

Quantum analog of Differential Privacy in term of Rényi divergence.

Léo COLISSON, supervised by Omar FAWZI, team MC2 of LIP

21/08/2016

Le contexte général

La Differential Privacy est un domaine lié à la cryptographie qui se veut de résoudre la problématique suivante : “Comment publier des bases de données utilisables pour des calculs statistiques sans compromettre la vie privée des personnes impliquées dans ces bases de données”. Les solutions proposées consistent à rajouter de l'aléa à la sortie des requêtes. Cependant, les algorithmes (requêtes) sont toujours classiques, et mon maître de stage soupçonne que dans le cas Quantique les théorèmes de composition ne fonctionnent plus aussi bien.

Pour étudier ce problème, nous avons besoin d'adapter la notion de Differential Privacy au cas Quantique, et pour cela il fallait trouver une formulation via la Divergence de Rényi, “distance” naturelle utilisée en informatique Quantique. Or 3 semaines avant mon stage un papier est sorti formalisant justement la Differential Privacy au travers de la Divergence de Rényi...

Le problème étudié

Un autre papier proposait également une troisième formulation de la Differential Privacy avec la KL-Divergence. Dans un premier temps j'ai donc été amené à étudier ces papier pour comprendre les choix qui avaient été fait (les re-formulation n'était pas équivalente à la Differential Privacy).

Une fois cette étude menée, j'ai pu m'intéresser à l'amélioration d'un théorème de composition de requêtes.

Enfin j'ai étudié le cas Quantique en proposant un équivalent Quantique à la Differential Privacy.

La contribution proposée

Pour résoudre ces 3 questions, j'ai été amené :

- A introduire ma propre définition de la Differential Privacy. Même si le résultat ne fut pas très concluant (définition équivalente à la Pure Differential Privacy), il m'a permis de comprendre pourquoi la définition proposée par le premier papier était judicieuse.
- Pour prouver les théorèmes dont j'avais besoin, j'ai écrit un programme Haskell contenant tout le formalisme lié à la Differential Privacy. Son but était de trouver des contre exemples aux théorèmes que je lui proposais. Il m'a été très utile tout au long du stage pour formuler mes conjectures.
- Lors de l'étude du théorème de composition j'ai réussi à trouver une borne meilleure que celle déjà existante en exploitant un lemme proposé par un autre papier couplée à un point de vue “divergence” de la Differential Privacy

- Enfin, sur le peu de temps qu'il me restait à la fin du stage j'ai défini un équivalent Quantique de la Differential Privacy et donné une preuve Quantique du lemme évoqué précédemment.

Les arguments en faveur de sa validité

La borne que j'ai trouvé pour le théorème "Advanced Composition Theorem" est meilleure que celle déjà existante, ce qui est en soit un point positif.

Au niveau de la version Quantique, la divergence que j'ai utilisé est de plus en plus utilisée car elle possède quelques propriétés intéressante (par exemple elle majore les autre divergences), et cette propriété m'a été utile pour faire une première démonstration, ce qui m'a conforté dans ma démarche.

Le bilan et les perspectives

J'ai pu améliorer un résultat déjà connu sur la borne d'un théorème, ce qui est déjà une avancée. Au niveau de la définition Quantique j'aurais aimé avoir plus de temps pour pousser un peu plus loin ma définition. Est-ce que les théorèmes de composition sont valides dans le cas Quantique? Quels exemples d'algorithmes Quantiques peut-on imaginer? Enfin, même si cette question n'était pas directement dans ma direction de recherche, j'aurais aimé arriver à quantifier la Differential Privacy d'un point de vue opérationnel afin d'avoir une vision moins théorique et un peu plus "quantifiée".

L3 Internship report: Quantum analog of Differential Privacy in term of Rényi divergence.

Léo COLISSON, supervised by Omar FAWZI, team MC2 of LIP

June, 2016

Abstract

Here is the report of the six weeks internship I did in 2016 in the MC2 team of the LIP, under the direction of Omar FAWZI. I worked on Differential Privacy, a cryptography concept whose aims is to provide algorithms that try to maximize the accuracy of queries on a statistical database while minimizing the chances of identifying it's records [Wik16]. The goal was to find a similar definition usable with quantum algorithms. To do so, I first tried to understand the link between several definitions of Differential Privacy, expressed in term of statistical distribution (it's the point of view of the initial definition), or in term of divergence (it's more natural to use this notion of distance in quantum algorithms). I introduced my own definition, which appears to be equivalent to pure Differential Privacy. Moreover I found a better bound for the advanced k-fold composition theorem. Finally I introduced a definition of Differential Privacy that applies on Quantum Algorithms.

1	The MC2 team of the LIP	4
1.1	The LIP	4
1.2	MC2	4
2	Several existing definitions...	5
2.1	Intuition	5
2.2	Initial formalization of Differential Privacy	6
2.3	Example of an algorithm	8
2.4	The Divergence point of view	8
2.4.1	Divergence	9
2.4.2	The (Mean) Concentrated Differential Privacy	10
2.4.3	The Zero Concentrated Differential Privacy	11
2.5	Properties of divergence	11
2.5.1	Basic properties	11
2.5.2	Additive property of δ	11
2.5.3	Horizontal Composition	12
3	My first work: studying the classical case	13
3.1	Understanding the Rényi Divergence	13
3.1.1	My conjecture program leads me to my own definition	13
3.1.2	The deception of the Triangle Inequality	13
3.2	The links between the different definitions	15
3.2.1	Link between <i>tCDP</i> and Pure Differential Privacy	15
3.2.2	Link between $(\mathfrak{M}, \mathfrak{C})$ -tCDP and (ξ, ρ) -zCDP	15
3.2.3	Link between $(\mathfrak{M}, \mathfrak{C})$ -tCDP and (μ, τ) -mCDP	16
3.3	The nice properties of these definitions	18
3.3.1	Post processing	18
3.3.2	k-fold composition	18
3.3.3	Use of the smooth divergence to find a better bound to k-fold com- position	19
3.3.4	Trade-off δ/ε	19

4	Extension of Differential Privacy to Quantum Algorithms	20
4.1	Definition	20
4.1.1	State	20
4.1.2	Divergence	20
4.2	Proof of the $\frac{1}{2}\varepsilon^2\alpha$ in the Quantum case	21
4.3	Interesting property	21
5	Conclusion	22
A	Proofs of above theorems	23
A.1	Proof of the Equivalence of distances (Lemma 1)	23
A.2	Proof of the Additive property of δ (Lemma 8)	23
A.3	Proof of the Impossible Triangle Inequality for Rényi Divergence (Theorem 11)	24
A.4	Proof of the Lemma 15	25
A.5	Proof of the lemma Upper bound of D_∞^δ (Lemma 19)	26
A.6	Proof of the New Advanced k-fold composition theorem (Theorem 20)	28
A.7	Proof of the Trade-off δ/ε theorem (Theorem 21)	30
A.8	Proof of the Quantum equivalent of Bounded divergence (Theorem 22)	31

CHAPTER 1

THE MC2 TEAM OF THE LIP

I did my L3 three weeks internship in the LIP, in the MC2 team (Modèles de calcul, Complexité, Combinatoire) under the direction of Omar FAWZI.

1.1 The LIP

The LIP (Laboratoire de l'Informatique du Parallélisme) gathers 57 permanent faculties and researchers, 40-50 PhD students and about 20 scientists working on temporary positions. There are also 12 engineers and helping people. Guillaume Hanrot is the currently chairman of the LIP, while the vice char is Isabelle Guérin-Lassous. The laboratory is situated in the Monod building in ENS Lyon, and is associated with the CNRS, the ENS Lyon, the INRIA, the UCB Lyon 1 and is part of MILYON (Laboratoire d'Excellence "Mathématiques et Informatique à Lyon").

They work on lot's of different topics, linked with computer and information sciences. They are also involved in inter-disciplinary projects. The LIP is divided into 7 teams :

- AriC (Arithmetic and Computing)
- Avalon (Algorithms and Software Architectures for Distributed and HPC Platforms)
- Compsys (Compilation and Embedded Computing Systems)
- DANTE (Dynamic Network)
- MC2 (Models of computation, Complexity, Combinatorics)
- PLUME (programs and proofs)
- ROMA (Resource Optimization : Models, Algorithms and Scheduling)

1.2 MC2

The MC2 (Modèles de calcul, Complexité, Combinatoire) is a LIP team working on complexity theory and on the relevant combinatorial structures. They try to study the limitations of efficient algorithms, as well as their features (sequential vs parallel, synchronous vs asynchronous, deterministic vs probabilistic or quantum). The team, led by Stéphan Thomassé, gathers a dozen of members, 4 doctorants, and 2 administration assistants.

2.1 Intuition

When you want to study a phenomenon, you often need to study databases, but sometimes some private data can be present in these databases. If you are the owner of such a database you may want to publish your database (for example to let researcher use it), but you do not want to injure the people involved in it. For example let's imagine you have a database showing that the people who smoke electronic cigarette often have cancer. If you publish your database "as it", without removing names, anyone could know that Mr XXX has cancer.

A naive approach would be to remove names. However, removing names isn't enough: for example, a famous attack has been able to recover the medical records of the governor of Massachusetts by linking an "anonymized" Netflix database (released for a competition) containing movies rates, with the Internet Movie Database (IMDb) and an anonymised medical encounter data [DR14]. Ideally, we would like that whatever an adversary knows, it couldn't learn more on somebody after looking the database.

The solution adopted by the social sciences when they try to have statistics on an embarrassing questions is to randomize the data. Here is their algorithm: the person that should answer to the embarrassing question toss a coin. If it's head, it answers sincerely. If it's tail, it toss another coin: if it's head it answer "Yes", else "No". By doing that, you can still do some statistical operations, but you cannot have information on one specific person.

