y., van Glabbeek, R. J., 2010. Characterising probabilistic processes logically - (extended abstract). In: Proceedings of LPAR-17. pp. 278–293.
- [26] Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P., 2003. Approximating labelled Markov processes. *Inf. Comput.* 184 (1), 160–200.
URL [https://doi.org/10.1016/S0890-5401\(03\)00051-8](https://doi.org/10.1016/S0890-5401(03)00051-8)
- [27] Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P., 2004. Metrics for labelled Markov processes. *Theor. Comput. Sci.* 318 (3), 323–354.
- [28] Desharnais, J., Jagadeesan, R., Gupta, V., Panangaden, P., 2002. The metric analogue of weak bisimulation for probabilistic processes. In: Proceedings of LICS 2002. pp. 413–422.
- [29] Dong, J., Roth, A., Su, W. J., 2019. Gaussian differential privacy. CoRR abs/1905.02383.
URL <http://arxiv.org/abs/1905.02383>
- [30] Du, W., Deng, Y., Gebler, D., 2016. Behavioural pseudometrics for nondeterministic probabilistic systems. In: Proceedings of SETTA 2016. Vol. 9984 of Lecture Notes in Computer Science. pp. 67–84.
- [31] Duchi, J. C., Jordan, M. I., Wainwright, M. J., 2013. Local privacy and statistical minimax rates. In: Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS). IEEE Computer Society, pp. 429–438.
- [32] Dwork, C., 2006. Differential privacy. In: Proceedings of ICALP 2006. Vol. 4052 of LNCS. pp. 1–12.
- [33] Erlingsson, Ú., Pihur, V., Korolova, A., 2014. 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). ACM, pp. 1054–1067.
- [34] Gaboardi, M., Haeberlen, A., Hsu, J., Narayan, A., Pierce, B. C., 2013. Linear dependent types for differential privacy. In: POPL. pp. 357–370.
- [35] Ghosh, A., Roughgarden, T., Sundararajan, M., 2009. Universally utility-maximizing privacy mechanisms. In: Proceedings of the 41st annual ACM Symposium on Theory of Computing (STOC). ACM, pp. 351–360.
- [36] Giacalone, A., Jou, C.-C., Smolka, S. A., 1990. Algebraic reasoning for probabilistic concurrent systems. In: Proceedings of IFIP Work, Conf. on Programming, Concepts and Methods. pp. 443–458.
- [37] Goguen, J. A., Meseguer, J., 1984. Unwinding and inference control. In: Proc. of the 1984 IEEE Symposium on Security and Privacy. pp. 75–87.
URL <https://doi.org/10.1109/SP.1984.10019>
- [38] Graf, S., Sifakis, J., 1986. A modal characterization of observational congruence on finite terms of CCS. *Information and Control* 68 (1-3), 125–145.
- [39] Hansson, H., Jonsson, B., 1994. A logic for reasoning about time and reliability. *FAC* 6 (5), 512–535.
- [40] Hennessy, M., Milner, R., 1985. Algebraic laws for nondeterminism and concurrency. *J. Assoc. Comput. Mach.* 32, 137–161.
- [41] Hermanns, H., Parma, A., Segala, R., Wachter, B., Zhang, L., 2011. Probabilistic logical characterization. *Inf. Comput.* 209 (2), 154–172.
- [42] Keller, R. M., 1976. Formal verification of parallel programs. *Commun. ACM* 19 (7), 371–384.
- [43] Kiefer, S., 2018. On computing the total variation distance of hidden markov models. In: Proceedings of ICALP 2018. Vol. 107 of LIPIcs. pp. 130:1–130:13.
URL <https://doi.org/10.4230/LIPIcs.ICALP.2018.130>
- [44] Kwiatkowska, M. Z., Norman, G., 1996. Probabilistic metric semantics for a simple language with recursion. In: Proceedings of MFCS’96. Vol. 1113 of Lecture Notes in Computer Science. pp. 419–430.
- [45] Larsen, K. G., 1990. Proof systems for satisfiability in Hennessy-Milner logic with recursion. *Theor. Comput. Sci.* 72 (2&3), 265–288.
URL [http://dx.doi.org/10.1016/0304-3975\(90\)90038-J](http://dx.doi.org/10.1016/0304-3975(90)90038-J)
- [46] Larsen, K. G., Skou, A., 1991. Bisimulation through probabilistic testing. *Inf. Comput.* 94 (1), 1–28.
- [47] Machanavajjhala, A., Kifer, D., Abowd, J. M., Gehrke, J., Vilhuber, L., 2008. Privacy: Theory meets practice on the map. In: Proceedings of ICDE 2008. pp. 277–286.
- [48] Mironov, I., 2017. Rényi differential privacy. In: Proceedings of CSF 2017. pp. 263–275.
URL <https://doi.org/10.1109/CSF.2017.11>
- [49] Narayanan, A., Shmatikov, V., 2009. De-anonymizing social networks. In: Proceedings of S&P 2009. pp. 173–187.
- [50] Reed, J., Pierce, B. C., 2010. Distance makes the types grow stronger: a calculus for differential privacy. In: Proc. of ICFP. ACM, pp. 157–168.
- [51] Sack, J., Zhang, L., 2012. A general framework for probabilistic characterizing formulae. In: Proceedings of VMCAI 2012. pp. 396–411.
- [52] Segala, R., 1995. Modeling and verification of randomized distributed real-time systems. Ph.D. thesis, MIT.
- [53] Song, L., Deng, Y., Cai, X., 2007. Towards automatic measurement of probabilistic processes. In: Proceedings of QSI 2007. pp. 50–59.
- [54] Tschantz, M. C., Kaynar, D., Datta, A., sep 2011. Formal verification of differential privacy for interactive systems (extended abstract). *ENTCS* 276, 61–79.
- [55] van Breugel, F., 2012. On behavioural pseudometrics and closure ordinals. *Inf. Process. Lett.* 112 (19), 715–718.
- [56] van Breugel, F., Worrell, J., 2005. A behavioural pseudometric for probabilistic transition systems. *Theor. Comput. Sci.* 331 (1), 115–142.
- [57] Xu, L., Chatzikokolakis, K., Lin, H., 2014. Metrics for differential privacy in concurrent systems. In: Proc. of FORTE. Vol. 8461 of LNCS. Springer, pp. 199–215.

