Beyond Differential Privacy: Composition Theorems and Relational Logic for f-divergences between Probabilistic Programs

Gilles Barthe and Federico Olmedo

IMDEA Software Institute, Madrid, Spain
{Gilles.Barthe,Federico.Olmedo}@imdea.org

Abstract. f-divergences form a class of measures of distance between probability distributions; they are widely used in areas such as information theory and signal processing. In this paper, we unveil a new connection between f-divergences and differential privacy, a confidentiality policy that provides strong privacy guarantees for private data-mining; specifically, we observe that the notion of α-distance used to characterize approximate differential privacy is an instance of the family of f-divergences. Building on this observation, we generalize to arbitrary f-divergences the sequential composition theorem of differential privacy. Then, we propose a relational program logic to prove upper bounds for the f-divergence between two probabilistic programs. Our results allow us to revisit the foundations of differential privacy under a new light, and to pave the way for applications that use different instances of f-divergences.

1 Introduction

Differential privacy [12] is a policy that provides strong privacy guarantees in private data analysis: informally, a randomized computation over a database D is differentially private if the private data of individuals contributing to D is protected against arbitrary adversaries with query access to D. Formally, let $\epsilon \geq 0$ and $0 \leq \delta \leq 1$: a randomized algorithm c is (ϵ, δ)-differentially private if its output distributions for any two neighbouring inputs x and y are (e^{ϵ}, δ)-close, i.e. for every event E:

$$\Pr[c(x) : E] \leq e^{\epsilon} \Pr[c(y) : E] + \delta$$

where $\Pr[c(x) : E]$ denotes the probability of event E in the distribution obtained by running c on input x. One key property of differential privacy is the existence of sequential and parallel composition theorems, which allows building differentially private computations from smaller blocks. In this paper, we focus on the first theorem, which states that the sequential composition of an (ϵ_1, δ_1)-differentially private algorithm with an (ϵ_2, δ_2)-differentially private one yields an $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$-differentially private algorithm.

f-divergences [2, 10] are convex functions that can be used to measure the distance between two distributions. The class of f-divergences includes many well-known notions

of distance, such as statistical distance, Kullback-Leibler divergence (relative entropy), or Hellinger distance. Over the years, f-divergences have found multiple applications in information theory, signal processing, pattern recognition, machine learning, and security. The practical motivation for this work is a recent application of f-divergences to cryptography: in [24], Steinberger uses Hellinger distance to improve the security analysis of key-alternating ciphers, a family of encryption schemes that encompasses the Advanced Encryption Standard AES.

Deductive verification of differentially private computations. In [6], we develop an approximate probabilistic Hoare logic, called apRHL, for reasoning about differential privacy of randomized computations. The logic manipulates judgments of the form:

$$c_1 \sim_{\alpha,\delta} c_2 : \Psi \Rightarrow \Phi$$

where c_1 and c_2 are probabilistic imperative programs, $\alpha \geq 1$, $0 \leq \delta \leq 1$ and Ψ and Φ are relations over states. As for its predecessor pRHL [5], the notion of valid judgment rests on a lifting operator that turns a relation R over states into a relation $\sim_R^{\alpha,\delta}$ over distributions of states: formally, the judgment above is valid iff for every pair of memories m_1 and m_2, $m_1 \mathrel{\Psi} m_2$ implies $(\llbracket c_1 \rrbracket\, m_1) \sim_\Phi^{\alpha,\delta} (\llbracket c_2 \rrbracket\, m_2)$. The definition of the lifting operator originates from probabilistic process algebra [15], and has close connections with flow networks and the Kantorovich metric [11].

apRHL judgments characterize differential privacy, in the sense that c is (ϵ, δ)-differentially private iff the apRHL judgment $c \sim_{e^\epsilon,\delta} \Psi : c \Rightarrow \equiv$ is valid, where Ψ is a logical characterization of adjacency—for instance, two lists of the same length are adjacent if they differ in a single element.

Problem statement and contributions. The goal of this paper is to lay the theoretical foundations for tool-supported reasoning about f-divergences between probabilistic computations. To achieve this goal, we start from [6] and take the following steps:

1. as a preliminary observation, we prove that the notion of α-distance used to characterize differential privacy is in fact an f-divergence;
2. we define a notion of composability of f-divergences and generalize the sequential composition theorem of differential privacy to composable divergences;
3. we generalize the notion of lifting used in apRHL to composable f-divergences;
4. we define fpRHL, a probabilistic relational Hoare logic for f-divergences, and prove its soundness.

Related work. The problem of computing the distance between two probabilistic computations has been addressed in different areas of computer science, including machine learning, stochastic systems, and security. We briefly point to some recent developments.

Methods for computing the distance between probabilistic automata have been studied by Cortes and co-authors [8, 9]; their work, which is motivated by machine-learning applications, considers the Kullback-Leibler divergence as well as the L_p distance.

Approximate bisimulation for probabilistic automata has been studied, among others, by Segala and Turrini [23] and by Tracol, Desharnais and Zhioua [25]. The survey [1] provides a more extensive account of the field.

In the field of security, approximate probabilistic bisimulation is closely connected to quantitative information flow of probabilistic computations, which has been studied e.g. by Di Pierro, Hankin and Wiklicky [20]. More recently, the connections between quantitative information flow and differential privacy have been explored e.g. by Barthe and Köpf [4], and by Alvim, Andrés, Chatzikokolakis and Palamidessi [3]. Moreover, several language-based methods have been developed for guaranteeing differential privacy; these methods are based on runtime verification, such as PINQ [17] or Airavat [22], type systems [21, 14], or deductive verification [7]. We refer to [19] for a survey of programming languages methods for differential privacy.