That's the point of view adopted by Differential Privacy: providing mathematical guaranties to ensure that it's not possible to obtain "too much" personal information from a database, while allowing people to do statistical operations on it without "too much" error. And to do so, you need randomization. Now, let's see how it works in a more formalized point of view.

2.2 Initial formalization of Differential Privacy

When you use a database, you need to send it queries in order to answer some questions you have: “What is the mean income of this company?”, “How many people in France are minor and have a cellphone?”. . . All these queries can be seen as algorithms that take in input a database and gives back a result. As seen above, it’s useful to randomize the algorithm in order to provide a good privacy, that’s why we need to introduce randomized algorithm.

Definition 1 (Probability Simplex, [DR14] p. 16). *Given a discrete set B , the probability simplex over B , denoted $\Delta(B)$ is defined to be*

$$\Delta(B) := \left\{ x \in \mathbb{R}^{|B|} \mid \forall i, x_i \geq 0, \text{ and } \sum_{i=1}^{|B|} x_i = 1 \right\} \quad (2.1)$$

Definition 2 (Randomized Algorithm, [DR14] p. 16). *A randomized algorithm \mathcal{M} with domain A and discrete range B is associated with a mapping $M : A \rightarrow \Delta(B)$. On input $a \in A$, the algorithm \mathcal{M} outputs $\mathcal{M}(a) = b$ with probability $(M(a))_b$ for each $b \in B$. The probability space is over the coin flips of the algorithm \mathcal{M} .*

Now, let’s formally define a database. A database can be seen in different ways, but here we will use the more mathematical definition. A database is a set of tuples, each tuple representing an entry (a row of a database, e.g. a person can be represented by the tuple $(18, 1, 0)$ to express that he is 18 years old, smoke, but do not have the cancer). We then just count the number of each tuple in an histogram like representation. Let \mathcal{X} be the set of these tuples, a database can then be represented as an element of $\mathbb{N}^{|\mathcal{X}|}$.

Now, we would like that for two very similar databases (with only one person who has a different entry for example), the result of the randomized algorithm stay the same. Indeed, if we have that, it wouldn’t be easy from a result to recover information on a specific person. We first need to define the notion of two “very similar” database by defining a distance between databases:

Definition 3 (Distance between Databases, [DR14] p. 17). *The l_1 norm of a database x denoted $\|x\|_1$ and is defined to be:*

$$\|x\|_1 = \sum_{i=1}^{|\mathcal{X}|} |x_i| \quad (2.2)$$

The l_1 distance between two databases x and x' is $\|x - x'\|_1$.

Now that everything is well defined, we can now introduce the Differential Privacy:

Definition 4 (Differential Privacy, [DR14] p. 17). *A randomized algorithm \mathcal{M} with domain $\mathbb{N}^{|\mathcal{X}|}$ is (ε, δ) -differentially private if for all $\mathcal{S} \subseteq \text{Range}(\mathcal{M})$ and for all $x, x' \in \mathbb{N}^{|\mathcal{X}|}$ such that $\|x - x'\|_1 \leq 1$,*

$$\Pr[\mathcal{M}(x) \in \mathcal{S}] \leq e^\varepsilon \Pr[\mathcal{M}(y) \in \mathcal{S}] + \delta \quad (2.3)$$

The special case when $\delta = 0$ is denoted ε -differentially private and we call it pure differential privacy.

Here the ε and the δ have two very different goals. You can see the ε as “you can be sure that the two distributions are close to a factor of $\varepsilon \dots$ ” and the δ is like “...but sometimes the difference can be huge, it just doesn’t occur more often than with a probability δ ”. That’s why, ideally (in pure differential privacy) we want $\delta = 0$. However it’s not always possible to achieve this goal, so we cannot avoid the δ .

We will also introduce the L_1 distance between two distributions:

Definition 5 (L_1 Distance). *Let’s define the distance between two distributions p and q on Ω :*

$$\Delta(P, Q) := \frac{1}{2} \int_{\mu \in \Omega} |p(\mu) - q(\mu)| d\mu \quad (2.4)$$

NB: The distance $\Delta(P, Q)$ definition is defined in a different manner from one author to the other. For example, the paper [DR14] used this definition (p. 44):

$$\Delta'(P, Q) = \max_S |Pr[P \in S] - Pr[Q \in S]| \quad (2.5)$$

However, all these distances are equivalent:

Lemma 1 (Equivalence of distances). *For all independent distribution P and Q*

$$\Delta(P, Q) = \Delta'(P, Q) \quad (2.6)$$

To see the **proof of this theorem**, please go in appendix.

We can now define the notion of ball:

Definition 6 (Ball). *Let’s define for a distribution P and $\delta > 0$ the ball of radius δ centered at P :*

$$\mathcal{B}_\delta(P) := \{P' : \Delta(P, P') \leq \delta\} \quad (2.7)$$

2.3 Example of an algorithm

Now let's see the example of an algorithm that provides differential privacy. There are plenty of algorithms, here we will study the Laplace Mechanism that transforms a deterministic algorithm $f : \mathbb{N}^{|\mathcal{X}|} \rightarrow \mathbb{R}^k$ into a $(\epsilon, 0)$ -differential privacy algorithm. The idea is to add Laplacian noise, depending on the "variations" of f . The formalized notion of variation is l_1 -sensitivity, which captures the magnitude by which a single individual's data can change the function f :

Definition 7 (l_1 -sensitivity [DR14]). *The l_1 -sensitivity of a function $f : \mathbb{N}^{|\mathcal{X}|} \rightarrow \mathbb{R}^k$ is*

$$\Delta f = \max_{\|x-x'\|_1 \leq 1} \|f(x) - f(x')\|_1 \quad (2.8)$$

Then, the noise added depends on the Laplace Distribution :

Definition 8 (The Laplace Distribution). *The Laplace Distribution (centered at 0) with scale b is the distribution with probability density function:*

$$Lap(x|b) = \frac{1}{2b} \exp\left(-\frac{|x|}{b}\right) \quad (2.9)$$

Definition 9 (The Laplace Mechanism). *Given any function $f : \mathbb{N}^{|\mathcal{X}|} \rightarrow \mathbb{R}^k$, the Laplace mechanism is defined as:*

$$\mathcal{M}_L(x, f(\cdot), \epsilon) = f(x) + (Y_1, \dots, Y_k) \quad (2.10)$$

where Y_i are i.i.d random variables drawn from $Lap(\Delta f / \epsilon)$

Theorem 2 ([DR14] p. 32). *The Laplace mechanism preserves $(\epsilon, 0)$ -differential privacy.*

Some theorems exist (see [DR14] p. 34) to bound the error made with this algorithm. For example, if you want to get the more common name in a list of 10000 names with an accuracy of 95% , you can show that the error is less than 12.2. If the population used for the study count 300,000,000 people it's pretty low !

2.4 The Divergence point of view

In quantum computer science, the natural "distance" between algorithms is formalized with entropy and Rényi divergence. Three weeks before my internship, Mark Bun and Thomas Steinke published a paper where they find an equivalent of differential privacy, named "Zero Concentrated Differential Privacy", and expressed in term of Rényi divergence [BS16]. Their definition was different from another definition of "(Mean) Concentrated Differential Privacy", proposed by Cynthia Dwork and Guy N. Rothblum [DR16]. Let's see some basic definitions about divergence, and then let's see the difference between Mean Differential Privacy and Concentrated Differential Privacy..

2.4.1 Divergence

We define the Rényi divergence as follows:

Definition 10 (Rényi divergence [vEH12]). *For a finite alphabet, and $\alpha \neq 1$, the Rényi divergence of a probability $P = (p_1, \dots, p_n)$ from another distribution $Q = (q_1, \dots, q_n)$ is*

$$D_\alpha(P||Q) = \frac{1}{\alpha - 1} \ln \sum_{i=1}^n p_i^\alpha q_i^{1-\alpha} \quad (2.11)$$

and in case of continuous distributions,

$$D_\alpha(P||Q) = \frac{1}{\alpha - 1} \ln \int p^\alpha q^{1-\alpha} d\mu \quad (2.12)$$

When you compute the limit for $\alpha \rightarrow +\infty$ you find the maximum divergence:

Definition 11 (Max Divergence, [vEH12], [DR14] p. 43). *The Max Divergence between two random variables Y and Z taking values from the same domain is defined to be:*

$$D_\infty = \ln(\text{ess sup}_P \frac{p}{q}) = \max_{S \subseteq \text{supp}(Y)} \left[\ln \frac{\Pr[Y \in S]}{\Pr[Z \in S]} \right] \quad (2.13)$$

We also define the δ -Approximate Max Divergence between Y and Z as follows:

$$D_\infty^\delta(Y||Z) = \max_{S \subseteq \text{Supp}(Y)/\Pr[Y \in S] \geq \delta} \left[\ln \frac{\Pr[Y \in S] - \delta}{\Pr[Z \in S]} \right] \quad (2.14)$$

Here are some useful characterizations of the divergence:

Theorem 3 (Divergence Characterization, [DR14] p. 44).

1. $D_\infty^\delta(Y||Z) \leq \varepsilon$ if and only if there exists a random variable Y' such that $\Delta(Y||Y') \leq \delta$ and $D_\infty(Y'||Z) \leq \varepsilon$
2. We have both $D_\infty^\delta(Y||Z) \leq \varepsilon$ and $D_\infty^\delta(Z||Y) \leq \varepsilon$ if and only if there exist random variables Y', Z' such that

$$\Delta(Y, Y') \leq \frac{\delta}{e^\varepsilon + 1}, \Delta(Z, Z') \leq \frac{\delta}{e^\varepsilon + 1} \text{ and } D_\infty(Y'||Z') \leq \varepsilon \quad (2.15)$$

And we can generate δ -approximate max divergence with a random α :

Definition 12 (Smooth divergence). *For any random variable P and Q defined on Ω and $\delta > 0$, the smooth divergence is defined as follows:*

$$D_\alpha^\delta(P||Q) := \frac{1}{\alpha - 1} \ln \inf_{P' \in \mathcal{B}_\delta(P)} \int_{x \in \Omega} P'(x)^\alpha Q(x)^{1-\alpha} dx$$

These characterization gives a new definition of equivalent to differential privacy:

Theorem 4 (Differential Privacy and Divergence). \mathcal{M} is (ϵ, δ) -differentially private if for all $x, x' / \|x - x'\| \leq 1$,

$$D_{\infty}^{\delta}(\mathcal{M}(x), \mathcal{M}(x')) \leq \epsilon \quad (2.16)$$

Others definitions has been introduced, like the concentrated differential privacy, that allow better accuracy than the classical differential privacy. Let see them.