Appendix A. Proofs of results in Section 4

Proof of Theorem 1. The implication $\mathbb{L}(s) = \mathbb{L}(t) \Rightarrow s \sim_{\text{Tr}} t$ can be proved by an easy induction over the structure of formulae in $\mathbb{L}$.

For $s \sim_{\text{Tr}} t \Rightarrow \mathbb{L}(s) = \mathbb{L}(t)$, we have

$$\begin{aligned} s \sim_{\text{Tr}} t &\text{ iff } \forall \alpha \in \mathcal{A}^*: \Pr(s, \alpha) = \Pr(t, \alpha) \\ &\text{ iff } \forall \Phi \in \mathbb{L}^1: \Pr(s, \text{Tr}(\Phi)) = \Pr(t, \text{Tr}(\Phi)) \\ &\text{ iff } \forall \Phi \in \mathbb{L}^1: \Pr(s, \text{Tr}(\Phi))\Phi \in \mathbb{L}(s) \Leftrightarrow \Pr(s, \text{Tr}(\Phi))\Phi \in \mathbb{L}(t) \\ &\text{ iff } \mathbb{L}(s) = \mathbb{L}(t). \end{aligned}$$

□

Proof of Proposition 1. The proof follows by noticing that since $\mathfrak{d}_{\mathcal{D}}^{\text{P}}$ is a metric over $\mathbb{L}^{\text{P}}$ and $\mathbb{L}(s)$ and $\mathbb{L}(t)$ are closed sets wrt. the topology induced by it, then $\mathbf{H}(\mathfrak{d}_{\mathcal{D}}^{\text{P}})(\mathbb{L}(s), \mathbb{L}(t)) = 0$ if and only if $\mathbb{L}(s) = \mathbb{L}(t)$. □

Proof of Proposition 2. In what follows, for each $\alpha \in \mathcal{A}^*$, let Φ_{α} denote the only formula in $\mathbb{L}^1$ such that $\text{Tr}(\Phi_{\alpha}) = \alpha$. We have

$$\begin{aligned} &\mathbf{d}_{\otimes}^T(s, t) \\ &= \mathbf{H}(\mathfrak{d}_{\otimes}^{\text{P}})(\mathbb{L}(s), \mathbb{L}(t)) \\ &= \max \left\{ \sup_{\Psi_s \in \mathbb{L}(s)} \inf_{\Psi_t \in \mathbb{L}(t)} \mathfrak{d}_{\otimes}^{\text{P}}(\Psi_s, \Psi_t); \sup_{\Psi_t \in \mathbb{L}(t)} \inf_{\Psi_s \in \mathbb{L}(s)} \mathfrak{d}_{\otimes}^{\text{P}}(\Psi_s, \Psi_t) \right\} \\ &= \sup_{\Psi_s \in \mathbb{L}(s)} \inf_{\Psi_t \in \mathbb{L}(t)} \mathfrak{d}_{\otimes}^{\text{P}}(\Psi_s, \Psi_t) \quad (\text{assume wlog.}) \\ &= \sup_{\Phi_{\alpha} \in \mathbb{L}^1} \inf_{\Phi_{\beta} \in \mathbb{L}^1} \mathfrak{d}_{\otimes}^{\text{P}}(\Pr(s, \alpha)\Phi_{\alpha}, \Pr(t, \beta)\Phi_{\beta}) \quad (\text{by def. of } \mathbb{L}^{\text{P}}) \\ &= \sup_{\alpha \in \mathcal{A}^*} \mathfrak{d}_{\otimes}^{\text{P}}(\Pr(s, \alpha)\Phi_{\alpha}, \Pr(t, \alpha)\Phi_{\alpha}) \quad (\mathfrak{d}_{\otimes}^{\text{P}} \leq \otimes_{\otimes}([0, 1])) \\ &= \sup_{\alpha \in \mathcal{A}^*} |\ln(\Pr(s, \alpha)) - \ln(\Pr(t, \alpha))| \quad (\text{by def. of } \mathfrak{d}_{\otimes}^{\text{P}}) \\ &= tv_{\otimes}(\mu_s, \mu_t) \quad (\text{by def. of } tv_{\otimes}, \mu_s, \mu_t). \end{aligned}$$