2 Mathematical Preliminaries

In this section we review the representation of distributions used in our development and recall the definition of f-divergences.

2.1 Probability Distributions

Throughout the presentation we consider distributions and sub-distributions over discrete sets only. A probability distribution (resp. sub-distribution) over a set A is an object $\mu : A \to [0,1]$ such that $\sum_{a\in A}\mu(a) = 1$ (resp. $\sum_{a\in A}\mu(a) \leq 1$). We let $\mathcal{D}(A)$ (resp. $\mathcal{D}_{\leq 1}(A)$) be the set of distributions (resp. sub-distributions) over A.

Distributions are closed under convex combinations: given distributions $(\mu_i)_{i\in\mathbb{N}}$ in $\mathcal{D}(A)$ and weights $(w_i)_{i\in\mathbb{N}}$ such that $\sum_{i\in\mathbb{N}} w_i = 1$ and $w_i \geq 0$ for all $i \in \mathbb{N}$, the convex combination $\sum_{i\in\mathbb{N}} w_i\,\mu_i$ is also a distribution over A. Thus, given $\mu \in \mathcal{D}(A)$ and $M : A \to \mathcal{D}(B)$, we define the distribution $\mathsf{bind}\;\mu\;M$ over B as $\sum_{a\in A}\mu(a)\;M(a)$. Likewise, sub-distributions are closed under convex combinations.

2.2 f-divergences

Let $\mathcal{F}$ be the set of non-negative convex functions $f : \mathbb{R}_0^+ \to \mathbb{R}_0^+$ such that f is continuous at 0 and $f(1) = 0$. Then each function in $\mathcal{F}$ induces a notion of distance between probability distributions as follows:

Definition 1 (f-divergence). *Given $f \in \mathcal{F}$, the* f-divergence $\Delta_f(\mu_1,\mu_2)$ *between two distributions μ_1 and μ_2 in $\mathcal{D}(A)$ is defined as:*

$$\Delta_f(\mu_1,\mu_2) \overset{def}{=} \sum_{a\in A}\mu_2(a) f\left(\frac{\mu_1(a)}{\mu_2(a)}\right)$$

The definition adopts the following conventions, which are used consistently throughout the paper:

$$0f\,(0/0) = 0 \qquad \textit{and} \qquad 0f\,(t/0) = t \lim_{x\to 0^+} x f\,(1/x) \quad \textit{if}\, t > 0$$

Moreover, if $\Delta_f(\mu_1,\mu_2) \leq \delta$ we say that μ_1 and μ_2 are (f,δ)-close.

f-divergence	**f**	**Simplified Form**
Statistical distance	$\mathsf{SD}(t) = \frac{1}{2}\lvert t-1\rvert$	$\sum_{a\in A}\frac{1}{2}\lvert\mu_1(a)-\mu_2(a)\rvert$
Kullback-Leibler[1]	$\mathsf{KL}(t) = t\ln(t) - t + 1$	$\sum_{a\in A}\mu_1(a)\ln\left(\frac{\mu_1(a)}{\mu_2(a)}\right)$
Hellinger distance	$\mathsf{HD}(t) = \frac{1}{2}(\sqrt{t}-1)^2$	$\sum_{a\in A}\frac{1}{2}\left(\sqrt{\mu_1(a)}-\sqrt{\mu_2(a)}\right)^2$

Fig. 1. Examples of f-divergences.

When defining f-divergences one usually allows f to take positive as well as negative values in $\mathbb{R}$. For technical reasons, however, we consider only non-negative functions. We now show that we can adopt this restriction without loss of generality.

Proposition 1. *Let $\mathcal{F}'$ be defined as $\mathcal{F}$, except that we allow $f \in \mathcal{F}'$ to take negative values. Then for every $f \in \mathcal{F}'$ there exists $g \in \mathcal{F}$ given by $g(t) = f(t) - f'_{-}(1)(t-1)$, such that $\Delta_f = \Delta_g$. (Here f'_{-} denotes the left derivative of f, whose existence can be guaranteed from the convexity of f.)*

The class of f-divergences includes several popular instances; these include statistical distance, relative entropy (also known as Kullback-Leibler divergence), and Hellinger distance. In Figure 1 we summarize the convex function used to define each of them and we also include a simplified form, useful to compute the divergence. (In case of negative functions, we previously apply the transformation mentioned in Proposition 1, so that we are consistent with our definition of f-divergences.)

In general, Δ_f does not define a metric. The symmetry axiom might be violated and the triangle inequality holds only if f equals a non-negative multiple of the statistical distance. The identity of indiscernibles does not hold in general, either.

3 A Sequential Composition Theorem for f-divergences

In this section we show that the notion of α-distance used to capture differential privacy is an f-divergence. Then we define the composition of f-divergences and show that the sequential composition theorem of differential privacy generalizes to this setting.