2.4.2 The (Mean) Concentrated Differential Privacy

Dwork and Rothblum introduced the Concentrated Differential Privacy, renamed as Mean Concentrated Differential Privacy by Mark Bun and Thomas Steinke. Here is how it is defined. First let's define what can be seen as the privacy loss:

Definition 13 (Privacy Loss Random Variable $L_{(Y||Z)}$). For two discrete random variables Y and Z , the privacy loss random variable $L_{(Y||Z)}$, whose range is \mathbb{R} , is distributed by drawing $y \sim Y$ and outputting $\ln \left(\frac{\Pr[Y=y]}{\Pr[Z=y]} \right)$.

Now we would like to compare the privacy loss to Gaussian. That's why we need to define Subgaussian random variable and Subgaussian Divergence:

Definition 14 (Subgaussian Random Variable [Kah60]). A random variable X is τ -subgaussian for a constant $\tau > 0$ if

$$\forall \lambda \in \mathbb{R}, \mathbb{E}[e^{\lambda X}] \leq e^{\frac{\lambda^2 \tau^2}{2}} \quad (2.17)$$

Definition 15 (Subgaussian Divergence and Indistinguishability). For two random variables Y and Z , we say that $D_{subG}(Y||Z) \preceq (\mu, \tau)$ if and only if

- $\mathbb{E}[L_{(Y||Z)}] \leq \mu$
- The centered distribution $(L_{(Y||Z)} - \mathbb{E}[L_{(Y||Z)}])$ is defined and subgaussian, and its subgaussian parameter is at most τ .

Now let's define the Mean Concentrated Differential Privacy (mCDP):

Definition 16 ((Mean) Concentrated Differential Privacy [DR16] p. 10). A randomized algorithm is \mathcal{M} is (μ, τ) -mean concentrated differentially (henceforth (μ, τ) -mCDP) private if for all pairs of adjacent databases x, x' we have $D_{subG}(\mathcal{M}(x)||\mathcal{M}(x')) \preceq (\mu, \tau)$

2.4.3 The Zero Concentrated Differential Privacy

After the Dwork paper on Concentrated Differential Privacy, Mark Bun and Thomas Steinke published another version of Differential Privacy:

Definition 17 (Zero-Concentrated Differential Privacy, [BS16]). *A randomized mechanism \mathcal{M} is (ξ, ρ) -zero-concentrated differentially private (henceforth (ξ, ρ) -zCDP) if, for all $x, x' \in \mathbb{N}^{|\mathcal{X}|}$ such that $\|x - x'\| \leq 1$ and all $\alpha \in (1, +\infty)$,*

$$D_\alpha(\mathcal{M}(x) \|\mathcal{M}(x')) \leq \xi + \rho\alpha \quad (2.18)$$

2.5 Properties of divergence

2.5.1 Basic properties

Here are some interesting properties of divergence that helped me later:

Theorem 5 (Divergence is positive, [vEH12] p. 8). *For all distributions P and Q , and for all order $\alpha \in [0, +\infty]$, $D_\alpha(P \|\mathcal{Q}) \geq 0$.*

Theorem 6 (Nondecreasing of Divergence, [vEH12] p. 6). *For $\alpha \in [0, +\infty]$, and P, Q two distributions, the Rényi divergence $D_\alpha(P \|\mathcal{Q})$ is nondecreasing in α .*

Since all finite distribution gives a bounded divergence, I thought that it was the general case... But it's not, as you can see in the next pictures : the Figure 3.1 shows a bounded divergence, while in Figure 2.1, the divergence isn't bounded.

Another very interesting property of divergence is expressed in [BS16] p. 14. It allowed me to find better bounds for the paper [DR16]:

Theorem 7 (Bounded divergence, [BS16] p. 14). *Let P and Q be probability distributions on Ω satisfying $D_\infty(P \|\mathcal{Q}) \leq \varepsilon$ and $D_\infty(Q \|\mathcal{P}) \leq \varepsilon$. Then for all $\alpha > 1$, $D_\alpha(P \|\mathcal{Q}) \leq \frac{1}{2}\varepsilon^2\alpha$.*

2.5.2 Additive property of δ

This lemma allows us to "chain" the distributions.

Lemma 8 (Additive property of δ). *If it exists x, x' two databases and Y, Y' two distributions such that $\Delta(\mathcal{M}(x) \|\mathcal{Y}) \leq \frac{\delta_0}{e^\varepsilon + 1}$, $\Delta(\mathcal{M}(x') \|\mathcal{Y}') \leq \frac{\delta_0}{e^\varepsilon + 1}$, $D_\infty^\delta(Y \|\mathcal{Y}') \leq \varepsilon_0$ and $D_\infty^\delta(Y' \|\mathcal{Y}) \leq \varepsilon_0$ then*

$$D_\infty^{\delta+\delta_0}(\mathcal{M}(x) \|\mathcal{M}(x')) \leq \varepsilon_0 \quad (2.19)$$

To see the (pretty straight forward) [proof of this theorem](#), please go in appendix.




Figure 2.1 – Divergence of two Gaussian functions such that $\mu_P = 4, \sigma_P = 4$ and $\mu_Q = 5, \sigma_Q = 5$.

2.5.3 Horizontal Composition

The divergence works well with horizontal composition:

Theorem 9 (Horizontal Composition). *If Y_0, Y'_0, Y_1, Y'_1 are independent random variable (which is the case when you run a randomized algorithm) and if*

$$D_\alpha(Y_0||Y'_0) \leq \varepsilon_0 \tag{2.20}$$

and

$$D_\alpha(Y_1||Y'_1) \leq \varepsilon_1 \tag{2.21}$$

then

$$D_\alpha((Y_0, Y_1)|| (Y'_0, Y'_1)) \leq \varepsilon_0 + \varepsilon_1 \tag{2.22}$$

Proof. Just write the definition, separate the factor terms using the fact that the Y_i are independent, apply the ln, and conclude. \square

CHAPTER 3

MY FIRST WORK: STUDYING THE CLASSICAL CASE

Each of these definitions adopts a different point of view. My work, in a first time, was to understand these definitions and try to compare them: are they equivalent ? If not, which one is the more precise ? Is there some cases where the mCDP is better than the zCDP ? Then I studied the composition theorems, and I found a better bound for the Advanced (k-fold) composition theorem that I present below.

3.1 Understanding the Rényi Divergence

3.1.1 My conjecture program leads me to my own definition

Understanding the Rényi Divergence was very important to fully understand the differences between the several definitions of differential privacy. The paper [vEH12] was very helpful for that, but it wasn't enough. To help me in my proofs I decided to code an Haskell program that can check for me the validity of a conjecture (you can see it as a counterexample finder). I implemented every concept defined above in it, to make it able to choose an arbitrary random algorithm, but also perform operations on simplex like computing the mean value of an algorithm. . . I also implemented some functions to find the best possible parameters for the mCDP/zCDP/tCDP functions.

When I ran my program to plot the divergence of a random program, I saw that the Rényi Divergence was always bounded. Since I didn't understand the goal of the zCDP linear approximation (why should I approximate with an linear function something which is bounded ?), I tried to implement my own definition of Concentrated Differential Privacy:

Definition 18 ($(\mathfrak{m}, \mathfrak{z})$ -tCDP). *A mechanism \mathcal{M} is $(\mathfrak{m}, \mathfrak{z})$ -tCDP if for all $\alpha \geq 1$, and for all x, x' such that $\|x - x'\| \leq 1$,*

$$D_\alpha(\mathcal{M}(x) \parallel \mathcal{M}(x')) \leq \min(\mathfrak{m}, \alpha \mathfrak{z}) \quad (3.1)$$

3.1.2 The deception of the Triangle Inequality

Before studying my definition, I wanted to find a kind of triangle inequality for Rényi Divergence. [BS16] propose p. 18 a kind of triangle inequality, based on the Hölder's inequality:

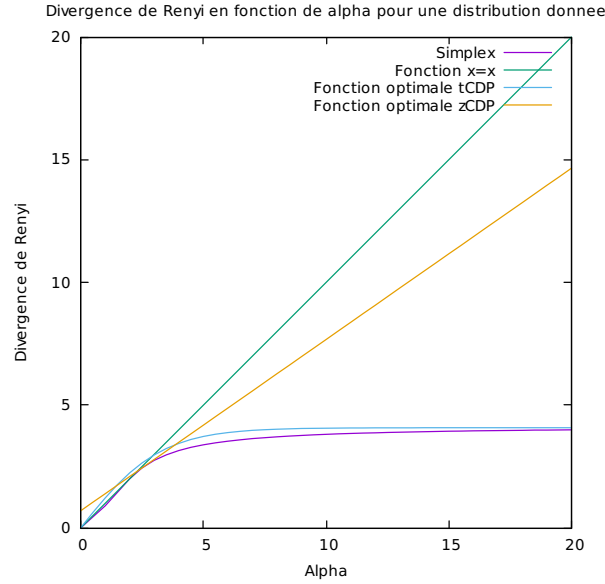


Figure 3.1 – Rényi divergence for a given distribution, with comparison of CDP

Theorem 10 (Triangle-like Inequality for Rényi Divergence, [BS16] p. 18). *Let P , Q , and R be probability distributions. Then for all $k, \alpha > 1$,*

$$D_\alpha(P\|Q) \leq \frac{k\alpha}{k\alpha - 1} D_{\frac{k\alpha-1}{k-1}}(P\|R) + D_{k\alpha}(R\|Q) \quad (3.2)$$

However this theorem does provide a real triangle inequality, so I wanted to know if it was possible to have a real triangle inequality for the Divergence. Unfortunately I showed that it wasn't possible, except in the case of $\alpha = +\infty$:

Theorem 11 (Impossible Triangle Inequality for Rényi Divergence). *It is not possible to have a triangle inequality for the divergence in the general case except if $\alpha = +\infty$.*

More formally, $\forall \alpha > 1, \exists (P_i)_{i \in \mathbb{N}}, (Q_i)_{i \in \mathbb{N}}, (R_i)_{i \in \mathbb{N}} / \exists m \in \mathbb{R} /$

$$\forall n \in \mathbb{N}, D_\alpha(P_n\|Q_n) \leq m, D_\alpha(Q_n\|R_n) \leq m$$

and

$$\lim_{n \rightarrow +\infty} D_\alpha(P_n\|Q_n) = +\infty$$

To see the [proof of this theorem](#), please go in appendix.