□

Appendix B. Proofs of results in Section 5

Proof of Proposition 3. Firstly we show that

$$\mathbf{H}(D_{\mathcal{D}})(L_1, L_2) \leq \sup_{\phi \in L} \left| \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \right|. \quad (\text{B.1})$$

We can assume wlog. that $\mathbf{H}(D_{\mathcal{D}})(L_1, L_2) = \sup_{\phi_1 \in L_1} \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi_1, \phi_2)$. Then

$$\begin{aligned} \sup_{\phi_1 \in L_1} \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi_1, \phi_2) &= \sup_{\phi_1 \in L_1} \left| \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi_1, \phi_2) - \inf_{\phi' \in L_1} D_{\mathcal{D}}(\phi_1, \phi') \right| \\ &\leq \sup_{\phi \in L} \left| \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \right| \end{aligned}$$

from which Equation (B.1) holds.

Next, we aim to show the converse inequality, namely

$$\mathbf{H}(D_{\mathcal{D}})(L_1, L_2) \geq \sup_{\phi \in L} \left| \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \right|. \quad (\text{B.2})$$

To this aim, we show that

$$\text{for each } \phi \in L \text{ it holds } \left| \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \right| \leq \mathbf{H}(D_{\mathcal{D}})(L_1, L_2). \quad (\text{B.3})$$

- Assume that $\phi \in L_1$. Then $\inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) = 0$ so that $\left| \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \right| = \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2)$. Moreover

$$\inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \leq \sup_{\phi_1 \in L_1} \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi_1, \phi_2) \leq \mathbf{H}(D_{\mathcal{D}})(L_1, L_2)$$

and Equation (B.3) follows in this case.

- The case of $\phi \in L_2$ is analogous and therefore Equation (B.3) follows also in this case.
- Finally, assume that $\phi \notin L_1 \cup L_2$. Without loss of generality, we can assume that $\inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) \geq \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2)$. Moreover, we recall that by definition of *infimum* it holds that for each $\varepsilon > 0$ there exists a formula $\phi_{\varepsilon} \in L_2$ such that

$$D_{\mathcal{D}}(\phi, \phi_{\varepsilon}) < \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) + \varepsilon. \quad (\text{B.4})$$

Analogously, for each $\varepsilon' > 0$ and for each $\phi_2 \in L_2$ there is a $\phi_{\varepsilon'} \in L_1$ such that

$$D_{\mathcal{D}}(\phi_2, \phi_{\varepsilon'}) < \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi_2, \phi_1) + \varepsilon'. \quad (\text{B.5})$$

Let us fix $\varepsilon, \varepsilon' > 0$. Then let $\phi_{\varepsilon} \in L_2$ be the formula realizing Equation (B.4), with respect to ϕ , and let $\tilde{\phi}_{\varepsilon'}$ be the formula in L_1 realizing Equation (B.4), with respect to this ϕ_{ε} . Therefore, we have

$$\begin{aligned} & \left| \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \right| \\ &= \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \\ &< \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - D_{\mathcal{D}}(\phi, \phi_{\varepsilon}) + \varepsilon && (\text{by Equation (B.4)}) \\ &< D_{\mathcal{D}}(\phi, \tilde{\phi}_{\varepsilon'}) - D_{\mathcal{D}}(\phi, \phi_{\varepsilon}) + \varepsilon \\ &\leq D_{\mathcal{D}}(\phi, \phi_{\varepsilon}) + D_{\mathcal{D}}(\phi_{\varepsilon}, \tilde{\phi}_{\varepsilon'}) - D_{\mathcal{D}}(\phi, \phi_{\varepsilon}) + \varepsilon && (\text{by triangle inequality}) \\ &= D_{\mathcal{D}}(\phi_{\varepsilon}, \tilde{\phi}_{\varepsilon'}) + \varepsilon \\ &< \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi_{\varepsilon}, \phi_1) + \varepsilon' + \varepsilon && (\text{by Equation (B.5)}) \\ &\leq \sup_{\phi_2 \in L_2} \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi_2, \phi_1) + \varepsilon' + \varepsilon \\ &\leq \mathbf{H}(D_{\mathcal{D}})(L_1, L_2) + \varepsilon' + \varepsilon. \end{aligned}$$

Summarizing, we have obtained that

$$\left| \inf_{\phi_1 \in L_1} D_{\mathcal{D}}(\phi, \phi_1) - \inf_{\phi_2 \in L_2} D_{\mathcal{D}}(\phi, \phi_2) \right| < \mathbf{H}(D_{\mathcal{D}})(L_1, L_2) + \varepsilon' + \varepsilon$$

and since this inequality holds for each ε and ε' , we can conclude that Equation (B.3) follows also in this case.

