Quantum Differential Privacy: An Information Theory Perspective

Christoph Hirche, Cambyse Rouzé, Daniel Stilck França

Quantum Differential Privacy: An Information Theory Perspective

Christoph Hirche, Cambyse Rouzé, Daniel Stilck França

***Abstract*—Differential privacy has been an exceptionally successful concept when it comes to providing provable security guarantees for classical computations. More recently, the concept was generalized to quantum computations. While classical computations are essentially noiseless and differential privacy is often achieved by artificially adding noise, near-term quantum computers are inherently noisy and it was observed that this leads to natural differential privacy as a feature.**

In this work we discuss quantum differential privacy in an information theoretic framework by casting it as a quantum divergence. A main advantage of this approach is that differential privacy becomes a property solely based on the output states of the computation, without the need to check it for every measurement. This leads to simpler proofs and generalized statements of its properties as well as several new bounds for both, general and specific, noise models. In particular, these include common representations of quantum circuits and quantum machine learning concepts. Here, we focus on the difference in the amount of noise required to achieve certain levels of differential privacy versus the amount that would make any computation useless. Finally, we also generalize the classical concepts of local differential privacy, Rényi differential privacy and the hypothesis testing interpretation to the quantum setting, providing several new properties and insights.

I. Introduction

PROCESSING data in some form is the core concept of most computational tasks. Nowadays, large data sets are being collected and processed for a variety of tasks ranging from medical studies to machine learning applications. With the accumulation of information always also come security concerns. If one presents the results of a study, will that allow the audience to conclude on the health of a particular individual? Social media companies constantly process our data using machine learning for advertisement purposes. How much do the results reveal about the underlying data set?

C. Hirche is with the Zentrum Mathematik, Technical University of Munich, 85748 Garching, Germany and the Centre for Quantum Technologies, National University of Singapore, Singapore

C. Rouzé is with the Munich Center for Quantum Science and Technology (MCQST), München, Germany

D. S. França is with QMATH, Department of Mathematical Sciences, University of Copenhagen, Universitetsparken 5, 2100 Copenhagen, Denmark

Differential privacy [16], [15] is a concept introduced in the classical computation setting to address these concerns. Vaguely speaking, differential privacy guarantees that the probability of an algorithm giving a certain outcome is roughly the same for any sufficiently similar input. This implies that good differential privacy makes it difficult for an observer to make precise statements about the data used.

With the growing interest in large-scale quantum computations and quantum machine learning applications, naturally, also security requirements are desired there. To that end, the concept of quantum differential privacy was introduced in [40] in the setting of quantum computing. That work was the starting point of a series of results connecting quantum differential privacy to gentle measurements [1], distributed quantum computing [23] and quantum machine learning [4], [32], [36], [12], [2], [3], [13]. Other authors also studied differential privacy in the classical to quantum regime [39].

In quantum differential privacy, we are interested in the properties of a quantum algorithm $\mathcal{A}$ represented by a general quantum channel. The input data is a quantum state ρ and someone observing the output should not be able to determine whether the input state was indeed ρ or a similar state σ. Similarity is usually defined via neighbouring states, denoted $\rho \sim \sigma$, according to some rule. Different notions of similarity have been proposed in the literature. For instance, a small trace distance [40] or convertibility by a local quantum channel [1]. For most of our discussion we will keep the definition unspecified and only sometimes fix it for examples. Now, a quantum algorithm $\mathcal{A}$ is (ϵ, δ)-differentially private, if for all measurements M and all neighbouring quantum states $\rho \sim \sigma$ we have

$$\operatorname{Tr}(M\mathcal{A}(\rho)) \leq e^{\epsilon} \operatorname{Tr}(M\mathcal{A}(\sigma)) + \delta, \tag{I.1}$$

see Definition III.1 for full details. In the classical setting, one way to achieve differential privacy is by taking an algorithm and adding noise to the output to obscure the input [17]. This idea was transferred to the quantum setting in [40], showing that concatenating an algorithm with sufficiently strong noise in form of a different quantum channel makes the algorithm differentially private. This

was further explored in [12] for layered algorithms that are affected by noise at every step, a model that describes many common scenarios such as quantum circuits and quantum machine learning algorithms implemented on noisy quantum computers. An interesting proposition following from [12] is that the noise present in near-term quantum devices, while presenting a computational difficulty, induces an inherent advantage by naturally making the computation differentially private. On the other hand it is of course also well known that such noise can make it impossible to run long computations, somewhat limiting the computational usefulness of near-term devices. One of the main goals of this work is to give upper bounds on the depth required to reach differential privacy and contrasting it to the computational limitations imposed by the noise. As we will see, differential privacy can be reached at a significantly shorter depth than that at which computationally prohibitive noise occurs.