3.2 The links between the different definitions

3.2.1 Link between $tCDP$ and Pure Differential Privacy

I finished to understand the interest of the linear approximation of divergence used in $zCDP$. My definition, that tries to bound the divergence was in fact equivalent to pure differential privacy. To show that, we will firstly demonstrate a stronger theorem that will be useful later, and then the theorem will be deduce as a corollary.

Theorem 12. For all $\delta \geq 0, \varepsilon > 0$, and for all distributions P and Q , we have

$$D_{\infty}^{\delta}(P||Q) \leq \varepsilon \Leftrightarrow \forall S, Pr[P \in S] \leq e^{\varepsilon} Pr[Q \in S] + \delta \quad (3.3)$$

Proof. The proof is pretty straight forward when you write the definitions, there is only a little case study. \square

Theorem 13. $(\mathfrak{M}, \mathfrak{X})$ - $tCDP$ is strictly equivalent to pure ε -differential privacy. More precisely, a $(\mathfrak{M}, \mathfrak{X})$ - $tCDP$ is \mathfrak{M} -differentially private, and an algorithm ε -differentially private is $(\varepsilon, \frac{\varepsilon^2}{2})$ - $tCDP$.

Proof. It's a corollary of **Theorem 12** for $\delta = 0$ and **Theorem 7**:

- If \mathcal{M} is $(\mathfrak{M}, \mathfrak{X})$ - $tCDP$, then for all $x, x' / \|x - x'\| \leq 1$, $D_{\infty}(\mathcal{M}(x)||\mathcal{M}(x')) \leq \mathfrak{M}$, so with **Theorem 12**, $\forall S, Pr[\mathcal{M}(x) \in S] \leq e^{\mathfrak{M}} Pr[\mathcal{M}(x') \in S]$, which is the definition of pure \mathfrak{M} -differential privacy
- If \mathcal{M} is ε -differentially private, then the **Theorem 12** tells us that for all $x, x' / \|x - x'\| \leq 1$ $D_{\infty}(\mathcal{M}(x)||\mathcal{M}(x')) \leq \varepsilon$. Since the divergence is nondecreasing (see **Theorem 6**), for all α and for all $x, x' / \|x - x'\| \leq 1$, $D_{\alpha}(\mathcal{M}(x)||\mathcal{M}(x')) \leq \varepsilon$. So the **Theorem 7** let us know that $D_{\alpha}(\mathcal{M}(x)||\mathcal{M}(x')) \leq \frac{\varepsilon^2 \alpha}{2}$. So $D_{\alpha}(\mathcal{M}(x)||\mathcal{M}(x')) \leq \min(\frac{\varepsilon^2 \alpha}{2}, \varepsilon)$, ie \mathcal{M} is $(\varepsilon, \frac{\varepsilon^2}{2})$ -differentially private.

\square

3.2.2 Link between $(\mathfrak{M}, \mathfrak{X})$ - $tCDP$ and (ξ, ρ) - $zCDP$

The link between $(\mathfrak{M}, \mathfrak{X})$ - $tCDP$ and (ξ, ρ) - $zCDP$ is pretty straight forward, since there definition are very closed :

Theorem 14. An algorithm $(\mathfrak{M}, \mathfrak{X})$ - $tCDP$ is $(0, \mathfrak{X})$ - $zCDP$. Reciprocally, An algorithm (ξ, ρ) - $zCDP$ is $(+\infty, \rho)$ - $tCDP$.

3.2.3 Link between $(\mathfrak{M}, \mathfrak{Z})$ -tCDP and (μ, τ) -mCDP

The link between $(\mathfrak{M}, \mathfrak{Z})$ -tCDP and (μ, τ) -mCDP is more complex. Let's begin with a lemma:

Lemma 15. For all $\lambda \in \mathbb{R}$ (where L is the Privacy Loss Random Variable (*Definition 13*) associated with an algorithm $(\mathfrak{M}, \mathfrak{Z})$ -tCDP),

$$\mathbb{E} \left[e^{\lambda(L - \mathbb{E}(L))} \right] \leq \exp \left(\frac{\lambda^2}{2} \left(\mathfrak{M} + \frac{\min(\mathfrak{M}, \mathfrak{Z})}{2} \right)^2 \right) \quad (3.4)$$

To see the **proof of this theorem**, please go in appendix.

Theorem 16. An algorithm $(\mathfrak{M}, \mathfrak{Z})$ -tCDP is $\left(\min(\mathfrak{M}, \mathfrak{Z}), \mathfrak{M} + \frac{\min(\mathfrak{M}, \mathfrak{Z})}{2} \right)$ -mCDP.

Proof. After the proof of the **Lemma 15**, this proof is immediate from the definition. \square