3.1 An f-divergence for Approximate Differential Privacy

In [6] we introduced the concept of α-distance to succinctly capture the notion of differentially private computations. Given $\alpha \geq 1$, the α-distance between distributions μ_1 and μ_2 in $\mathcal{D}(A)$ is defined as

$$\Delta_\alpha(\mu_1,\mu_2) \overset{\text{def}}{=} \max_{E\subseteq A} d_\alpha(\mu_1(E),\mu_2(E))$$

[1] Rigorously speaking, the function used for defining the Kullback-Leibler divergence should be given by $f(t) = t\ln(t) + t - 1$ if $t > 0$ and $f(t) = 1$ if $t = 0$ to guarantee its continuity at 0.

where $d_\alpha(a,b) \overset{\text{def}}{=} \max\{a - \alpha b, 0\}$. (This definition slightly departs from that of [6], in the sense that we consider an asymmetric version of the α-distance. The original version, symmetric, corresponds to taking $d_\alpha(a,b) \overset{\text{def}}{=} \max\{a - \alpha b, b - \alpha a, 0\}$). Now we can recast the definition of differential privacy in terms of the α-distance and say that a randomized computation c is (ϵ, δ)-*differentially private* iff $\Delta_{e^\epsilon}(c(x), c(y)) \leq \delta$ for any two adjacent inputs x and y.

Our composition result of f-divergences builds on the observation that α-distance is an instance of the class of f-divergences.

Proposition 2. *For every $\alpha \geq 1$, the α-distance $\Delta_\alpha(\mu_1, \mu_2)$ coincides with the f-divergence $\Delta_{\mathsf{AD}_\alpha}(\mu_1, \mu_2)$ associated to function $\mathsf{AD}_\alpha(t) \overset{\text{def}}{=} \max\{t - \alpha, 0\}$.*

3.2 Composition

One key property of f-divergences is a monotonicity result referred to as the *data processing inequality* [18]. In our setting, it is captured by the following proposition:

Proposition 3. *Let μ_1,$\mu_2 \in \mathcal{D}(A)$, $M : A \to \mathcal{D}(B)$ and $f \in \mathcal{F}$. Then*

$$\Delta_f(\mathsf{bind}\ \mu_1\ M, \mathsf{bind}\ \mu_2\ M) \leq \Delta_f(\mu_1, \mu_2)$$

In comparison, the sequential composition theorem for differential privacy [16] is captured by the following theorem.

Theorem 1. *Let μ_1,$\mu_2 \in \mathcal{D}(A)$, $M_1, M_2 : A \to \mathcal{D}(B)$ and $\alpha, \alpha' \geq 1$. Then*

$$\Delta_{\alpha\alpha'}(\mathsf{bind}\ \mu_1\ M_1, \mathsf{bind}\ \mu_2\ M_2) \leq \Delta_\alpha(\mu_1, \mu_2) + \max_a \Delta_{\alpha'}(M_1(a), M_2(a))$$

Note that the data processing inequality for α-distance corresponds to the composition theorem for the degenerate case where M_1 and M_2 are equal. The goal of this paragraph is to generalize the sequential composition theorem to f-divergences. To this end, we first define a notion of composability between f-divergences.

Definition 2 (f-divergence composability). *Let $f_1, f_2, f_3 \in \mathcal{F}$. We say that (f_1, f_2) is* f_3-composable *iff for all $\mu_1, \mu_2 \in \mathcal{D}(A)$ and $M_1, M_2 : A \to \mathcal{D}(B)$, there exists $\mu_3 \in \mathcal{D}(A)$ such that*

$$\Delta_{f_3}(\mathsf{bind}\ \mu_1\ M_1, \mathsf{bind}\ \mu_2\ M_2) \leq \Delta_{f_1}(\mu_1, \mu_2) + \sum_{a \in A} \mu_3(a) \Delta_{f_2}(M_1(a), M_2(a))$$

Our notion of composability is connected to the notion of additive information measures from [13, Ch. 5]. To justify the connection, we first present an adaptation of their definition to our setting.

Definition 3 (f-divergence additivity). *Let $f_1, f_2, f_3 \in \mathcal{F}$. We say that (f_1, f_2) is* f_3-additive *iff for all distributions $\mu_1, \mu_2 \in \mathcal{D}(A)$ and $\mu_1', \mu_2' \in \mathcal{D}(B)$,*

$$\Delta_{f_3}(\mu_1 \times \mu_1', \mu_2 \times \mu_2') \leq \Delta_{f_1}(\mu_1, \mu_2) + \Delta_{f_2}(\mu_1', \mu_2')$$

Here, $\mu \times \mu'$ denotes the product distribution of μ and μ', i.e. $(\mu \times \mu')(a, b) \overset{\text{def}}{=} \mu(a)\mu'(b)$.

It is easily seen that composability entails additivity.

Proposition 4. *Let $f_1, f_2, f_3 \in \mathcal{F}$ such that (f_1, f_2) is f_3-composable. Then (f_1, f_2) is f_3-additive.*

The f-divergences from Figure 1 present good behaviour under composition. The statistical distance, Hellinger distance and the Kullback-Leibler divergence are composable w.r.t. themselves. Moreover, α-divergences are composable.

Proposition 5.

- $(\mathsf{SD}, \mathsf{SD})$ *is* SD*-composable;*
- $(\mathsf{KL}, \mathsf{KL})$ *is* KL*-composable;*
- $(\mathsf{HD}, \mathsf{HD})$ *is* HD*-composable;*
- $(\mathsf{AD}_{\alpha_1}, \mathsf{AD}_{\alpha_2})$ *is* $\mathsf{AD}_{\alpha_1\alpha_2}$*-composable for every* $\alpha_1, \alpha_2 \geq 1$.