To this end, we introduce an information theoretic approach to quantum differential privacy that allows us to cast the requirement in terms of a quantum divergence. This in turn lets us make several observations that are solely based on the properties of the divergence, giving a fruitful new approach to discussing quantum differential privacy. The divergence of choice is the quantum hockey-stick divergence introduced in [33] and the general framework follows classical work presented among others in [6], [5] using tools such as contraction coefficients to bound differential privacy in iterative algorithms.

The outline and main contributions of this work are as follows.

- In Section II we revisit the quantum hockey-stick divergence and show several new properties. In particular we consider its associated contraction coefficient and give a simple expression for it that significantly reduces its computational complexity. These results lay the foundation for what follows, but should also be of independent interest.
- In Section III we discuss quantum differential privacy and how to cast it in terms of the hockey-stick divergence and a variation of the smooth $\max$-relative entropy. Based on this we then give properties and implications of quantum differential privacy including a general bound showing exponential decay of δ with the algorithm depth. This also gives an operational interpretation of the hockey-stick divergence.
- In Section IV we discuss specific noise models including global and local depolarizing noise which we contrast with each other but also with the induced computational limitations. In particular, we get a separation for depolarizing noise where the trace distance decays exponentially, but good differential privacy is reached after a finite number of steps. Furthermore we discuss more general models such as arbitrary qubit noise.
- In Section V We introduce quantum generalisations of local differential privacy and Rényi differential privacy: two often discussed extensions of the standard differential privacy definition. Finally, we discuss the hypothesis testing interpretation of quantum differential privacy and derive a useful trade-off between ϵ and δ from it.

Notations: A (classical or quantum) system R is associated with a finite-dimensional Hilbert space $\mathcal{H}_R$. Let $\mathcal{P}(\mathcal{H}_R)$ be the set of positive semidefinite linear operators acting on $\mathcal{H}_R$. A quantum state ρ_R on R is a positive semidefinite linear operator with unit trace acting on $\mathcal{H}_R$, denoted $\rho_R \in \mathcal{S}_=(\mathcal{H}_R)$. The set of subnormalized states is denoted $\mathcal{S}_\leq(\mathcal{H}_R)$. A state ρ_R of rank 1 is called pure, and we may choose a normalized vector $|\psi\rangle_R \in \mathcal{H}_R$ satisfying $\rho_R = |\psi\rangle\langle\psi|_R$. Otherwise, ρ_R is called a mixed state. By the spectral theorem, every mixed state can be written as a convex combination of pure states. For a pure state $|\phi\rangle$ we may use the shorthand $\phi \equiv |\phi\rangle\langle\phi|$. For a classical system X there is a distinguished orthonormal basis $\{|x\rangle\}_{x=1}^{\dim \mathcal{H}_X}$ of $\mathcal{H}_X$ diagonalizing every state on X.

A quantum channel $\mathcal{N}\colon A \to B$ is a linear completely positive and trace-preserving map from the operators on $\mathcal{H}_A$ to the operators on $\mathcal{H}_B$. Given a quantum channel $\mathcal{M}$ with input and output dimension d, its Choi matrix is defined as

$$C_{\mathcal{M}} := \sum_{i,j\in[d]} \mathcal{M}(|i\rangle\langle j|) \otimes |i\rangle\langle j|\,,$$

where $\{|i\rangle : i \in [d]\}$ is the standard basis of $\mathbb{C}^d$. In the following we will usually drop the indices as the systems are clear from context. A measurement is an operator $0 \leq M \leq \mathbb{1}$ and a collection of measurements such that $\sum_i M_i = \mathbb{1}$ is called a POVM.

II. The quantum hockey-stick divergence

In this section we discuss the main technical tools needed for our investigation of quantum differential privacy. The quantum hockey-stick divergence was first introduced in [33], in the context of exploring strong converse bounds for the quantum capacity, as

$$E_\gamma(\rho\|\sigma) := \operatorname{Tr}(\rho - \gamma\sigma)^+, \tag{II.1}$$

for $\gamma \geq 1$. Here X^+ denotes the positive part of eigen-decomposition of a hermitian matrix $X = X^+ - X^-$. In [33] it was noted that this quantity is closely related to the trace norm via

$$E_\gamma(\rho\|\sigma) = \frac{1}{2}\|\rho - \gamma\sigma\|_1 + \frac{1}{2}(\operatorname{Tr}(\rho) - \gamma\operatorname{Tr}(\sigma)), \tag{II.2}$$

so, if $\rho, \sigma \in \mathcal{S}_=(\mathcal{H})$, $E_1(\rho\|\sigma) = \frac{1}{2}\|\rho - \sigma\|_1$ equals the trace distance. As the trace distance has many desirable properties, we are tempted to hope that similar properties also hold for the quantum hockey-stick divergence. For instance, the trace distance is invariant under unitaries and Eq. (II.2) immediately implies that the same holds true for the hockey-stick divergence.