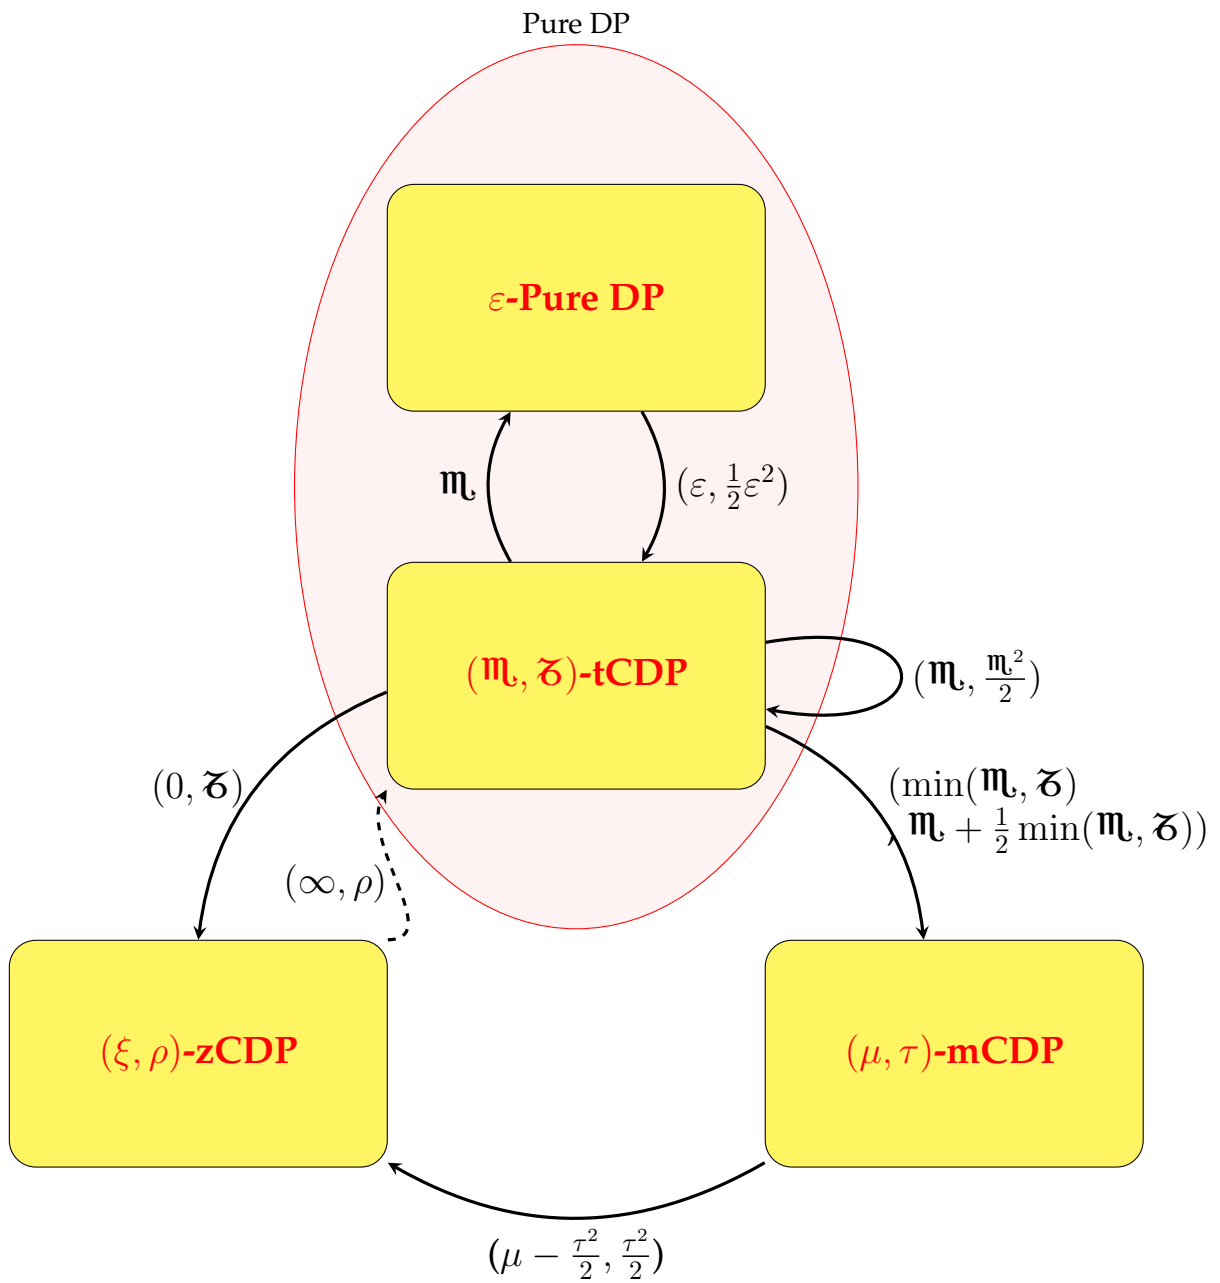


Figure 3.2 – Summary of the links between definitions

3.3 The nice properties of these definitions

Now, after studying the links between the different definitions, I wanted to study the application. There are some property that you would like to ensure with the Differential Privacy. For example you would like that the definition can resist to post-processing, group (I didn't improved the results for groups, so I won't present it here), and k-fold composition. For the k-fold composition, I found a better bound than the one found by Dwork.

3.3.1 Post processing

When you run an algorithm (ϵ, δ) -differentially private, you would like that even after applying to the result new computes, the result still stay (ϵ, δ) -differentially private. That's the goal of this theorem to ensure that:

Theorem 17 (Post-Processing). *Let $\mathcal{M} : \mathbb{N}^{|\mathcal{X}|} \rightarrow \mathbb{R}$ be a randomized algorithm that is (ϵ, δ) -differentially private. Let $f : R \rightarrow R'$ be an arbitrary mapping. Then $f \circ \mathcal{M} : \mathbb{N}^{|\mathcal{M}|} \rightarrow R'$ is (ϵ, δ) -differentially private.*

This concept, pretty natural is already the best one we can expect, and apply for all the different Differential Privacy definitions that we will see.

3.3.2 k-fold composition

We would like to model (and avoid) an attack where the attacker A can choose the requests "in real time". Formally the protocol allowed looks like (where \mathcal{F} is a given set of database access mechanisms):

For $i=1 \dots k$

- A outputs two adjacent databases x_i^0 and x_i^1 , a mechanism $M_i \in \mathcal{F}$ and parameters w_i .
- A receives $y_i \in_R \mathcal{M}_i(w_i, x_{i,b})$

Definition 19 (k-fold composition). *We say that the family \mathcal{F} of database access mechanisms satisfies ϵ -differential privacy under k-fold adaptive composition if for every adversary A , we have $D_\infty(V^0, V^1) \leq \epsilon$ where V^b denotes the view of A in k-fold composition experiment above.*

(ϵ, δ) -differential privacy under k-fold adaptive composition instead requires that $D_\infty^\delta(V^0, V^1) \leq \epsilon$

Here is the main theorem on composition of (ϵ, δ) -differentially private algorithm. It has been updated several times, here was the last version:

Theorem 18 (Advanced (k-fold) Composition, [DR16] p. 2). For all $\varepsilon, \delta, \delta' \geq 0$, the class of (ε, δ) -differentially private mechanisms satisfies $(\varepsilon', k\delta + \delta')$ -differential privacy under k-fold adaptive composition for

$$\varepsilon' = \sqrt{2k \ln(1/\delta')} \varepsilon + k\varepsilon(e^\varepsilon - 1)/2 \quad (3.5)$$

3.3.3 Use of the smooth divergence to find a better bound to k-fold composition

Here I present my proof based on smooth entropy that provide a better bound to the Advanced k-fold composition theorem than the one proposed by Dwork ([Theorem 18](#)). Before, I just need to introduce a little lemma that was inspired by [RW04] (a similar lemma was written for entropy).

Lemma 19 (Upper bound of D_∞^δ). For any random variable P and Q , $\alpha \geq 1$, and $0 < \delta \leq \Delta(P, Q)$

$$D_\infty^\delta(P||Q) \leq D_\alpha(P||Q) - \frac{\ln(\delta)}{\alpha - 1} \quad (3.6)$$

(Note that if $\delta > \Delta(P, Q)$, $D_{+\infty}^\delta(P||Q) = 0$)

To see the [proof of this theorem](#), please go in appendix.

Theorem 20 (New Advanced k-fold composition : better bounds). Let's $\mathcal{M}_1 \dots \mathcal{M}_n$ be n independent algorithms (ε, δ) -differentially private (here independent means that there is no share of random bit between runs). Let $x_1 \dots x_n$ and $x'_1 \dots x'_n$, such as $\forall i, x_i$ and x'_i are neighbour databases. Then the composition of all queries is $\left(\frac{k\varepsilon^2}{2} + \sqrt{2k \ln(1/\delta')} \varepsilon, n\delta + \delta'\right)$ -differentially private. In other words we have :

$$D_\infty^{k\delta + \delta'}(\mathcal{M}_1(x_1) \dots \mathcal{M}_n(x_n) || \mathcal{M}_1(x'_1) \dots \mathcal{M}_n(x'_n)) \quad (3.7)$$

To see the [proof of this theorem](#), please go in appendix.

3.3.4 Trade-off δ/ε

I also proved another theorem by using the [Lemma 19](#). I didn't find yet an application, but it provides a kind of trade-off between ε and δ : if you accept to loose in δ , you can get a better ε .

Theorem 21 (Trade-off δ/ε). Let P and Q be two random variables such that $D_\infty(P||Q) \leq \varepsilon$ and $D_\infty(Q||P) \leq \varepsilon$, and let $0 < \delta \leq \Delta(P, Q)$ (if $\delta > \Delta(P, Q)$, then $D_\infty^\delta(P||Q) = 0$). Then

$$D_\infty^\delta(P||Q) \leq \frac{1}{2} \varepsilon^2 \left(1 + \frac{\varepsilon}{\sqrt{2 \ln(1/\delta)}} \right) \quad (3.8)$$

To see the [proof of this theorem](#), please go in appendix.

CHAPTER 4

EXTENSION OF DIFFERENTIAL PRIVACY TO QUANTUM ALGORITHMS

The initial goal of my internship was to define differential privacy for Quantum Algorithms. I didn't have many time left to do that but I tried to find a few equivalents between the two worlds at the end of my internship.

4.1 Definition

4.1.1 State

In Quantum Algorithms, we work with Quantum state. The more general view of a Quantum state is called the Density Matrix. It represents both the statistical distribution (classical case) and the pure quantum randomness. A density matrix is an operator ("matrix") which is

- positive semidefinite
- hermitian
- with a trace equal to 1

It is represented like this (with $\forall i, \lambda_i > 0$ and $\sum_i \lambda_i = 1$):

$$\rho = \sum_i \lambda_i |\psi_i\rangle\langle\psi_i| \quad (4.1)$$

To have an intuition, when the state is only classical, the matrix is diagonal and has the probabilities on the diagonal (you can notice that the trace is indeed equal to 1). Now let's try to define a Quantum divergence of divergence.

4.1.2 Divergence

In Quantum world the divergence isn't really easy to define, because of one major problem: the operator aren't commutative anymore. So the product $P^\alpha Q^{1-\alpha}$ isn't well defined... A lot of different definitions of Rényi divergence exist, as explained in [Tom15], but I will use one with a special property: the max-divergence is a majorant of all divergence.

Definition 20 (Maximum Quantum Rényi Divergence [Tom15]). For any $\rho, \sigma \in \mathcal{P}(A)$ we define the quantum max-divergence as

$$\mathbb{D}_\alpha(\rho\|\sigma) = \frac{1}{\alpha - 1} \ln \text{Tr}(\sigma^{\frac{1}{2}}(\sigma^{-\frac{1}{2}}\rho\sigma^{-\frac{1}{2}})^\alpha\sigma^{\frac{1}{2}}) \quad (4.2)$$

The limit for $\alpha \rightarrow +\infty$ is

$$\mathbb{D}_\infty(\rho\|\sigma) = \inf\{\lambda : \rho \leq \exp(\lambda)\sigma\} \quad (4.3)$$

For any Rényi divergence \mathbb{D}'_α we have

$$\mathbb{D}'_\alpha(\rho\|\sigma) \leq \mathbb{D}_\alpha(\rho\|\sigma) \quad (4.4)$$

We can then use this Quantum definition of divergence with classical differential privacy definitions. Here is the Quantum equivalent of Zero-Concentrated Differential Privacy:

Definition 21 (Quantum Zero-Concentrated Differential Privacy). A quantum randomized mechanism \mathcal{M} is (ξ, ρ) -quantum-zero-concentrated differentially private (henceforth (ξ, ρ) -qzCDP) if, for all $x, x' \in \mathbb{N}^{|\mathcal{X}|}$ such that $\|x - x'\| \leq 1$ and all $\alpha \in (1, +\infty)$,

$$\mathbb{D}_\alpha(\mathcal{M}(x)\|\mathcal{M}(x')) \leq \xi + \rho\alpha \quad (4.5)$$

4.2 Proof of the $\frac{1}{2}\varepsilon^2\alpha$ in the Quantum case

In the Quantum case you can also use a Quantum equivalent of the [Theorem 7](#):

Theorem 22 (Quantum bounded divergence). Let ρ and σ be probability distributions on $\mathcal{P}(A)$ satisfying $\mathbb{D}_\infty(\rho\|\sigma) \leq \varepsilon$ and $\mathbb{D}_\infty(\sigma\|\rho) \leq \varepsilon$. Then for all $\alpha > 1$

$$\mathbb{D}_\alpha(\rho\|\sigma) \leq \frac{1}{2}\varepsilon^2\alpha \quad (4.6)$$

To see the [proof of this theorem](#), please go in appendix.