Equation (B.1) and Equation (B.2) taken together prove the thesis. $\square$

Proof of Lemma 1. Immediate from Definition 10, Definition 11 and Proposition 3. $\square$

Proof of Proposition 4. We expand only the case of the additive total variation distance. The general case can be obtained by proving the two inequalities separately.

Since in this case $\mathcal{O}_{\mathcal{D}}([0, 1]) = 1$, we have

$$\begin{aligned}
\sup_{\Phi \in \mathbb{L}^1} |\llbracket \Phi \rrbracket(s) - \llbracket \Phi \rrbracket(t)| &= \sup_{\Phi \in \mathbb{L}^1} \left| \inf_{\Psi_s \in \mathbb{L}(s)} \mathfrak{d}^{\mathsf{P}}(\Phi, \Psi_s) - \inf_{\Psi_t \in \mathbb{L}(t)} \mathfrak{d}^{\mathsf{P}}(\Phi, \Psi_t) \right| \\
&= \sup_{\alpha \in \mathcal{A}^*} \left| \inf_{\Psi_s \in \mathbb{L}(s)} \mathfrak{d}^{\mathsf{P}}(\Phi_{\alpha}, \Psi_s) - \inf_{\Psi_t \in \mathbb{L}(t)} \mathfrak{d}^{\mathsf{P}}(\Phi_{\alpha}, \Psi_t) \right| \\
&= \sup_{\alpha \in \mathcal{A}^*} |\mathfrak{d}^{\mathsf{P}}(\Phi_{\alpha}, \Pr(s, \alpha)\Phi_{\alpha}) - \mathfrak{d}^{\mathsf{P}}(\Phi_{\alpha}, \Pr(t, \alpha)\Phi_{\alpha})| \\
&= \sup_{\alpha \in \mathcal{A}^*} ||1 - \Pr(s, \alpha)| - |1 - \Pr(t, \alpha)|| \\
&= \sup_{\alpha \in \mathcal{A}^*} |\Pr(s, \alpha) - \Pr(t, \alpha)| \\
&= tv(\mu_s, \mu_t)
\end{aligned}$$

where:

- the second step follows by letting Φ_{α} be the formula in $\mathbb{L}^1$ s.t. $\text{Tr}(\Phi_{\alpha}) = \alpha$;
- the third step follows by definition of $\mathfrak{d}^{\mathsf{P}}$ and the fact that $|r_1 - r_2| \leq 1$ for all $r_1, r_2 \in [0, 1]$;
- the sixth step follows by definition of tv .

$\square$

Appendix C. Proofs of results in Section 6

Proof of Proposition 6. First of all we notice that the up-to- k bisimilarity metrics constitute an ascending chain of $\mathbf{B}_V$

$$\mathbf{B}(\mathbf{0}) \preceq \mathbf{B}^2(\mathbf{0}) \preceq \dots \preceq \mathbf{B}^n(\mathbf{0}) \preceq \dots$$

Moreover, since we are considering image-finite processes and distributions with finite support, functional $\mathbf{B}_V$ is monotone, continuous and its closure ordinal is ω . Thus we can infer that $\lim_{k \rightarrow \infty} \mathbf{d}_{\lambda, V}^k = \mathbf{d}_{\lambda, V}^{\omega} = \sup_{k \in \mathbb{N}} \mathbf{B}_V^k(\mathbf{0})$ and that $\mathbf{d}_{\lambda, V}^{\omega}$ is a fixed point of $\mathbf{B}_V$. By an easy induction over $k \in \mathbb{N}$, we can prove that $\mathbf{d}_{\lambda, V} \geq \mathbf{d}_{\lambda, V}^k$ for all $k \in \mathbb{N}$. In particular $\mathbf{d}_{\lambda, V} \geq \mathbf{d}_{\lambda, V}^{\omega}$. Hence, by uniqueness of the least fixed point, we can conclude that $\mathbf{d}_{\lambda, V}^{\omega} = \mathbf{d}_{\lambda, V}$. $\square$

Appendix D. Proofs of the results in Section 8

Proof of Lemma 2. The proof follows by induction over $k \in \mathbb{N}$ and over the structure of formulae combined with the monotonicity of the generalized Kantorovich lifting proved in [15]. $\square$

Proof of Proposition 7. Assume any $\varphi, \varphi' \in \mathcal{L}^s$. By Lemma 2 we have $\mathfrak{d}_{\lambda, V}^k(\varphi, \varphi') \leq \mathfrak{d}_{\lambda, V}^{k+1}(\varphi, \varphi')$ and moreover, since we are in the hypothesis of Remark 1, $\mathfrak{d}_{\lambda, V}^k$ is bounded for each $k \in \mathbb{N}$. Hence $(\mathfrak{d}_{\lambda, V}^k(\varphi, \varphi'))_{k \in \mathbb{N}}$ is a bounded non decreasing sequence of pseudometrics. This ensures that $\lim_{k \rightarrow \infty} \mathfrak{d}_{\lambda, V}^k(\varphi, \varphi')$ exists. We