Another important example will be the following, relating the divergence to a maximization over measurements.

Lemma II.1. *An alternative expression for the hockey-stick divergence for $\rho, \sigma \in \mathcal{P}(\mathcal{H})$ is given by*

$$E_\gamma(\rho\|\sigma) = \max_{0 \le \Lambda \le \mathbb{1}} \operatorname{Tr}\{\Lambda(\rho - \gamma\sigma)\}. \tag{II.3}$$

Proof. The proof is similar to the standard argument for the trace distance. Let $X = \rho - \gamma\sigma$ and (X^+, X^-) its decomposition into positive and negative parts such that $X = X^+ - X^-$. For a general operator $0 \le \Lambda \le \mathbb{1}$ one easily sees

$$\begin{aligned} &\operatorname{Tr}\{\Lambda(\rho - \gamma\sigma)\} = \operatorname{Tr}\{\Lambda(X^+ - X^-)\} \\ &\le \operatorname{Tr}\{\Lambda X^+\} \le \operatorname{Tr} X^+ = E_\gamma(\rho\|\sigma). \end{aligned}$$

It remains to show that equality in the above can be achieved by some measurement. For that simply pick Π_{X^+}, the projector onto the support of X^+ and observe that

$$\begin{aligned} &\operatorname{Tr}\{\Pi_{X^+}(\rho - \gamma\sigma)\} = \operatorname{Tr}\{\Pi_{X^+}(X^+ - X^-)\} \\ &= \operatorname{Tr}\{\Pi_{X^+} X^+\} = \operatorname{Tr} X^+ = E_\gamma(\rho\|\sigma). \end{aligned}$$

This concludes the proof. □

A property that was already shown in [33] is that E_γ, like any good divergence, obeys data-processing, meaning for any quantum channel $\mathcal{N}$ we have

$$E_\gamma(\mathcal{N}(\rho)\|\mathcal{N}(\sigma)) \le E_\gamma(\rho\|\sigma),$$

which even holds for any $\rho, \sigma \in \mathcal{P}(\mathcal{H})$, see [33, Lemma 4]. This makes it meaningful to define its contraction coefficient as

$$\eta_\gamma(\mathcal{N}) := \sup_{\rho,\sigma \in \mathcal{S}_=(\mathcal{H})} \frac{E_\gamma(\mathcal{N}(\rho)\|\mathcal{N}(\sigma))}{E_\gamma(\rho\|\sigma)}, \tag{II.4}$$

where the optimization is over $\rho, \sigma \in \mathcal{S}_=(\mathcal{H})$, and obviously $0 \le \eta_\gamma(\mathcal{N}) \le 1$. For a recent overview over contraction coefficients and their properties see [20].

Interestingly, the contraction coefficient for the trace distance simplifies significantly. Instead of having to optimize over arbitrary initial states ρ, σ, it suffices to only consider orthogonal pure states [30]. We will prove now that an analogous result holds also for the hockey-stick divergence. This generalizes both the proof for the trace distance in [30] and that for the classical hockey-stick divergence in [6, Theorem 3].

Theorem II.2. *The hockey-stick divergence contraction coefficient can be equivalently expressed as*

$$\eta_\gamma(\mathcal{N}) = \sup_{|\varphi\rangle \perp |\psi\rangle} E_\gamma(\mathcal{N}(|\varphi\rangle\langle\varphi|)\|\mathcal{N}(|\psi\rangle\langle\psi|)). \tag{II.5}$$

Proof. Let $X = \rho - \gamma\sigma$ and (X^+, X^-) its decomposition into positive and negative parts such that $X = X^+ - X^-$. Note that by definition and Equation (II.2) we have

$$\begin{aligned} E_\gamma(\rho\|\sigma) &= \operatorname{Tr} X^+ = \frac{1}{2}\|X\|_1 + \frac{1}{2}(1-\gamma), \\ E_\gamma(\mathcal{N}(\rho)\|\mathcal{N}(\sigma)) &= \frac{1}{2}\|\mathcal{N}(X)\|_1 + \frac{1}{2}(1-\gamma). \end{aligned} \tag{II.6}$$