The sequential composition theorem of differential privacy extends naturally to the class of composable divergences.

Theorem 2. *Let $f_1, f_2, f_3 \in \mathcal{F}$. If (f_1, f_2) is f_3-composable, then for all $\mu_1, \mu_2 \in \mathcal{D}(A)$ and all $M_1, M_2 : A \to \mathcal{D}(B)$,*

$$\Delta_{f_3}(\mathsf{bind}\ \mu_1\ M_1, \mathsf{bind}\ \mu_2\ M_2) \leq \Delta_{f_1}(\mu_1, \mu_2) + \max_a \Delta_{f_2}(M_1(a), M_2(a))$$

Theorem 2 will be the cornerstone for deriving the sequential composition rule of fpRHL. (As an intermediate step, we first show that the composition result extends to relation liftings.)

4 Lifting

The definition of valid apRHL judgment rests on the notion of lifting. As a last step before defining our relational logic, we extend the notion of lifting to f-divergences. One key difference between our definition and that of [6] is that the former uses two witnesses, rather than one. In the remainder, we let $\mathsf{supp}\,(\mu)$ denote the set of elements $a \in A$ such that $\mu(a) > 0$. Moreover, given $\mu \in \mathcal{D}(A \times B)$, we define $\pi_1(\mu)$ and $\pi_2(\mu)$ by the clauses $\pi_1(\mu)(a) = \sum_{b \in B} \mu(a, b)$ and $\pi_2(\mu)(b) = \sum_{a \in A} \mu(a, b)$.

Definition 4 (Lifting). *Let $f \in \mathcal{F}$ and $\delta \in \mathbb{R}_0^+$. Then (f, δ)*-lifting *$\sim_R^{f,\delta}$ of a relation $R \subseteq A \times B$ is defined as follows: given $\mu_1 \in \mathcal{D}(A)$ and $\mu_2 \in \mathcal{D}(B)$, $\mu_1 \sim_R^{f,\delta} \mu_2$ iff there exist $\mu_L, \mu_R \in \mathcal{D}(A \times B)$ such that: i) $\mathsf{supp}\,(\mu_L) \subseteq R$; ii) $\mathsf{supp}\,(\mu_R) \subseteq R$; iii) $\pi_1(\mu_L) = \mu_1$; iv) $\pi_2(\mu_R) = \mu_2$ and v) $\Delta_f(\mu_L, \mu_R) \leq \delta$. The distributions μ_L and μ_R are called the left and right witnesses for the lifting, respectively.*

A pleasing consequence of our definition is that the witnesses for relating two distributions are themselves distributions, rather than sub-distributions; this is in contrast with our earlier definition from [6], where witnesses for the equality relation are necessarily sub-distributions. Moreover, our definition is logically equivalent to the original one from [15], provided $\delta = 0$, and f satisfies the identity of indiscernibles. In the case of statistical distance and α-distance, our definition also has a precise mathematical relationship with (an asymmetric variant of) the lifting used in [6].

Proposition 6. *Let $\alpha \geq 1$, $\mu_1 \in \mathcal{D}(A)$ and $\mu_2 \in \mathcal{D}(B)$. If $\mu_1 \sim_R^{\mathsf{AD}_\alpha,\delta} \mu_2$ then there exists a sub-distribution $\mu \in \mathcal{D}(A \times B)$ such that: i) $\mathsf{supp}\,(\mu) \subseteq R$; ii) $\pi_1(\mu) \leq \mu_1$; iii) $\pi_2(\mu) \leq \mu_2$ and iv) $\Delta_\alpha(\mu_1, \pi_1\mu) \leq \delta$, where $\leq$ denotes the natural pointwise order on the space of sub-distributions, i.e. $\mu \leq \mu'$ iff $\mu(a) \leq \mu'(a)$ for all a.*

We briefly review some key properties of liftings. The first result characterizes liftings over equivalence relations, and will be used to show that f-divergences can be characterized by our logic.

Proposition 7 (Lifting of equivalence relations). *Let R be an equivalence relation over A and let $\mu_1, \mu_2 \in \mathcal{D}(A)$. Then,*

$$\mu_1 \sim_R^{f,\delta} \mu_2 \iff \Delta_f(\mu_1/R, \mu_2/R) \leq \delta,$$

where μ/R is a distribution over the quotient set A/R, defined as $(\mu/R)([a]) \overset{def}{=} \mu([a])$. In particular, if R is the equality relation $\equiv$, we have

$$\mu_1 \sim_{\equiv}^{f,\delta} \mu_2 \iff \Delta_f(\mu_1, \mu_2) \leq \delta$$

Our next result allows deriving probability claims from lifting judgments. Given $R \subseteq A \times B$ we say that the subsets $A_0 \subseteq A$ and $B_0 \subseteq B$ are *R-equivalent*, and write $A_0 =_R B_0$, iff for every $a \in A$ and $b \in B$, $a \mathrel{R} b$ implies $a \in A_0 \iff b \in B_0$.

Proposition 8 (Fundamental property of lifting). *Let $\mu_1 \in \mathcal{D}(A)$, $\mu_2 \in \mathcal{D}(B)$, and $R \subseteq A \times B$. Then, for any two events $A_0 \subseteq A$ and $B_0 \subseteq B$,*

$$\mu_1 \sim_R^{f,\delta} \mu_2 \land A_0 =_R B_0 \implies \mu_2(B_0)\, f\left(\frac{\mu_1(A_0)}{\mu_2(B_0)}\right) \leq \delta$$

Our final result generalizes the sequential composition theorem from the previous section to arbitrary liftings.