4.3 Interesting property

I would like to use the following property (taken from [Tom15]) to proof a similar result to the [Theorem 20](#). However because of lack of time I didn't finish it (since the definition are different (not the same distance...) it's not perfectly straight forward).

Theorem 23 ([Tom15], prop 6.5 p. 99). Let $\rho \in \mathcal{S}(A)$, $0 < \varepsilon < 1$ and $\alpha \in (1, +\infty)$. Then

$$D_{max}^\varepsilon(\rho\|\sigma) \leq \mathbb{D}'_\alpha(\rho\|\sigma) + \frac{g(\varepsilon)}{\alpha - 1} \quad (4.7)$$

where $g(\varepsilon) = -\log(1 - \sqrt{1 - \varepsilon^2})$ and \mathbb{D}'_α is any quantum Rényi divergence.

CHAPTER 5

CONCLUSION

To conclude, during this internship I learned lot's of things in this field of mathematics. I got a better intuition about Rényi divergence and I understood the main differences between the different definitions of differential privacy expressed in term of divergence. I also found a better bound for the Advanced (k-fold) composition theorem.

These definition can be extended to Quantum world, and I think that the proofs can be easily translated from classical world to quantum world, as shown in [Theorem 22](#).

I also learned more practical things, such as how to be well organized: it's easy to get lost in it's note when you are trying to proof something during several days. Differential privacy was for me a totally unknown field, and the subject was crossing both cryptography and quantum computer science, which was very interesting for me. For the first time I felt like if I were a researcher, it was a great experiment ! I would have wanted to spend more time on Quantum divergence, but 6 weeks are quite short to study everything you want. . .

I just want to conclude my report with a special thanks to my supervisor Omar FAWZI, who was always present to help me during this 6 weeks internship and who always guided me toward good idea. Thank you !

APPENDIX A

PROOFS OF ABOVE THEOREMS

A.1 Proof of the Equivalence of distances (**Lemma 1**)

Proof of Lemma 1.

$$\Delta'(P\|Q) = \max_{S \subset \Omega} |Pr[P \in S] - Pr[Q \in S]| \quad (\text{A.1})$$

$$= \max_{S \subset \Omega} \left| \int_S p(\mu) d\mu - \int_S q(\mu) d\mu \right| \quad (\text{A.2})$$

$$= \max_{S \subset \Omega} \left| \int_S p(\mu) - q(\mu) d\mu \right| \quad (\text{A.3})$$

$$= \max_{S \subset \Omega} \left| \int_{\mu \in S: p(\mu) - q(\mu) \geq 0} |p(\mu) - q(\mu)| d\mu - \int_{\mu \in S: p(\mu) - q(\mu) < 0} |q(\mu) - p(\mu)| d\mu \right| \quad (\text{A.4})$$

$$= \max \left(\int_{\mu \in \Omega: p(\mu) - q(\mu) \geq 0} |p(\mu) - q(\mu)| d\mu, \int_{\mu \in \Omega: p(\mu) - q(\mu) < 0} |q(\mu) - p(\mu)| d\mu \right) \quad (\text{A.5})$$

By posing $A = \int_{\mu \in \Omega: p(\mu) - q(\mu) \geq 0} |p(\mu) - q(\mu)| d\mu$ and $B = \int_{\mu \in \Omega: p(\mu) - q(\mu) < 0} |q(\mu) - p(\mu)| d\mu$ we have $\Delta'(P\|Q) = \max(A, B)$ But $A - B = 0$, ie $A = B$, so

$$\Delta'(P\|Q) = \frac{1}{2}(A + B) \quad (\text{A.6})$$

ie

$$\Delta'(P\|Q) = \frac{1}{2} \left(\int_{\mu \in \Omega} |p(\mu) - q(\mu)| d\mu \right) = \Delta(P\|Q) \quad (\text{A.7})$$

□

A.2 Proof of the Additive property of δ (**Lemma 8**)

Proof of Lemma 8. Let's suppose that exists x, x' two databases and Y, Y' two distributions such that $\Delta(\mathcal{M}(x)\|Y) \leq \frac{\delta_0}{e^\varepsilon + 1}$, $\Delta(\mathcal{M}(x')\|Y') \leq \frac{\delta_0}{e^\varepsilon + 1}$, $D_\infty^\delta(Y\|Y') \leq \varepsilon_0$ and $D_\infty^\delta(Y'\|Y) \leq \varepsilon_0$. The characterization theorem (**Theorem 3**) gives us that $\exists Z, Z'$ such that

- $D_\infty(Z, Z') \leq \varepsilon_0$
- $\Delta(Y\|Z) \leq \frac{\delta}{e^{\varepsilon_0+1}}$
- $\Delta(Y'\|Z') \leq \frac{\delta}{e^{\varepsilon_0+1}}$

Still using the characterization theorem (in the other sens this time), we just need to show that

$$\Delta(\mathcal{M}(x)\|Z) \leq \frac{\delta + \delta_0}{e^{\varepsilon_0+1}} \quad (\text{A.8})$$

(the same property with $\mathcal{M}(x)$ is exactly the same demonstration) Let's show it:

$$\Delta(\mathcal{M}(x)\|Z) = \max_S |Pr[\mathcal{M}(x) \in S] - Pr[Z \in S]| \quad (\text{A.9})$$

$$= \max_S |Pr[\mathcal{M}(x) \in S] - Pr[Y \in S] + Pr[Y \in S] - Pr[Z \in S]| \quad (\text{A.10})$$

$$\leq \max_S |Pr[\mathcal{M}(x) \in S] - Pr[Y \in S]| + |Pr[Y \in S] - Pr[Z \in S]| \quad (\text{A.11})$$

$$\leq \frac{\delta + \delta_0}{e^{\varepsilon_0+1}} \quad (\text{A.12})$$

□

A.3 Proof of the Impossible Triangle Inequality for Rényi Divergence (**Theorem 11**)

*Proof of the **Theorem 11**.* Let $\alpha \in (1, +\infty)$, $n \in \mathbb{N}^*$. Let P_n, Q_n, R_n the probability families distributed as:

$$Pr[P_n = 0] = \frac{1}{n^{\frac{\alpha-1}{\alpha}}}; \quad Pr[P_n = 1] = 1 - \frac{1}{n^{\frac{\alpha-1}{\alpha}}} \quad (\text{A.13})$$

$$Pr[Q_n = 0] = \frac{1}{n}; \quad Pr[Q_n = 1] = 1 - \frac{1}{n} \quad (\text{A.14})$$

$$Pr[R_n = 0] = \frac{1}{n^{\frac{\alpha}{\alpha-1}}}; \quad Pr[R_n = 1] = 1 - \frac{1}{n^{\frac{\alpha}{\alpha-1}}} \quad (\text{A.15})$$

Since $\alpha > 1$, $\exists M \in \mathbb{N}$ such that $\forall n \geq M$, $(Pr[P_n = 0], Pr[Q_n = 0], Pr[R_n = 0]) \in [0, 1]^3$. So if needed the sequences can be shifted to be well defined.

Then, we just need to see that for all n

$$D_\alpha(P_n\|Q_n) = \frac{1}{\alpha-1} \ln \left(\left(\frac{1}{n^{\frac{\alpha-1}{\alpha}}} \right)^\alpha \left(\frac{1}{n} \right)^{1-\alpha} + \left(1 - \frac{1}{n^{\frac{\alpha-1}{\alpha}}} \right)^\alpha \left(1 - \frac{1}{n} \right)^{1-\alpha} \right) \quad (\text{A.16})$$

is bounded by a constant m (the left part is equal to 1, and the right part has 1 for limit in $+\infty$ and is continuous so it is bounded). You can show the same thing with $D_\alpha(Q_n\|R_n)$. However, the divergence $D_\alpha(P_n\|R_n)$ isn't bounded :

$$D_\alpha(P_n\|R_n) \leq \frac{1}{\alpha-1} \ln \left(\left(\frac{1}{n^{\frac{\alpha-1}{\alpha}}} \right)^\alpha \left(\frac{1}{n^{\frac{\alpha}{\alpha-1}}} \right)^{1-\alpha} \right) \quad (\text{A.17})$$

$$\leq \frac{1}{\alpha-1} \ln(n) \quad (\text{A.18})$$

So

$$\lim_{n \rightarrow \infty} D_\alpha(P_n\|R_n) = +\infty \quad (\text{A.19})$$

so the triangle inequality isn't possible for all $\alpha \in (1, +\infty)$

□

A.4 Proof of the **Lemma 15**

*Proof of the **Lemma 15**.* Let's $\mathcal{M} : \mathbb{N}^{|\mathcal{X}|} \rightarrow B$ be a $(\mathfrak{m}, \mathfrak{z})$ -differentially private algorithm. To proof this theorem I used the Hoeffding's lemma :

Lemma 24 (Hoeffding's lemma). *Let X be any real-valued random variable with expected value $\mathbb{E}(X) = 0$ and such that $a \leq X \leq b$ almost surely. Then for all $\lambda \in \mathbb{R}$,*

$$\mathbb{E} \left[e^{\lambda X} \right] \leq \exp \left(\frac{\lambda^2 (b-a)^2}{8} \right). \quad (\text{A.20})$$

We would like to use it with

$$X = L - \mathbb{E}(L) \quad (\text{A.21})$$

Let's bound it : Firstly, let's bound $|L|$ on each entry $b \in B$. Let x, x' be two neighbour databases.

$$L_b = \ln \frac{\Pr[\mathcal{M}(x) = b]}{\Pr[\mathcal{M}(x') = b]} \quad (\text{A.22})$$

$$\leq \ln \left(\sup_i \frac{\Pr[\mathcal{M}(x) = b_i]}{\Pr[\mathcal{M}'(x) = b_i]} \right) \quad (\text{A.23})$$

$$= D_\infty(\mathcal{M}(x), \mathcal{M}(x')) \quad (\text{A.24})$$

$$\leq \mathfrak{m} \quad (\text{A.25})$$

In the same idea

$$-L_b = -\ln \frac{\Pr[\mathcal{M}(x) = b]}{\Pr[\mathcal{M}(x') = b]} \quad (\text{A.26})$$

$$= \ln \frac{\Pr[\mathcal{M}(x') = b]}{\Pr[\mathcal{M}(x) = b]} \quad (\text{A.27})$$

so with the same argument we have $\boxed{\forall b, |L_b| \leq \mathfrak{m}_L}$.