Further, define $\hat{X}^+ = \frac{X^+}{\operatorname{Tr} X^+}$ and $\hat{X}^- = \frac{X^-}{\operatorname{Tr} X^-}$ with spectral decompositions

$$\begin{aligned} \hat{X}^+ &= \sum_m p_m |m\rangle\langle m|, \\ \hat{X}^- &= \sum_n q_n |n\rangle\langle n|. \end{aligned}$$

With these definitions, we observe

$$\begin{aligned} &\|\mathcal{N}(X)\|_1 \qquad\qquad \text{(II.7)} \\ &= \|\mathcal{N}(X^+) - \mathcal{N}(X^-)\|_1 \\ &= \|\operatorname{Tr}(X^+)\mathcal{N}(\hat{X}^+) - \operatorname{Tr}(X^-)\mathcal{N}(\hat{X}^-)\|_1 \\ &= \Big\| \operatorname{Tr}(X^+)\mathcal{N}(\sum_m p_m |m\rangle\langle m|) \\ &\qquad - \operatorname{Tr}(X^-)\mathcal{N}(\sum_n q_n |n\rangle\langle n|) \Big\|_1 \\ &= \Big\| \sum_{m,n} p_m q_n \left(\operatorname{Tr}(X^+)\mathcal{N}(|m\rangle\langle m|) \right. \\ &\qquad \left. - \ \operatorname{Tr}(X^-)\mathcal{N}(|n\rangle\langle n|) \right) \Big\|_1 \\ &\le \sum_{m,n} p_m q_n \left\| \operatorname{Tr}(X^+)\mathcal{N}(|m\rangle\langle m|) \right. \\ &\qquad \left. - \operatorname{Tr}(X^-)\mathcal{N}(|n\rangle\langle n|) \right\|_1 \\ &\le \sup_{m,n} \|\operatorname{Tr}(X^+)\mathcal{N}(|m\rangle\langle m|) - \operatorname{Tr}(X^-)\mathcal{N}(|n\rangle\langle n|)\|_1. \end{aligned} \tag{II.8}$$

We continue with

$$\begin{aligned} &\|\operatorname{Tr}(X^+)\mathcal{N}(|m\rangle\langle m|) - \operatorname{Tr}(X^-)\mathcal{N}(|n\rangle\langle n|)\|_1 \\ \le &\|\operatorname{Tr}(X^+)\mathcal{N}(|m\rangle\langle m|) - \gamma \operatorname{Tr}(X^+)\mathcal{N}(|n\rangle\langle n|)\|_1 \\ &+ \|\gamma \operatorname{Tr}(X^+)\mathcal{N}(|n\rangle\langle n|) - \operatorname{Tr}(X^-)\mathcal{N}(|n\rangle\langle n|)\|_1 \\ = &\operatorname{Tr}(X^+)\|\mathcal{N}(|m\rangle\langle m|) - \gamma\mathcal{N}(|n\rangle\langle n|)\|_1 \\ &+ \|((\gamma-1)\operatorname{Tr}(X^+) + (1-\gamma))\mathcal{N}(|n\rangle\langle n|)\|_1 \\ = &\operatorname{Tr}(X^+)\|\mathcal{N}(|m\rangle\langle m|) - \gamma\mathcal{N}(|n\rangle\langle n|)\|_1 \\ &- ((\gamma-1)\operatorname{Tr}(X^+) + (1-\gamma)) \end{aligned}$$

$$
\begin{aligned}
&= \operatorname{Tr}(X^+)(\|\mathcal{N}(|m\rangle\langle m|) - \gamma\mathcal{N}(|n\rangle\langle n|)\|_1 \\
&\quad - (\gamma - 1)) - (1 - \gamma) \\
&= 2\operatorname{Tr}(X^+)E_\gamma(\mathcal{N}(|m\rangle\langle m|)\|\mathcal{N}(|n\rangle\langle n|)) - (1 - \gamma) \\
&= 2E_\gamma(\rho\|\sigma)E_\gamma(\mathcal{N}(|m\rangle\langle m|)\|\mathcal{N}(|n\rangle\langle n|)) - (1 - \gamma),
\end{aligned}
$$

where the first inequality is by triangle inequality, the first equality because $\operatorname{Tr}(X^+) \geq 0$ and the second because $((\gamma - 1)\operatorname{Tr}(X^+) + (1 - \gamma)) \leq 0$ since $\operatorname{Tr}(X^+) = E_\gamma(\rho\|\sigma) \leq 1$. The remaining steps consist of shuffling terms and applying definitions. Plugging this back into Equation (II.8), we get