Proposition 9 (Lifting composition). *Let $f_1, f_2, f_3 \in \mathcal{F}$ such that (f_1, f_2) is f_3-composable. Moreover let $\mu_1 \in \mathcal{D}(A)$, $\mu_2 \in \mathcal{D}(B)$, $M_1 : A \to \mathcal{D}(A')$ and $M_2 : B \to \mathcal{D}(B')$. If $\mu_1 \sim_{R_1}^{f_1,\delta_1} \mu_2$ and $M_1(a) \sim_{R_2}^{f_2,\delta_2} M_2(b)$ for all a and b such that $a \mathrel{R} b$, then*

$$(\mathsf{bind}\ \mu_1\ M_1) \sim_{R_2}^{f_3,\delta_1+\delta_2} (\mathsf{bind}\ \mu_2\ M_2)$$

5 A Relational Logic for f-divergences

Building on the results of the previous section, we define a relational logic, called fpRHL, for proving upper bounds for the f-divergence between probabilistic computations written in a simple imperative language.

5.1 Programming Language

We consider programs written in a probabilistic imperative language pWHILE. The syntax of the programming language is defined inductively as follows:

$$
\begin{array}{rll}
\mathcal{C} ::= & \mathsf{skip} & \text{nop} \\
\mid & \mathcal{V} \leftarrow \mathcal{E} & \text{deterministic assignment} \\
\mid & \mathcal{V} \xleftarrow{\$} \mathcal{DE} & \text{random assignment} \\
\mid & \mathsf{if}\ \mathcal{E}\ \mathsf{then}\ \mathcal{C}\ \mathsf{else}\ \mathcal{C} & \text{conditional} \\
\mid & \mathsf{while}\ \mathcal{E}\ \mathsf{do}\ \mathcal{C} & \text{while loop} \\
\mid & \mathcal{C};\ \mathcal{C} & \text{sequence}
\end{array}
$$

Here $\mathcal{V}$ is a set of variables, $\mathcal{E}$ is a set of deterministic expressions, and $\mathcal{DE}$ is a set of expressions that denote distributions from which values are sampled in random assignments. Program states or memories are mappings from variables to values. More precisely, memories map a variable v of type T to a value in its interpretation $[\![T]\!]$. We use $\mathcal{M}$ to denote the set of memories. Programs are interpreted as functions from initial memories to sub-distributions over memories. The semantics, which is given in Figure 2, is based on two evaluation functions $[\![\cdot]\!]_{\mathcal{E}}$ and $[\![\cdot]\!]_{\mathcal{DE}}$ for expressions and distribution expressions; these functions respectively map memories to values and memories to sub-distributions of values. Moreover, the definition uses the operator unit, which maps every $a \in A$ to the unique distribution over A that assigns probability 1 to a and probability 0 to every other element of A, and the null distribution μ_0, that assigns probability 0 to all elements of A. Note that the semantics of programs is a map from memories to sub-distributions over memories. Sub-distributions, rather than distributions, are used to model probabilistic non-termination. However, for the sake of simplicity, in the current development of the logic, we only consider programs that terminate with probability 1 on all inputs and leave the general case for future work.

$$
\begin{array}{lcl}
[\![\mathsf{skip}]\!]\ m & = & \mathsf{unit}\ m \\
[\![c;\ c']\!]\ m & = & \mathsf{bind}\ ([\![c]\!]\ m)\ [\![c']\!] \\
[\![x \leftarrow e]\!]\ m & = & \mathsf{unit}\ (m\,\{[\![e]\!]_{\mathcal{E}}\ m/x\}) \\
[\![x \xleftarrow{\$} \mu]\!]\ m & = & \mathsf{bind}\ ([\![\mu]\!]_{\mathcal{DE}}\ m)\ (\lambda v.\ \mathsf{unit}\ (m\,\{v/x\})) \\
[\![\mathsf{if}\ e\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2]\!]\ m & = & \mathit{if}\ ([\![e]\!]_{\mathcal{E}}\ m = \mathsf{true})\ \mathit{then}\ ([\![c_1]\!]\ m)\ \mathit{else}\ ([\![c_2]\!]\ m) \\
[\![\mathsf{while}\ e\ \mathsf{do}\ c]\!]\ m & = & \lambda f.\ \sup_{n\in\mathbb{N}}\ ([\![[\mathsf{while}\ e\ \mathsf{do}\ c]_n]\!]\ m\ f)
\end{array}
$$

$$
\text{where}\quad
\begin{array}{lcl}
[\mathsf{while}\ e\ \mathsf{do}\ c]_0 & = & \mathit{if}\ ([\![e]\!]_{\mathcal{E}}\ m = \mathsf{true})\ \mathit{then}\ (\mathsf{unit}\ m)\ \mathit{else}\ \mu_0 \\
[\mathsf{while}\ e\ \mathsf{do}\ c]_{n+1} & = & \mathsf{if}\ e\ \mathsf{then}\ c;\ [\mathsf{while}\ e\ \mathsf{do}\ c]_n
\end{array}
$$

Fig. 2. Semantics of programs.