conclude that $\mathfrak{d}_{\lambda,V}$ is well-defined. $\square$

Proof of Proposition 8. The proof of the inclusion $\equiv_{\mathcal{L}} \subseteq \ker(\mathfrak{d}_{\lambda,V})$ follows by induction over the structure of formulae. The only interesting case, that we expand here, is the inductive step related to the diamond modality. So, let $\varphi = \langle a \rangle \psi$. Then, by Definition 22, $\varphi' \equiv_{\mathcal{L}} \varphi$ only if $\varphi' = \bigwedge_{j \in J} \langle a \rangle \psi_j$ with $\psi_j \equiv_{\mathcal{L}}^\dagger \psi$ for all $j \in J$. For simplicity, we consider the case of $|J| = 1$. The case of $|J| > 1$ follows as an easy generalization. Let $\varphi' = \langle a \rangle \psi'$ with $\psi' \equiv_{\mathcal{L}}^\dagger \psi$. Consequently, if we assume $\psi = \bigoplus_{i \in I} r_i \varphi_i$, by Definition 12 we get that ψ' is of the form $\psi' = \bigoplus_{i \in I} r_i \varphi'_i$ with $\varphi'_i \equiv_{\mathcal{L}} \varphi_i$ for all $i \in I$. Then we have

$$\begin{aligned}
& \mathfrak{d}_{\lambda,V}(\varphi, \varphi') \\
&= \mathfrak{d}_{\lambda,V}(\langle a \rangle \psi, \langle a \rangle \psi') \\
&= \lambda \cdot \mathbf{K}_V(\mathfrak{d}_{\lambda,V})(\psi, \psi') \\
&= \lambda \cdot \sup_f d_V \left(\sum_{i \in I} r_i f(\varphi_i), \sum_{i \in I} r_i f(\varphi'_i) \right) \quad \text{subject to } d_V(f(\varphi_i), f(\varphi'_i)) \leq \mathfrak{d}_{\lambda,V}(\varphi_i, \varphi'_i) \ \forall i \in I \\
&= \lambda \cdot \sup_f d_V \left(\sum_{i \in I} r_i f(\varphi_i), \sum_{i \in I} r_i f(\varphi'_i) \right) \quad \text{subject to } d_V(f(\varphi_i), f(\varphi'_i)) \leq 0 \ \forall i \in I \\
&= \lambda \cdot \sup_f d_V \left(\sum_{i \in I} r_i f(\varphi_i), \sum_{i \in I} r_i f(\varphi'_i) \right) \quad \text{subject to } f(\varphi_i) = f(\varphi'_i) \ \forall i \in I \\
&= 0
\end{aligned}$$

where:

- the third step follows by the evaluation of the Kantorovich metric in terms of the equivalent linear program;
- the forth step follows by the inductive hypothesis, for which $\varphi_i \equiv_{\mathcal{L}} \varphi'_i$ implies $\mathfrak{d}_{\lambda,V}(\varphi_i, \varphi'_i) = 0$ for all $i \in I$;
- the fifth and sixth steps follow by d_V being a metric.

The proof of the opposite inclusion $\ker(\mathfrak{d}_{\lambda,V}) \subseteq \equiv_{\mathcal{L}}$ follows from the same arguments used in the technical report version of [15] to prove that the kernel of each bisimulation metric is a bisimulation [15, Proposition 4]. Thus, we just sketch here the main reasoning and we refer the interested reader to [15] for a detailed proof.

Our purpose is to prove that $\ker(\mathfrak{d}_{\lambda,V})$ is a post-fixed point of the functional E introduced in Definition 21. To this aim, we need first to ensure that the lifting of the kernel of $\mathfrak{d}_{\lambda,V}$ to a relation over distribution formulae coincides with the kernel of the generalized Kantorovich metric $\mathbf{K}_V$, that we use to evaluate the distance over formulae in $\mathcal{L}^d$. This result corresponds to [15, Lemma 1] and it holds under the requirement of (V, d_V) having a geodesic. This is to guarantee that in V there are no isolated points that can cause the distance d_V (and thus $\mathbf{K}_V$) to explode. Having this preliminary result, we can proceed by induction over the structure of formulae to prove that $\ker(\mathfrak{d}_{\lambda,V}) \subseteq E(\ker(\mathfrak{d}_{\lambda,V}))$. The non-trivial inductive step of the diamond modality will follow by $\ker(\mathfrak{d}_{\lambda,V})^\dagger = \ker(\mathbf{K}_V(\mathfrak{d}_{\lambda,V}))$. $\square$

Proof of Theorem 5. We aim at proving that

$$\ell_{\lambda,V}^k(s, t) = \mathbf{d}_{\lambda,V}^k(s, t) \text{ for all } k \in \mathbb{N} \text{ and for all } s, t \in \mathcal{S} \quad (\text{D.1})$$