$$
\begin{aligned}
&\|\mathcal{N}(X)\|_1 \\
&\leq \sup_{m,n} 2E_\gamma(\rho\|\sigma)E_\gamma(\mathcal{N}(|m\rangle\langle m|)\|\mathcal{N}(|n\rangle\langle n|)) - (1 - \gamma)
\end{aligned}
$$

and again plugging this into Equation (II.6), we have

$$
\begin{aligned}
&E_\gamma(\mathcal{N}(\rho)\|\mathcal{N}(\sigma)) \\
&\leq \frac{1}{2}\left\{\sup_{m,n} 2E_\gamma(\rho\|\sigma)E_\gamma(\mathcal{N}(|m\rangle\langle m|)\|\mathcal{N}(|n\rangle\langle n|))\right. \\
&\quad \left. - (1 - \gamma)\right\} + \frac{1}{2}(1 - \gamma) \\
&= E_\gamma(\rho\|\sigma) \sup_{m,n} E_\gamma(\mathcal{N}(|m\rangle\langle m|)\|\mathcal{N}(|n\rangle\langle n|)).
\end{aligned}
$$

From here it follows directly that

$$
\eta_\gamma(\mathcal{N}) \leq \sup_{|\varphi\rangle \perp |\psi\rangle} E_\gamma(\mathcal{N}(|\varphi\rangle\langle\varphi|)\|\mathcal{N}(|\psi\rangle\langle\psi|)).
$$

It remains to show that the inequality is indeed achieved. For that simply note that for all orthogonal $|\varphi\rangle$, $|\psi\rangle$, we have

$$
E_\gamma(|\varphi\rangle\langle\varphi|\||\psi\rangle\langle\psi|) = 1
$$

and therefore

$$
\begin{aligned}
\eta_\gamma(\mathcal{N}) &\geq \sup_{|\varphi\rangle \perp |\psi\rangle} \frac{E_\gamma(\mathcal{N}(|\varphi\rangle\langle\varphi|)\|\mathcal{N}(|\psi\rangle\langle\psi|))}{E_\gamma(|\varphi\rangle\langle\varphi|\||\psi\rangle\langle\psi|)} \\
&= \sup_{|\varphi\rangle \perp |\psi\rangle} E_\gamma(\mathcal{N}(|\varphi\rangle\langle\varphi|)\|\mathcal{N}(|\psi\rangle\langle\psi|)),
\end{aligned}
$$

where the lower bound follows from the fact that we are restricting the supremum to a smaller set than in the definition of the contraction coefficient. Furthermore, for orthogonal pure states one can check that $E_\gamma(|\varphi\rangle\langle\varphi|\||\psi\rangle\langle\psi|) = 1$. This concludes the proof. □

Next we prove a Fuchs-van-de-Graaf type inequality that reduces to the well-known

$$
\frac{1}{2}\|\rho - \sigma\|_1 \leq \sqrt{1 - F(\rho, \sigma)}
$$

for $\gamma = 1$, where $F(\rho, \sigma) := \|\sqrt{\rho}\sqrt{\sigma}\|_1^2$ is the quantum fidelity.

Lemma II.3. *For $\gamma \geq 1$ and $\rho, \sigma \in \mathcal{P}(\mathcal{H})$, we have*

$$
\begin{aligned}
&E_\gamma(\rho\|\sigma) \qquad\qquad \text{(II.9)} \\
&\leq \frac{1}{2}\sqrt{(\operatorname{Tr}(\rho + \gamma\sigma))^2 - 4\gamma(\operatorname{Tr}\rho)^2(\operatorname{Tr}\sigma)^2 F(\hat{\rho}, \hat{\sigma})} \\
&\quad + \frac{\operatorname{Tr}(\rho - \gamma\sigma)}{2}, \qquad\qquad \text{(II.10)}
\end{aligned}
$$

where $\hat{\rho} = \frac{\rho}{\operatorname{Tr}\rho}$ and $\hat{\sigma} = \frac{\sigma}{\operatorname{Tr}\sigma}$. For $\rho, \sigma \in \mathcal{S}_=(\mathcal{H})$ this simplifies to

$$
E_\gamma(\rho\|\sigma) \leq \frac{1}{2}\sqrt{(1 + \gamma)^2 - 4\gamma F(\rho, \sigma)} + \frac{(1 - \gamma)}{2}. \qquad \text{(II.11)}
$$

Proof. Using [9, Supplementary Lemma 3] as stated in Lemma A.1 we get

$$
\|\rho - \gamma\sigma\|_1^2 + 4\|\sqrt{\rho}\sqrt{\gamma\sigma}\|_1^2 \leq (\operatorname{Tr}[\rho + \gamma\sigma])^2,
$$

which immediately implies

$$
\begin{aligned}
&\|\rho - \gamma\sigma\|_1 \\
&\leq \sqrt{(\operatorname{Tr}(\rho + \gamma\sigma))^2 - 4\gamma(\operatorname{Tr}\rho)^2(\operatorname{Tr}\sigma)^2 F(\hat{\rho}, \hat{\sigma})}.
\end{aligned}
$$