5.2 Judgments

fpRHL judgments are of the form $c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \Phi$, where c_1 and c_2 are programs, Ψ and Φ are relational assertions, $f \in \mathcal{F}$ and $\delta \in \mathbb{R}_0^+$. Relational assertions are first-order formulae over generalized expressions, i.e. expressions in which variables are tagged with a $\langle 1\rangle$ or $\langle 2\rangle$. Relational expressions are interpreted as formulae over pairs of memories, and the tag on a variable is used to indicate whether its interpretation should be taken in the first or second memory. For instance, the relational assertion $x\langle 1\rangle = x\langle 2\rangle$ states that the values of x coincide in the first and second memories. More generally, we use $\equiv$ to denote the relational assertion that states that the values of all variables coincide in the first and second memories.

An fpRHL judgment is valid iff for every pair of memories related by the pre-condition Ψ, the corresponding pair of output distributions is related by the (f,δ)-lifting of the post-condition Φ.

Definition 5 (Validity in fpRHL). *A judgment $c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \Phi$ is* valid*, written $\models c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \Phi$, iff*

$$\forall m_1, m_2 \bullet m_1 \, \Psi \, m_2 \implies (\llbracket c_1 \rrbracket \, m_1) \sim_{\Phi}^{f,\delta} (\llbracket c_2 \rrbracket \, m_2)$$

fpRHL judgments provide a characterization of f-divergence. Concretely, judgments with the identity relation as post-condition can be used to derive (f,δ)-closeness results.

Proposition 10. *If $\models c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \equiv$, then for all memories m_1, m_2,*

$$m_1 \, \Psi \, m_2 \implies \Delta_f(\llbracket c_1 \rrbracket \, m_1, \llbracket c_2 \rrbracket \, m_2) \leq \delta$$

Moreover, fpRHL characterizes continuity properties of probabilistic programs. We assume a continuity model in which programs are executed on random inputs, i.e. distributions of initial memories, and we use f-divergences as metrics to compare program inputs and outputs.

Proposition 11. *Let $f_1, f_2, f_3 \in \mathcal{F}$ such that (f_1, f_2) is f_3-composable. If we have $\models c_1 \sim_{f_2,\delta_2} c_2 : \equiv \Rightarrow \equiv$, then for any two distributions of initial memories μ_1 and μ_2,*

$$\Delta_{f_1}(\mu_1, \mu_2) \leq \delta_1 \implies \Delta_{f_3}(\mathsf{bind}\ \mu_1 \ \llbracket c_1 \rrbracket, \mathsf{bind}\ \mu_2 \ \llbracket c_2 \rrbracket) \leq \delta_1 + \delta_2$$

Finally, we can use judgments with arbitrary post-conditions to relate the probabilities of single events in two programs. This is used, e.g. in the context of game-based cryptographic proofs.

Proposition 12. *If $\models c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \Phi$, then for all memories m_1, m_2 and events E_1, E_2,*

$$m_1 \, \Psi \, m_2 \wedge E_1 =_{\Phi} E_2 \implies \Pr\left[c_2(m_2) : E_2\right] \, f\left(\frac{\Pr\left[c_1(m_1) : E_1\right]}{\Pr\left[c_2(m_2) : E_2\right]}\right) \leq \delta$$

5.3 Proof System

Figure 3 presents a set of core rules for reasoning about the validity of an fpRHL judgment. All the rules are transpositions of rules from apRHL [6]. However, fpRHL rules do no directly generalize their counterparts from apRHL. This is because both logics admit symmetric and asymmetric versions, but apRHL and fpRHL are opposite

$$\frac{\forall m_1, m_2 \bullet m_1\ \mathit{\Psi}\ m_2 \implies (m_1\{[\![e_1]\!]\ m_1/x_1\})\ \mathit{\Phi}\ (m_2\{[\![e_2]\!]\ m_2/x_2\})}{\vdash x_1 \leftarrow e_1 \sim_{f,0} x_2 \leftarrow e_2 : \mathit{\Psi} \Rightarrow \mathit{\Phi}}[\mathsf{assn}]$$

$$\frac{\forall m_1, m_2 \bullet m_1\ \mathit{\Psi}\ m_2 \implies \mathit{\Delta}_f([\![\mu_1]\!]_{\mathcal{DE}}\ m_1, [\![\mu_2]\!]_{\mathcal{DE}}\ m_2) \leq \delta}{\vdash x_1 \xleftarrow{\$} \mu_1 \sim_{f,\delta} x_2 \xleftarrow{\$} \mu_2 : \mathit{\Psi} \Rightarrow x_1\langle 1\rangle = x_2\langle 2\rangle}[\mathsf{rand}]$$

$$\frac{\begin{array}{c}\mathit{\Psi} \implies b\langle 1\rangle \equiv b'\langle 2\rangle \\ \vdash c_1 \sim_{f,\delta} c_1' : \mathit{\Psi} \wedge b\langle 1\rangle \Rightarrow \mathit{\Phi} \qquad \vdash c_2 \sim_{f,\delta} c_2' : \mathit{\Psi} \wedge \neg b\langle 1\rangle \Rightarrow \mathit{\Phi}\end{array}}{\vdash \mathsf{if}\ b\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2 \sim_{f,\delta} \mathsf{if}\ b'\ \mathsf{then}\ c_1'\ \mathsf{else}\ c_2' : \mathit{\Psi} \Rightarrow \mathit{\Phi}}[\mathsf{cond}]$$