To this end, we proceed by induction over k . Consider the base case $k = 0$. We have

$$\begin{aligned}
& \mathbf{d}_{\lambda,V}^0(s, t) \\
&= 0 \quad \text{(by definition of } \mathbf{d}_{\lambda,V}^0) \\
&= \mathfrak{d}_{\lambda,V}^0(\top, \top) \quad \text{(by definition of } \mathfrak{d}_{\lambda,V}^0)
\end{aligned}$$

$$\begin{aligned}
&= \mathfrak{d}_{\lambda,V}^0(\varphi_s^0, \varphi_t^0) && \text{(by definition of } \varphi_s^0, \varphi_t^0) \\
&= \ell_{\lambda,V}^0(s, t) && \text{(by definition of } \ell_{\lambda,V}^0).
\end{aligned}$$

Consider the inductive step $k > 0$. By the inductive hypothesis, for all $s', t' \in \mathcal{S}$ it holds that $\ell_{\lambda,V}^k(s', t') = \mathbf{d}_{\lambda,V}^k(s', t')$. Given any process s , we define

$$\phi_s^k = \bigwedge_{(s,a,\pi) \in \rightarrow} \langle a \rangle \psi_\pi^k \quad \text{and} \quad \theta_s = \bigwedge_{b \notin \text{init}(s)} \neg \langle b \rangle \top$$

so that the up-to- $(k+1)$ mimicking formula of s can be rewritten as

$$\varphi_s^{k+1} = \phi_s^k \wedge \theta_s.$$

Hence, for any pair of processes $s, t \in \mathcal{S}$ we have

$$\begin{aligned}
&\ell_{\lambda,V}^{k+1}(s, t) \\
&= \mathfrak{d}_{\lambda,V}^{k+1}(\varphi_s^{k+1}, \varphi_t^{k+1}) \\
&= \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k \wedge \theta_s, \phi_t^k \wedge \theta_t) \\
&= \max \left\{ \begin{array}{l} \sup \left\{ \begin{array}{l} \inf \{ \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k), \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \theta_t) \} \\ \inf \{ \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \phi_t^k), \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t) \} \end{array} \right\}, \\ \sup \left\{ \begin{array}{l} \inf \{ \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k), \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \phi_t^k) \} \\ \inf \{ \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \theta_t), \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t) \} \end{array} \right\} \end{array} \right\} \\
&= \max \left\{ \begin{array}{l} \sup \left\{ \begin{array}{l} \inf \{ \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k), \odot_{d_V}(V) \} \\ \inf \{ \odot_{d_V}(V), \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t) \} \end{array} \right\}, \\ \sup \left\{ \begin{array}{l} \inf \{ \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k), \odot_{d_V}(V) \} \\ \inf \{ \odot_{d_V}(V), \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t) \} \end{array} \right\} \end{array} \right\} \\
&= \max \{ \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k), \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t) \}
\end{aligned}$$

where the second last equality follows by

$$\begin{aligned}
&\mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \theta_t) \\
&= \mathfrak{d}_{\lambda,V}^{k+1} \left(\bigwedge_{(s,a,\pi) \in \rightarrow} \langle a \rangle \psi_\pi^k, \bigwedge_{b \notin \text{init}(t)} \neg \langle b \rangle \top \right) \\
&= \max \left\{ \begin{array}{l} \sup_{(s,a,\pi) \in \rightarrow} \inf_{b \notin \text{init}(t)} \mathfrak{d}_{\lambda,V}^{k+1}(\langle a \rangle \psi_\pi^k, \neg \langle b \rangle \top), \\ \sup_{b \notin \text{init}(t)} \inf_{(s,a,\pi) \in \rightarrow} \mathfrak{d}_{\lambda,V}^{k+1}(\langle a \rangle \psi_\pi^k, \neg \langle b \rangle \top) \end{array} \right\} \\
&= \max \left\{ \begin{array}{l} \sup_{(s,a,\pi) \in \rightarrow} \inf_{b \notin \text{init}(t)} \{ \odot_{d_V}(V) \}, \\ \sup_{b \notin \text{init}(t)} \inf_{(s,a,\pi) \in \rightarrow} \{ \odot_{d_V}(V) \} \end{array} \right\} \\
&= \odot_{d_V}(V)
\end{aligned}$$

and, analogously, $\mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \phi_t^k) = \odot_{d_V}(V)$.

Summarizing, Equation D.1 becomes

$$\mathbf{d}_{\lambda,V}^{k+1}(s, t) = \max \{ \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k), \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t) \}. \quad (\text{D.2})$$

To show Equation D.2 we distinguish two cases: either $\text{init}(s) \neq \text{init}(t)$ or $\text{init}(s) = \text{init}(t)$. Let us start with the case $\text{init}(s) \neq \text{init}(t)$. Without loss of generality, assume that there is some action $\hat{b}$ with $\hat{b} \in \text{init}(s) \setminus \text{init}(t)$. The case $\hat{b} \in \text{init}(t) \setminus \text{init}(s)$ is analogous. Under this assumption, we have