Plugging this into Equation (II.2) gives the desired result. □

This also gives us a bound on the contraction coefficient as

$$
\begin{aligned}
&\eta_\gamma(\mathcal{N}) \qquad\qquad \text{(II.12)} \\
&\leq \sup_{|\varphi\rangle \perp |\psi\rangle} \frac{1}{2}\sqrt{(1 + \gamma)^2 - 4\gamma F(\mathcal{N}(|\varphi\rangle\langle\varphi|), \mathcal{N}(|\psi\rangle\langle\psi|))} \\
&\quad + \frac{(1 - \gamma)}{2}. \qquad\qquad \text{(II.13)}
\end{aligned}
$$

We can also bound the hockey-stick divergence for any $\gamma \geq 1$ directly by the trace-distance, leading to an alternative way of bounding the contraction coefficients. This generalizes the analog classical result in [5].

Lemma II.4. *For $\gamma \geq 1$ and $\rho, \sigma \in \mathcal{S}_=(\mathcal{H})$, we have*

$$
1 - \gamma(1 - \frac{1}{2}\|\rho - \sigma\|_1) \leq E_\gamma(\rho\|\sigma) \leq \frac{1}{2}\|\rho - \sigma\|_1, \qquad \text{(II.14)}
$$

implying

$$
1 - \gamma(1 - \eta_1(\mathcal{N})) \leq \eta_\gamma(\mathcal{N}) \leq \eta_1(\mathcal{N}), \qquad \text{(II.15)}
$$

Proof. The only thing needed to prove is the first inequality and the remaining statement follows easily. To that end

we observe

$$\begin{aligned}
&\gamma\frac{1}{2}\|\rho-\sigma\|_1\\
&=\gamma\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\rho-\sigma)\\
&=\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\gamma\rho-\gamma\sigma+\rho-\rho)\\
&\le\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\rho-\gamma\sigma)+\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\gamma\rho-\rho)\\
&=E_\gamma(\rho\|\sigma)+\gamma-1.
\end{aligned}$$

The second inequality of the statement follows by definition and the third and fourth by using the simplified expression for the contraction coefficients in Theorem II.2. □

Finally, we list some potentially useful properties of E_γ. Some of these are generalizations of classical results proven in [24].

Proposition II.5. *We have the following properties:*

- *(Triangle inequality) For $\gamma_1,\gamma_2\ge 1$ and $\rho,\sigma\in\mathcal{P}(\mathcal{H})$, we have*

$$E_{\gamma_1\gamma_2}(\rho\|\sigma)\le E_{\gamma_1}(\rho\|\tau)+\gamma_1 E_{\gamma_2}(\tau\|\sigma). \tag{II.16}$$

- *(Strong convexity) Let $\gamma_1,\gamma_2\ge 1$, $\rho=\sum_x p(x)\rho_x$ and $\sigma=\sum_x q(x)\sigma_x$ with $\rho_x,\sigma_x\in\mathcal{P}(\mathcal{H})$, we have*

$$E_{\gamma_1\gamma_2}(\rho\|\sigma)\le\sum_x p(x)E_{\gamma_1}(\rho_x\|\sigma_x)+\gamma_1 E_{\gamma_2}(\tilde{p}\|\tilde{q}), \tag{II.17}$$

  where $\tilde{p}$ and $\tilde{q}$ are non-normalized distributions $\tilde{p}(x)=p(x)\operatorname{Tr}\sigma_x$ and $\tilde{q}(x)=q(x)\operatorname{Tr}\sigma_x$, respectively. This also implies convexity and joint convexity.
- *(Stability) For $\gamma\ge 1$ and $\rho,\sigma,\tau\in\mathcal{P}(\mathcal{H})$, $\tau\neq 0$, we have*

$$E_\gamma(\rho\otimes\tau\|\sigma\otimes\tau)=\operatorname{Tr}[\tau]\,E_\gamma(\rho\|\sigma). \tag{II.18}$$