$$\frac{\begin{array}{c}(f_1, \ldots, f_n) \text{ composable and monotonic} \\ \mathit{\Theta} \stackrel{\text{def}}{=} b\langle 1\rangle \equiv b'\langle 2\rangle \qquad \mathit{\Psi} \wedge e\langle 1\rangle \leq 0 \implies \neg b\langle 1\rangle \\ \vdash c \sim_{f_1,\delta} c' : \mathit{\Psi} \wedge b\langle 1\rangle \wedge b'\langle 2\rangle \wedge e\langle 1\rangle = k \Rightarrow \mathit{\Psi} \wedge \mathit{\Theta} \wedge e\langle 1\rangle < k\end{array}}{\vdash \mathsf{while}\ b\ \mathsf{do}\ c \sim_{f_n,n\delta} \mathsf{while}\ b'\ \mathsf{do}\ c' : \mathit{\Psi} \wedge \mathit{\Theta} \wedge e\langle 1\rangle \leq n \Rightarrow \mathit{\Psi} \wedge \neg b\langle 1\rangle \wedge \neg b'\langle 2\rangle}[\mathsf{while}]$$

$$\frac{}{\vdash \mathsf{skip} \sim_{f,0} \mathsf{skip} : \mathit{\Psi} \Rightarrow \mathit{\Psi}}[\mathsf{skip}] \qquad \frac{\begin{array}{c}(f_1, f_2) \text{ is } f_3\text{-composable} \\ \vdash c_1 \sim_{f_1,\delta_1} c_2 : \mathit{\Psi} \Rightarrow \mathit{\Phi}' \quad \vdash c_1' \sim_{f_2,\delta_2} c_2' : \mathit{\Phi}' \Rightarrow \mathit{\Phi}\end{array}}{\vdash c_1; c_1' \sim_{f_3,\delta_1+\delta_2} c_2; c_2' : \mathit{\Psi} \Rightarrow \mathit{\Phi}}[\mathsf{seq}]$$

$$\frac{\begin{array}{c}\vdash c_1 \sim_{f,\delta} c_2 : \mathit{\Psi} \wedge \mathit{\Theta} \Rightarrow \mathit{\Phi} \\ \vdash c_1 \sim_{f,\delta} c_2 : \mathit{\Psi} \wedge \neg\mathit{\Theta} \Rightarrow \mathit{\Phi}\end{array}}{\vdash c_1 \sim_{f,\delta} c_2 : \mathit{\Psi} \Rightarrow \mathit{\Phi}}[\mathsf{case}] \qquad \frac{\begin{array}{c}\vdash c_1 \sim_{f',\delta'} c_2 : \mathit{\Psi}' \Rightarrow \mathit{\Phi}' \\ \mathit{\Psi} \Rightarrow \mathit{\Psi}' \quad \mathit{\Phi}' \Rightarrow \mathit{\Phi} \quad f \leq f' \quad \delta' \leq \delta\end{array}}{\vdash c_1 \sim_{f,\delta} c_2 : \mathit{\Psi} \Rightarrow \mathit{\Phi}}[\mathsf{weak}]$$

Fig. 3. Core proof rules.

variants: fpRHL is asymmetric and apRHL is symmetric. Refer to Section 5.4 for a discussion about the symmetric version of fpRHL.

We briefly describe some main rules, and refer the reader to [6] for a longer description about each of them. Rule [seq] relates two sequential compositions and is a direct consequence from the lifting composition (see Proposition 9). Rule [while] relates two loops that terminate in lockstep. The bound depends on the maximal number of iterations of the loops, and we assume given a loop variant e that decreases at each iteration, and is initially upper bounded by some constant n. We briefly explain the side conditions: $(f_1, \ldots, f_n)$ is composable iff (f_i, f_1) is f_{i+1}-composable for every $1 \leq i < n$. Moreover, $(f_1, \ldots, f_n)$ is monotonic iff $f_i \leq f_{i+1}$ for $1 \leq i < n$. Note that the rule is given for $n \geq 2$; specialized rules exist for $n = 0$ and $n = 1$. This rule readily specializes to reason about (ϵ, δ)-differential privacy by taking $f_i = \mathsf{AD}_{\alpha^i}$, where $\alpha = e^\epsilon$.

If an fpRHL judgment is derivable using the rules of Figure 3, then it is valid. Formally,

Proposition 13 (Soundness). *If* $\vdash c_1 \sim_{f,\delta} c_2 : \mathit{\Psi} \Rightarrow \mathit{\Phi}$ *then* $\models c_1 \sim_{f,\delta} c_2 : \mathit{\Psi} \Rightarrow \mathit{\Phi}$.

5.4 Symmetric Logic

One can also define a symmetric version of the logic by adding as an additional clause in the definition of the lift relation that $\Delta_f(\mu_R, \mu_L) \leq \delta$. An instance of this logic is the symmetric apRHL logic from [6]. All rules remain unchanged, except for the random sampling rule that now requires the additional inequality to be checked in the premise of the rule.

6 Conclusion

This paper makes two contributions: first, it unveils a connection between differential privacy and f-divergences. Second, it lays the foundations for reasoning about f-divergences between randomized computations. As future work, we intend to implement support for fpRHL in EasyCrypt [4], and formalize the results from [24]. We also intend to investigate the connection between our notion of lifting and flow networks.

Acknowledgments. This work was partially funded by the European Projects FP7-256980 NESSoS and FP7-229599 AMAROUT, Spanish project TIN2009-14599 DESAFIOS 10 and Madrid Regional project S2009TIC-1465 PROMETIDOS.