We also need a bound on $\mathbb{E}(L)$:

$$\mathbb{E}(L) = \int p \ln \frac{p}{q} d\mu \quad (\text{A.28})$$

$$= D_1(\mathcal{M}(x) \parallel \mathcal{M}(x')) \quad (\text{A.29})$$

$$\leq \min(\mathfrak{m}_L, \mathfrak{z}) \quad (\text{A.30})$$

$$(\text{A.31})$$

Since $\mathbb{E}(L)$ is a divergence, it is also positive, so $\boxed{0 \leq \mathbb{E}(L) \leq \min(\mathfrak{m}_L, \mathfrak{z})}$ Now, by using the two boxed result, we have

$$- \mathfrak{m}_L - \min(\mathfrak{m}_L, \mathfrak{z}) \leq L - \mathbb{E}(L) \leq \mathfrak{m}_L \quad (\text{A.32})$$

so using Hoeffding's lemma we have

$$\mathbb{E} \left[e^{\lambda L - \mathbb{E}(L)} \right] \leq \exp \left(\frac{\lambda^2 (2\mathfrak{m}_L + \min(\mathfrak{m}_L, \mathfrak{z}))^2}{8} \right) \quad (\text{A.33})$$

which conclude the proof. \square

A.5 Proof of the lemma Upper bound of D_∞^δ (**Lemma 19**)

Proof of Lemma 19. Let's begin by noting that if Q is null on a null set, we can just remove this null set from Ω and do the proof as it, the result will stay true with the whole Ω . Now, if Q is null on a non null set, then $D_\alpha(P \parallel Q) = +\infty$ so the result remains true. In the following we can therefore consider that $\forall x \in \Omega, Q(x) \neq 0$.

Let $s_m \geq 1$ and $S = \left\{ x / \frac{P(x)}{Q(x)} \geq s_m \right\}$.

We can define a distribution P' such that $\forall x \in \Omega$,

$$\frac{P'(x)}{Q(x)} \leq s_m \quad (\text{A.34})$$

and

$$\Delta(P, P') = \int_{x \in S} P(x) - P'(x) \quad (\text{A.35})$$

and $\forall x \in S, P'(x) \leq P(x)$. Indeed, the two distributions

$$P_1(x) := s_m Q(x) \text{ if } x \in S, 0 \text{ otherwise} \quad (\text{A.36})$$

$$P_2(x) := s_m Q(x) \quad (\text{A.37})$$

are such that

$$\int_{x \in \Omega} P_1(x) dx \leq 1 \leq \int_{x \in \Omega} P_2(x) dx \quad (\text{A.38})$$

(the inequalities are true because $s_m \geq 1$). So to build P' just define $P'(x) = s_m Q(x)$ if $x \in \Omega$ (this ensure $\Delta(P, P') = \int_{x \in S} P(x) - P'(x)$), and else find values between P_1 and P_2 to ensure that P' is well normed.

We then have

$$\int_{x \in \Omega} P(x)^\alpha Q(x)^{1-\alpha} dx = \int_{x \in \Omega} P(x)^\alpha Q(x)^{1-\alpha} dx \quad (\text{A.39})$$

$$= \int_{x \in \Omega} P(x) \left(\frac{P(x)}{Q(x)} \right)^{\alpha-1} dx \quad (\text{A.40})$$

$$\geq \int_{x \in S} P(x) \left(\frac{P(x)}{Q(x)} \right)^{\alpha-1} dx \quad (\text{A.41})$$

$$\geq s_m^{\alpha-1} \int_{x \in S} P(x) dx \quad (\text{A.42})$$

$$\geq s_m^{\alpha-1} \int_{x \in S} P(x) - P'(x) dx \quad (\text{A.43})$$

$$= s_m^{\alpha-1} \Delta(P, P') \quad (\text{A.44})$$

$$(\text{A.45})$$

We deduce of this inequality the following statement:

$$\frac{1}{\alpha-1} \ln \left(\int_{x \in \Omega} P(x)^\alpha Q(x)^{1-\alpha} dx \right) \geq \frac{1}{\alpha-1} \ln(s_m^{\alpha-1} \Delta(P, P')) \quad (\text{A.46})$$

ie

$$D_\alpha(P||Q) \geq \ln(s_m) + \frac{\ln(\Delta(P, P'))}{\alpha-1} \quad (\text{A.47})$$

Because $\forall x, \frac{P'(x)}{Q(x)} \leq s_m, \ln(s_m) \geq D_\infty(P'||Q)$ so

$$D_\alpha(P||Q) \geq D_\infty(P'||Q) + \frac{\ln(\Delta(P, P'))}{\alpha-1} \quad (\text{A.48})$$

ie

$$D_\infty(P'||Q) \leq D_\alpha(P||Q) - \frac{\ln(\Delta(P, P'))}{\alpha-1} \quad (\text{A.49})$$

And by definition of D_∞^δ , we have

$$D_\infty^\delta(P||Q) \leq D_\alpha(P||Q) - \frac{\ln(\Delta(P, P'))}{\alpha - 1} \quad (\text{A.50})$$

Now, let's show that we can choose P' such that $\Delta(P, P') = \delta$: we would then have

$$D_\infty^\delta(P||Q) \leq D_\alpha(P||Q) - \frac{\ln(\delta)}{\alpha - 1} \quad (\text{A.51})$$

which would end the proof.

To show that P' exists, let's see that the function

$$\begin{aligned} f : [1, +\infty) &\longrightarrow \mathbb{R} \\ s &\longmapsto \Delta(P, P_s) \end{aligned}$$

where P_s is a distribution such that $\forall x, \frac{P_s(x)}{Q(x)} \leq s$ and, by writing $S_s := \{x/\frac{P(x)}{Q(x)} \geq s\}$, $\Delta(P, P_s) = \int_{x \in S_s} P(x) - P_s(x) dx$ (we have already shown at the beginning of the proof that such a distribution could be defined). Then, you can notice that:

- f is continue
- $f(1) = 0$. Indeed, $\forall x, P_s(x) \leq Q(x)$, so in order to have P_s normed to 1, you must have $P_s = Q$, that means $f(1) := \Delta(P_s, Q) = 0$
- $\lim_{s \rightarrow +\infty} f(s) = \Delta(P, Q)$. Indeed, because Q can be consider as non null,, $\lim_{s \rightarrow +\infty} |S_s| = 0$. We then have $\lim_{s \rightarrow +\infty} \Delta(P, P_s) = \lim_{s \rightarrow +\infty} \int_{x \in S_s} P(x) - P'(x) = 0$, so $\lim_{s \rightarrow +\infty} f(s) := \lim_{s \rightarrow +\infty} \Delta(P_s, Q) = \Delta(P, Q)$.

We have $\delta \leq \Delta(P, Q)$, so the intermediate value theorem let us find a s_m such that $f(s_m) := \Delta(P, P_{s_m}) = \delta$.

□

A.6 Proof of the New Advanced k-fold composition theorem (**Theorem 20**)

Proof of Theorem 20. Let's $\mathcal{M}_1 \dots \mathcal{M}_n$ be n independent algorithms (ε, δ) -differentially private (here independent means that there is no share of random bit between runs). Let $x_1 \dots x_n$ and $x'_1 \dots x'_n$ n neighbour databases. For all i there is

$$D_\infty^\delta(\mathcal{M}_i(x_i)||\mathcal{M}_i(x'_i)) \leq \varepsilon \quad (\text{A.52})$$

By characterization $\forall i, \exists P_i, Q_i$ such that $\forall i$,

$$D_\infty^\delta(P_i \| Q_i) \leq \varepsilon \quad (\text{A.53})$$

$$\Delta(\mathcal{M}_i(x), P_i) \leq \frac{\delta}{e^\varepsilon + 1} \quad (\text{A.54})$$

$$\Delta(\mathcal{M}_i(x'), Q_i) \leq \frac{\delta}{e^\varepsilon + 1} \quad (\text{A.55})$$

so with the [Theorem 7](#) we have for all $\alpha > 1$:

$$D_\alpha(P_i \| Q_i) \leq \frac{1}{2} \alpha \varepsilon^2 \quad (\text{A.56})$$

You have then with the [Theorem 9](#) :

$$D_\alpha(P_1 \dots P_n \| Q_1 \dots Q_n) \leq \frac{k\alpha\varepsilon^2}{2} \quad (\text{A.57})$$

Let $\delta' \in (0, \Delta((P_1 \dots P_n), Q_1 \dots Q_n)]$. We can then apply the [Lemma 19](#) :

$$D_\infty^{\delta'}(P_1 \dots P_n \| Q_1 \dots Q_n) \leq D_\alpha(P_1 \dots P_n \| Q_1 \dots Q_n) - \frac{\ln(\delta')}{\alpha - 1} \quad (\text{A.58})$$

We use [A.57](#) :

$$D_\infty^{\delta'}(P_1 \dots P_n \| Q_1 \dots Q_n) \leq \frac{k\alpha\varepsilon^2}{2} - \frac{\ln(\delta')}{\alpha - 1} \quad (\text{A.59})$$