$$\begin{aligned}
& \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t) \\
&= \max\left\{ \sup_{b \notin \text{init}(s)} \inf_{b' \notin \text{init}(t)} \mathfrak{d}_{\lambda,V}^{k+1}(\neg\langle b \rangle \top, \neg\langle b' \rangle \top), \sup_{b' \notin \text{init}(t)} \inf_{b \notin \text{init}(s)} \mathfrak{d}_{\lambda,V}^{k+1}(\neg\langle b \rangle \top, \neg\langle b' \rangle \top) \right\} \\
&\geq \max\left\{ \sup_{b \notin \text{init}(s)} \inf_{b' \notin \text{init}(t)} \mathfrak{d}_{\lambda,V}^{k+1}(\neg\langle b \rangle \top, \neg\langle b' \rangle \top), \inf_{b \notin \text{init}(s)} \mathfrak{d}_{\lambda,V}^{k+1}(\neg\langle b \rangle \top, \neg\langle \hat{b} \rangle \top) \right\} \\
&= \max\left\{ \sup_{b \notin \text{init}(s)} \inf_{b' \notin \text{init}(t)} \mathfrak{d}_{\lambda,V}^{k+1}(\neg\langle b \rangle \top, \neg\langle b' \rangle \top), \mathcal{O}_{d_V}(V) \right\} \\
&= \mathcal{O}_{d_V}(V)
\end{aligned}$$

where the second last equality follows from $t \xrightarrow{\hat{b}}$ whereas $s \not\xrightarrow{\hat{b}}$ does not hold. Thus Equation D.2 instantiates as

$$\mathbf{d}_{\lambda,V}^{k+1}(s, t) = \max\{\mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k), \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t)\} = \max\{\mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k), \mathcal{O}_{d_V}(V)\} = \mathcal{O}_{d_V}(V)$$

which holds by $\text{init}(s) \neq \text{init}(t)$.

The second case is $\text{init}(s) = \text{init}(t)$. We prove first that $\mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t) = 0$. We have

$$\begin{aligned}
& \mathfrak{d}_{\lambda,V}^{k+1}(\theta_s, \theta_t) \\
&= \max\left\{ \sup_{b \notin \text{init}(s)} \inf_{b' \notin \text{init}(t)} \mathfrak{d}_{\lambda,V}^{k+1}(\neg\langle b \rangle \top, \neg\langle b' \rangle \top), \sup_{b' \notin \text{init}(t)} \inf_{b \notin \text{init}(s)} \mathfrak{d}_{\lambda,V}^{k+1}(\neg\langle b \rangle \top, \neg\langle b' \rangle \top) \right\} \\
&\leq \max\left\{ \sup_{b \notin \text{init}(s)} \mathfrak{d}_{\lambda,V}^{k+1}(\neg\langle b \rangle \top, \neg\langle b \rangle \top), \sup_{b' \notin \text{init}(t)} \mathfrak{d}_{\lambda,V}^{k+1}(\neg\langle b' \rangle \top, \neg\langle b' \rangle \top) \right\} \\
&= \max\left\{ \sup_{b \notin \text{init}(s)} 0, \sup_{b' \notin \text{init}(t)} 0 \right\} \\
&= 0.
\end{aligned}$$

Therefore, Equation D.2 becomes

$$\mathbf{d}_{\lambda,V}^{k+1}(s, t) = \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k) \tag{D.3}$$

which follows by

$$\begin{aligned}
& \mathfrak{d}_{\lambda,V}^{k+1}(\phi_s^k, \phi_t^k) \\
&= \mathfrak{d}_{\lambda,V}^{k+1}\left(\bigwedge_{(s,a,\pi_s) \in \rightarrow} \langle a \rangle \psi_{\pi_s}^k, \bigwedge_{(t,a,\pi_t) \in \rightarrow} \langle a' \rangle \psi_{\pi_t}^k\right) \\
&= \max \left\{ \begin{array}{l} \sup_{(s,a,\pi_s) \in \rightarrow} \inf_{(t,a,\pi_t) \in \rightarrow} \mathfrak{d}_{\lambda,V}^{k+1}(\langle a \rangle \psi_{\pi_s}^k, \langle a' \rangle \psi_{\pi_t}^k), \\ \sup_{(t,a,\pi_t) \in \rightarrow} \inf_{(s,a,\pi_s) \in \rightarrow} \mathfrak{d}_{\lambda,V}^{k+1}(\langle a \rangle \psi_{\pi_s}^k, \langle a' \rangle \psi_{\pi_t}^k) \end{array} \right\} \\
&= \max \left\{ \begin{array}{l} \sup_{a \in \mathcal{A}} \sup_{(s,a,\pi_s) \in \rightarrow} \inf_{(t,a,\pi_t) \in \rightarrow} \lambda \cdot \mathbf{K}_V(\mathfrak{d}_{\lambda,V}^k)(\psi_{\pi_s}^k, \psi_{\pi_t}^k), \\ \sup_{a \in \mathcal{A}} \sup_{(t,a,\pi_t) \in \rightarrow} \inf_{(s,a,\pi_s) \in \rightarrow} \lambda \cdot \mathbf{K}_V(\mathfrak{d}_{\lambda,V}^k)(\psi_{\pi_s}^k, \psi_{\pi_t}^k) \end{array} \right\} \\
&= \max \left\{ \begin{array}{l} \sup_{a \in \mathcal{A}} \sup_{(s,a,\pi_s) \in \rightarrow} \inf_{(t,a,\pi_t) \in \rightarrow} \lambda \cdot \mathbf{K}_V(\mathbf{d}_{\lambda,V}^k)(\pi_s, \pi_t), \\ \sup_{a \in \mathcal{A}} \sup_{(t,a,\pi_t) \in \rightarrow} \inf_{(s,a,\pi_s) \in \rightarrow} \lambda \cdot \mathbf{K}_V(\mathbf{d}_{\lambda,V}^k)(\pi_s, \pi_t) \end{array} \right\} \\
&= \mathbf{d}_{\lambda,V}^{k+1}(s, t)
\end{aligned}$$