References

1. Abate, A.: Approximation metrics based on probabilistic bisimulations for general state-space markov processes: a survey. Electronic Notes in Theoretical Computer Sciences (2012), in Print
2. Ali, S.M., Silvey, S.D.: A general class of coefficients of divergence of one distribution from another. Journal of the Royal Statistical Society. Series B (Methodological) 28(1), 131–142 (1966)
3. Alvim, S., M., Andres, E., M., Chatzikokolakis, K., Palamidessi, C.: On the relation between Differential Privacy and Quantitative Information Flow. In: 38th International Colloquium on Automata, Languages and Programming - ICALP 2011. Lecture Notes in Computer Science, vol. 6756, pp. 60–76. Springer (2011)
4. Barthe, G., Grégoire, B., Heraud, S., Zanella-Béguelin, S.: Computer-aided security proofs for the working cryptographer. In: Advances in Cryptology – CRYPTO 2011. Lecture Notes in Computer Science, vol. 6841, pp. 71–90. Springer, Heidelberg (2011)
5. Barthe, G., Grégoire, B., Zanella-Béguelin, S.: Formal certification of code-based cryptographic proofs. In: 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009. pp. 90–101. ACM, New York (2009)
6. Barthe, G., Köpf, B., Olmedo, F., Zanella-Béguelin, S.: Probabilistic relational reasoning for differential privacy. In: 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012. pp. 97–110. ACM, New York (2012)
7. Chaudhuri, S., Gulwani, S., Lublinerman, R., Navidpour, S.: Proving programs robust. In: 19th ACM SIGSOFT Symposium on the Foundations of Software Engineering and 13rd European Software Engineering Conference, ESEC/FSE 2011. pp. 102–112. ACM, New York (2011)

8. Cortes, C., Mohri, M., Rastogi, A.: Lp distance and equivalence of probabilistic automata. Int. J. Found. Comput. Sci. 18(4), 761–779 (2007)
9. Cortes, C., Mohri, M., Rastogi, A., Riley, M.: On the computation of the relative entropy of probabilistic automata. Int. J. Found. Comput. Sci. 19(1), 219–242 (2008)
10. Csiszár, I.: Eine informationstheoretische ungleichung und ihre anwendung auf den beweis der ergodizitat von markoffschen ketten. Publications of the Mathematical Institute of the Hungarian Academy of Science 8, 85–108 (1963)
11. Deng, Y., Du, W.: Logical, metric, and algorithmic characterisations of probabilistic bisimulation. Tech. Rep. CMU-CS-11-110, Carnegie Mellon University (March 2011)
12. Dwork, C.: Differential privacy. In: 33rd International Colloquium on Automata, Languages and Programming, ICALP 2006. Lecture Notes in Computer Science, vol. 4052, pp. 1–12. Springer, Heidelberg (2006)
13. Ebanks, B., Sahoo, P., Sander, W.: Characterizations of Information Measures. World Scientific (1998)
14. Gaboardi, M., Haeberlen, A., Hsu, J., Narayan, A., Pierce, B.C.: Linear dependent types for differential privacy. In: 40th ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages, POPL 2013. pp. 357–370. ACM, New York (2013)
15. Jonsson, B., Yi, W., Larsen, K.G.: Probabilistic extensions of process algebras. In: Bergstra, J., Ponse, A., Smolka, S. (eds.) Handbook of Process Algebra, pp. 685–710. Elsevier, Amsterdam (2001)
16. McSherry, F.: Privacy integrated queries: an extensible platform for privacy-preserving data analysis. Commun. ACM 53(9), 89–97 (September 2010)
17. McSherry, F.D.: Privacy integrated queries: an extensible platform for privacy-preserving data analysis. In: 35th SIGMOD International Conference on Management of Data, SIGMOD 2009. pp. 19–30. ACM, New York (2009)
18. Pardo, M., Vajda, I.: About distances of discrete distributions satisfying the data processing theorem of information theory. Information Theory, IEEE Transactions on 43(4), 1288 – 1293 (jul 1997)
19. Pierce, B.C.: Differential privacy in the programming languages community (2012), invited tutorial at DIMACS Workshop on Recent Work on Differential Privacy across Computer Science
20. Pierro, A.D., Hankin, C., Wiklicky, H.: Measuring the confinement of probabilistic systems. Theor. Comput. Sci. 340(1), 3–56 (2005)
21. Reed, J., Pierce, B.C.: Distance makes the types grow stronger: a calculus for differential privacy. In: 15th ACM SIGPLAN International Conference on Functional programming, ICFP 2010. pp. 157–168. ACM, New York (2010)
22. Roy, I., Setty, S.T.V., Kilzer, A., Shmatikov, V., Witchel, E.: Airavat: security and privacy for MapReduce. In: 7th USENIX Conference on Networked Systems Design and Implementation, NSDI 2010. pp. 297–312. USENIX Association, Berkeley (2010)
23. Segala, R., Turrini, A.: Approximated computationally bounded simulation relations for probabilistic automata. In: 20th IEEE Computer Security Foundations Symposium, CSF 2007. pp. 140–156. IEEE Computer Society (2007)
24. Steinberger, J.: Improved security bounds for key-alternating ciphers via hellinger distance. Cryptology ePrint Archive, Report 2012/481 (2012), http://eprint.iacr.org/
25. Tracol, M., Desharnais, J., Zhioua, A.: Computing distances between probabilistic automata. In: Proceedings of QAPL. EPTCS, vol. 57, pp. 148–162 (2011)