- *(Subadditivity) For $\gamma_1,\gamma_2\ge 1$ and $\rho_1,\sigma_1\in\mathcal{P}(\mathcal{H}_1)$, $\rho_2,\sigma_2\in\mathcal{P}(\mathcal{H}_2)$, we have*

$$\begin{aligned}
&E_{\gamma_1\gamma_2}(\rho_1\otimes\rho_2\|\sigma_1\otimes\sigma_2)\\
&\le\operatorname{Tr}[\rho_2]E_{\gamma_1}(\rho_1\|\sigma_1)+\operatorname{Tr}[\sigma_1]\gamma_1 E_{\gamma_2}(\rho_2\|\sigma_2),
\end{aligned} \tag{II.19}$$

$$\begin{aligned}
&E_{\gamma_1\gamma_2}(\rho_1\otimes\rho_2\|\sigma_1\otimes\sigma_2)\\
&\le\operatorname{Tr}[\rho_1]E_{\gamma_1}(\rho_2\|\sigma_2)+\operatorname{Tr}[\sigma_2]\gamma_1 E_{\gamma_2}(\rho_1\|\sigma_1).
\end{aligned} \tag{II.20}$$

- *(Symmetry) For $\gamma\ge 1$ and $\rho,\sigma\in\mathcal{P}(\mathcal{H})$, we have*

$$E_\gamma(\rho\|\sigma)=\gamma E_{\frac{1}{\gamma}}(\sigma\|\rho)+(\operatorname{Tr}(\rho)-\gamma\operatorname{Tr}(\sigma)). \tag{II.21}$$

- *(Trace bound) For $\gamma\ge 1$ and $\rho,\sigma,\tau\in\mathcal{P}(\mathcal{H})$, we have*

$$\begin{aligned}
&E_\gamma(\rho\|\sigma)+E_\gamma(\rho\|\tau) &\text{(II.22)}\\
&\ge\frac{\gamma}{2}\|\tau-\sigma\|_1+(\operatorname{Tr}(\rho)-\gamma\operatorname{Tr}(\tau)). &\text{(II.23)}
\end{aligned}$$

Proof. The first statement is easily seen as follows,

$$\begin{aligned}
&E_{\gamma_1\gamma_2}(\rho\|\sigma)\\
&=\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\rho-\gamma_1\gamma_2\sigma)\\
&=\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\rho-\gamma_1\tau+\gamma_1\tau-\gamma_1\gamma_2\sigma)\\
&\le\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\rho-\gamma_1\tau)+\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\gamma_1\tau-\gamma_1\gamma_2\sigma)\\
&\le E_{\gamma_1}(\rho\|\tau)+\gamma_1 E_{\gamma_2}(\tau\|\sigma).
\end{aligned}$$

The strong convexity follows similarly by considering $\tau=\sum_x p(x)\sigma_x$. Stability follows by first considering the statement for the state $\tau'=\tau/\operatorname{Tr}[\tau]$. Applying data-processing twice, once for the partial trace and once for $\mathcal{N}(\rho)=\rho\otimes\tau'$ gives the statement for states. The generalization then follows by noting that the underlying quantity is absolutely homogeneous in τ. Subadditivity follows by triangle inequality and stability, picking first $\tau=\sigma_1\otimes\rho_2$ and then $\tau=\rho_1\otimes\sigma_2$. For the last two items we need the observation that

$$\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\rho-\gamma\sigma)=\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}(\mathbb{1}-\Lambda)(\rho-\gamma\sigma).$$

From this symmetry follows immediately and the trace bound by observing

$$\begin{aligned}
&E_\gamma(\rho\|\sigma)+E_\gamma(\rho\|\tau)\\
&=\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\rho-\gamma\sigma)+\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}(\mathbb{1}-\Lambda)(\rho-\gamma\tau)\\
&\ge\max_{0\le\Lambda\le\mathbb{1}}\operatorname{Tr}\Lambda(\rho-\gamma\sigma-\rho+\gamma\tau)+(\operatorname{Tr}(\rho)-\gamma\operatorname{Tr}(\tau))\\
&=\frac{\gamma}{2}\|\tau-\sigma\|_1+(\operatorname{Tr}(\rho)-\gamma\operatorname{Tr}(\tau)).
\end{aligned}$$

□

Next we connect the hockey-stick divergence to the smooth max-relative entropy. This is similar to [40, Lemma 1], however based on different definitions. Also, besides avoiding explicit use of measurements, we provide a considerably simpler proof. For a classical analogue of both results, see [17], [24].