where the forth equality can be proved as follows. By definition we have that

$$\mathbf{K}_V(\mathfrak{d}_{\lambda,V}^k)(\psi_{\pi_s}^k, \psi_{\pi_t}^k) = \sup_f d_V \left(\sum_{s' \in \text{supp}(\pi_s)} \pi_s(s') f(\varphi_{s'}^k), \sum_{t' \in \text{supp}(\pi_t)} \pi_t(t') f(\varphi_{t'}^k) \right)$$

where the supremum is evaluated over all functions f such that $d_V(f(\varphi_{s'}^k), f(\varphi_{t'}^k)) \leq \mathfrak{d}_{\lambda,V}^k(\varphi_{s'}^k, \varphi_{t'}^k)$ for all $s' \in \text{supp}(\pi_s), t' \in \text{supp}(\pi_t)$. By definition of supremum, for each $\varepsilon > 0$, there is a particular function f_ε satisfying this constraint such that

$$\mathbf{K}_V(\mathfrak{d}_{\lambda,V}^k)(\psi_{\pi_s}^k, \psi_{\pi_t}^k) < d_V \left(\sum_{s' \in \text{supp}(\pi_s)} \pi_s(s') f_\varepsilon(\varphi_{s'}^k), \sum_{t' \in \text{supp}(\pi_t)} \pi_t(t') f_\varepsilon(\varphi_{t'}^k) \right) + \varepsilon.$$

Let $g: \mathcal{S} \rightarrow V$ be defined as $g(s) = f_\varepsilon(\varphi_s^k)$ for all $s \in \mathcal{S}$. Then, since by the inductive hypothesis we have $\mathfrak{d}_{\lambda,V}^k(\varphi_{s'}^k, \varphi_{t'}^k) = \mathbf{d}_{\lambda,V}^k(s', t')$ for all $s', t' \in \mathcal{S}$, by the choice of g we obtain that $d_V(g(s'), g(t')) \leq \mathbf{d}_{\lambda,V}^k(s', t')$ for all $s' \in \text{supp}(\pi_s), t' \in \text{supp}(\pi_t)$. Therefore, g is one of the 1-Lipschitz functions on which $\mathbf{K}_V(\mathbf{d}_{\lambda,V}^k)(\pi_s, \pi_t)$ is evaluated and thus we can draw that

$$\begin{aligned} & \mathbf{K}_V(\mathfrak{d}_{\lambda,V}^k)(\psi_{\pi_s}^k, \psi_{\pi_t}^k) \\ & < d_V \left(\sum_{s' \in \text{supp}(\pi_s)} \pi_s(s') f_\varepsilon(\varphi_{s'}^k), \sum_{t' \in \text{supp}(\pi_t)} \pi_t(t') f_\varepsilon(\varphi_{t'}^k) \right) + \varepsilon \\ & = d_V \left(\sum_{s' \in \text{supp}(\pi_s)} \pi_s(s') g(s'), \sum_{t' \in \text{supp}(\pi_t)} \pi_t(t') g(t') \right) + \varepsilon \\ & \leq \mathbf{K}_V(\mathbf{d}_{\lambda,V}^k)(\pi_s, \pi_t) + \varepsilon. \end{aligned}$$

Since the inequality holds for all $\varepsilon > 0$, we can conclude that

$$\mathbf{K}_V(\mathfrak{d}_{\lambda,V}^k)(\psi_{\pi_s}^k, \psi_{\pi_t}^k) \leq \mathbf{K}_V(\mathbf{d}_{\lambda,V}^k)(\pi_s, \pi_t).$$

Similarly, we can prove that

$$\mathbf{K}_V(\mathbf{d}_{\lambda,V}^k)(\pi_s, \pi_t) \leq \mathbf{K}_V(\mathfrak{d}_{\lambda,V}^k)(\psi_{\pi_s}^k, \psi_{\pi_t}^k)$$

and thus the equality at the forth step above follows. $\square$

Appendix E. Proofs of results in Section 9

Proof of Proposition 9. In [15, Theorem 2] it was proved that for all (V, d_V) we have $\mathbf{K}_V(\mathbf{dm}_{\otimes_V})(\pi_s, \pi_t) \leq \mathbf{d}_{\lambda,V}(s, t)$. Therefore, the proof follows by this result, our characterization Theorem 6 and the equality $tv_{\otimes} = \mathbf{K}_{\otimes}(\mathbf{dm}_{\otimes})$. $\square$