Now, since this is true for all α , let's minimize it in α to find a better bound. Let

$$f(\alpha) = \frac{k\alpha\varepsilon^2}{2} - \frac{\ln(\delta')}{\alpha - 1} \quad (\text{A.60})$$

We derive it :

$$f'(\alpha) = \frac{k\varepsilon^2}{2} + \frac{\ln(\delta')}{(\alpha - 1)^2} \quad (\text{A.61})$$

The minimum is obtained when $f'(\alpha) = 0$:

$$\frac{k\varepsilon^2}{2} + \frac{\ln(\delta')}{(\alpha - 1)^2} = 0 \quad (\text{A.62})$$

So

$$(\alpha - 1)^2 = \frac{2 \ln(1/\delta')}{k\varepsilon^2} \quad (\text{A.63})$$

ie

$$\alpha = 1 + \frac{\sqrt{2 \ln(1/\delta')}}{\sqrt{k\varepsilon}} \quad (\text{A.64})$$

Let's inject it in A.59:

$$D_{\infty}^{\delta'}(P_1 \dots P_n \| Q_1 \dots Q_n) \leq \frac{k \left(1 + \frac{\sqrt{2 \ln(1/\delta')}}{\sqrt{k\varepsilon}}\right) \varepsilon^2}{2} - \frac{\ln(\delta')}{\left(1 + \frac{\sqrt{2 \ln(1/\delta')}}{\sqrt{k\varepsilon}}\right) - 1} \quad (\text{A.65})$$

$$= \frac{k\varepsilon^2}{2} + \sqrt{2k \ln(1/\delta')} \varepsilon \quad (\text{A.66})$$

Now, the Lemma 8 about the additive property of δ let us conclude :

$$D_{\infty}^{n\delta+\delta'}(\mathcal{M}_1(x_1) \dots \mathcal{M}_n(x_n) \| \mathcal{M}'_1(x'_1) \dots \mathcal{M}'_n(x'_n)) \leq \frac{k\varepsilon^2}{2} + \sqrt{2k \ln(1/\delta')} \varepsilon \quad (\text{A.67})$$

□

A.7 Proof of the Trade-off δ/ε theorem (Theorem 21)

Proof of Theorem 21. We will use the result of the Lemma 19. Let P and Q be two random variables such that $D_{\infty}(P\|Q) \leq \varepsilon$ and $D_{\infty}(Q\|P) \leq \varepsilon$, and let $0 < \delta \leq \Delta(P, Q)$. We then have for all $\alpha > 1$:

$$D_{\infty}^{\delta}(P\|Q) \leq D_{\alpha}(P\|Q) - \frac{\ln(\delta)}{\alpha - 1} \quad (\text{A.68})$$

and because $D_{\alpha}(P\|Q) \leq \frac{1}{2}\varepsilon^2\alpha$, we have

$$D_{\infty}^{\delta}(P\|Q) \leq \frac{1}{2}\varepsilon^2\alpha - \frac{\ln(\delta)}{\alpha - 1} \quad (\text{A.69})$$

Let's find the α that minimize this quantity, by derivating the above expression:

$$\frac{d}{d\alpha} \left(\frac{1}{2}\varepsilon^2\alpha - \frac{\ln(\delta)}{\alpha - 1} \right) = \frac{1}{2}\varepsilon^2 + \frac{\ln(\delta)}{(\alpha - 1)^2} \quad (\text{A.70})$$

This is null if and only if

$$(\alpha_m - 1)^2 = -\frac{\varepsilon^2}{2 \ln(\delta)} \quad (\text{A.71})$$

If we suppose $\delta < 1$, we have

$$\alpha_m = 1 + \sqrt{\frac{\varepsilon^2}{2 \ln(1/\delta)}} = 1 + \frac{\varepsilon}{\sqrt{2 \ln(1/\delta)}} \quad (\text{A.72})$$

This is greater than 1, so the minimum is

$$D_{\infty}^{\delta}(P\|Q) \leq \frac{1}{2}\varepsilon^2 \left(1 + \frac{\varepsilon}{\sqrt{2 \ln(1/\delta)}} \right) \quad (\text{A.73})$$

□

A.8 Proof of the Quantum equivalent of Bounded divergence (Theorem 22)

Proof of Theorem 22. Let P and Q be probability distributions on $\mathcal{P}(A)$ satisfying

$$\mathbb{D}_\infty(P\|Q) \leq \varepsilon \tag{A.74}$$

and $\mathbb{D}_\infty(Q\|P) \leq \varepsilon$. Let $\alpha > 1$. As explained in [Tom15], we can bound \mathbb{D}_α by the max-divergence. Let's rewrite the unitary operator on the eigenvectors basis : $\sigma^{-\frac{1}{2}}\rho\sigma^{-\frac{1}{2}} = \sum_i t_i |\psi_i\rangle\langle\psi_i|$. The Equation A.74 gives us $\forall i, t_i \leq e^\varepsilon$. Moreover, by definition of \mathbb{D}_∞ we have

$$\sigma \leq e^\varepsilon \rho \tag{A.75}$$

By multiplying by $\sigma^{-\frac{1}{2}}$ on the right and on the left, we have :

$$e^{-\varepsilon} \leq \sigma^{-\frac{1}{2}}\rho\sigma^{-\frac{1}{2}} \tag{A.76}$$

so with the same argument as above, we also have $\forall i, t_i \geq e^{-\varepsilon}$. So $\forall i, \exists \lambda_i$ such that

$$t_i = \lambda_i e^\varepsilon + (1 - \lambda_i) e^{-\varepsilon} \tag{A.77}$$

Let's inject it in the definition of \mathbb{D}_α :

$$e^{(\alpha-1)\mathbb{D}_\alpha(\rho\|\sigma)} \leq \text{Tr}[\sigma^{\frac{1}{2}}(\sigma^{-\frac{1}{2}}\rho\sigma^{-\frac{1}{2}})^\alpha\sigma^{\frac{1}{2}}] \tag{A.78}$$

$$= \text{Tr}[\sigma^{\frac{1}{2}}(\sum_i (\lambda_i e^\varepsilon + (1 - \lambda_i) e^{-\varepsilon}) |\psi_i\rangle\langle\psi_i|)^\alpha\sigma^{\frac{1}{2}}] \tag{A.79}$$

$$= \text{Tr}[\sigma^{\frac{1}{2}} \sum_i (\lambda_i e^\varepsilon + (1 - \lambda_i) e^{-\varepsilon})^\alpha |\psi_i\rangle\langle\psi_i| \sigma^{\frac{1}{2}}] \text{ (orthonormal basis)} \tag{A.80}$$

$$= \sum_i (\lambda_i e^\varepsilon + (1 - \lambda_i) e^{-\varepsilon})^\alpha \text{Tr}[\sigma |\psi_i\rangle\langle\psi_i|] \tag{A.81}$$

$$\leq \sum_i \lambda_i e^{\varepsilon\alpha} + (1 - \lambda_i) e^{-\varepsilon\alpha} \text{Tr}[\sigma |\psi_i\rangle\langle\psi_i|] \text{ (Jensen)} \tag{A.82}$$

$$= e^{\alpha\varepsilon} \sum_i \lambda_i \text{Tr}[\sigma |\psi_i\rangle\langle\psi_i|] + e^{-\alpha\varepsilon} \sum_i (1 - \lambda_i) \text{Tr}[\sigma |\psi_i\rangle\langle\psi_i|] \tag{A.83}$$

By posing $A = \sum_i \lambda_i \text{Tr}[\sigma |\psi_i\rangle\langle\psi_i|]$ and $B = \sum_i (1 - \lambda_i) \text{Tr}[\sigma |\psi_i\rangle\langle\psi_i|]$ we then have

$$e^{(\alpha-1)\mathbb{D}_\alpha(\rho\|\sigma)} \leq e^{\alpha\varepsilon} A + e^{-\alpha\varepsilon} B \tag{A.84}$$

and $A + B = 1$. By evaluating [A.84](#) with $\alpha = 1$, we have two equations for two unknown variables (A and B): we can solve it into

$$A = \frac{1 - e^{-\varepsilon}}{e^\varepsilon - e^{-\varepsilon}} \tag{A.85}$$

$$B = \frac{e^{-\varepsilon} - 1}{e^\varepsilon - e^{-\varepsilon}} \tag{A.86}$$

$$\tag{A.87}$$

The end of the proof is like the classical case in [\[BS16, p. 14\]](#): we inject the above A and B in [A.84](#), express the exp in term of sinh, and conclude with the inequality

$$\frac{\sinh(x) - \sinh(y)}{\sinh(x - y)} \leq e^{\frac{1}{2}xy} \tag{A.88}$$

We then have

$$\mathbb{D}_\alpha(\rho \parallel \sigma) \leq \frac{1}{2} \varepsilon^2 \alpha \tag{A.89}$$

□

- [BS16] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. *CoRR*, abs/1605.02065, 2016.
- [DR14] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [DR16] Cynthia Dwork and Guy N. Rothblum. Concentrated differential privacy. *CoRR*, abs/1603.01887, 2016.
- [Kah60] J. Kahane. Propriétés locales des fonctions à séries de fourier aléatoires. *Studia Mathematica*, 19(1):1–25, 1960.
- [RW04] R. Renner and S. Wolf. Smooth renyi entropy and applications. In *Information Theory, 2004. ISIT 2004. Proceedings. International Symposium on*, pages 233–, June 2004.
- [Tom15] Marco Tomamichel. Quantum information processing with finite resources - mathematical foundations. 2015.
- [vEH12] Tim van Erven and Peter Harremoës. Rényi divergence and kullback-leibler divergence. *CoRR*, abs/1206.2459, 2012.
- [Wik16] Wikipedia. Differential privacy — Wikipedia, the free encyclopedia, 2016. [Online; accessed June-2016].