Lemma II.6. *For $\gamma\ge 1$ we have,*

$$D^\epsilon_{\max}(\rho\|\sigma)\le\log\gamma\quad\Leftrightarrow\quad E_\gamma(\rho\|\sigma)\le\epsilon, \tag{II.24}$$

where $D^\epsilon_{\max}(\rho\|\sigma)=\inf_{\bar{\rho}\in B^\epsilon(\rho)}D_{\max}(\bar{\rho}\|\sigma)$ is the smooth max*-relative entropy with $D_{\max}(\rho\|\sigma)=\inf\{\lambda:\rho\le e^\lambda\sigma\}$ and $B^\epsilon(\rho)=\{\bar{\rho}:\bar{\rho}\in\mathcal{P}(\mathcal{H})\wedge E_1(\rho,\bar{\rho})\le\epsilon\}$.*

Proof. First, we show the $\Rightarrow$ direction. Assume that $D^\epsilon_{\max}(\rho\|\sigma)\le\log\gamma$, this implies that there exists a $\bar{\rho}$

such that $E_1(\rho,\bar{\rho}) \leq \epsilon$ and $\bar{\rho} \leq \gamma\sigma$. From this and the triangle inequality in Proposition II.5 we immediately get

$$\begin{aligned} E_\gamma(\rho\|\sigma) &\leq E_1(\rho\|\bar{\rho}) + E_\gamma(\bar{\rho}\|\sigma) \\ &\leq \epsilon + 0. \end{aligned}$$

We now show the $\Leftarrow$ direction. Fix $\bar{\rho} = \gamma\sigma$. As $\gamma \geq 1$ and σ is a positive operator, this is a positive operator. Furthermore, we have that

$$\begin{aligned} E_1(\rho,\bar{\rho}) &= \operatorname{Tr}(\rho - \bar{\rho})^+ \\ &= \operatorname{Tr}(\rho - \gamma\sigma)^+ \\ &= E_\gamma(\rho,\sigma) \leq \epsilon\,, \end{aligned}$$

where in the last inequality we used our hypothesis. Thus, $\bar{\rho} \in B^\epsilon(\rho)$. Now observe that $D_{\max}(\bar{\rho}\|\sigma) = \log\gamma$ by construction, from which it follows that $D^\epsilon_{\max}(\rho\|\sigma) \leq \log\gamma$. We therefore see that $E_\gamma(\rho,\sigma) \leq \epsilon$ implies $D^\epsilon_{\max}(\rho\|\sigma) \leq \log\gamma$. □

Note that $D^\epsilon_{\max}$ does not correspond to the usual definition of the smooth max-relative entropy, as we do not constrain the optimization to normalized or subnormalized operators and use E_1 as our distance measure of choice. It does however generalize the definition used in the proof of the analog classical lemma, see [24, Definition 8]. Also, [40, Lemma 1] proves a similar statement optimizing over normalized operators that are however not necessarily positive. It remains open for now whether both conditions can be achieved simultaneously in the above proof.

In the next section, we will proceed to apply the above results to quantum differential privacy.

III. Quantum differential privacy

There are several similar definitions of quantum differential privacy in the literature that apply to different settings. Following [40], we will formulate ours in a way that it applies to an arbitrary quantum algorithm $\mathcal{A}$, i.e. a completely positive, trace-preserving map. The general idea is that if we apply the algorithm to a state from a fixed database, say $\mathcal{D}$, then a malicious party gaining access to the output should not be able to distinguish by any measurement whether the used input was a certain state or one of its immediate neighbors in the database. Classically the motivation is usually to consider a set of databases that are neighbors if they differ in only one entry, e.g. an observer of a medical trial should not be able to determine the results of an individual participant.

In the quantum literature, several definitions of the neighboring status are used, e.g. closeness in trace distance or reachability by a single local operation. Leaving the exact choice of definition open for now, we denote two states being neighbors by $\sigma \sim \rho$. We can now state the definition of quantum differential privacy.

Definition III.1. *Let $\mathcal{D}$ be a set of quantum states and $\mathcal{A}$ be a quantum algorithm (i.e. a* CPTP *map). We call $\mathcal{A}$ (ϵ,δ)-differentially private if for all measurements $0 \leq M \leq \mathbb{1}$ and all $\rho,\sigma \in \mathcal{D}$ such that $\rho \sim \sigma$, we have*

$$\operatorname{Tr}(M\mathcal{A}(\rho)) \leq e^\epsilon \operatorname{Tr}(M\mathcal{A}(\sigma)) + \delta. \tag{III.1}$$

We simply call $\mathcal{A}$ ϵ-differentially private if $\mathcal{A}$ is $(\epsilon,0)$-differentially private.

This definition is a rather direct generalization of classical differential privacy. Indeed, if we only consider diagonal states and projectors P in the computational basis in the definition above and interpret $\operatorname{Tr}(P\mathcal{A}(\sigma))$ as the probability of the measurement outcome lying in a given set, we obtain exactly the definition in [15, Definition 2.4]. But for quantum states we need to optimize over all possible basis and can even consider POVMs. However, we will now use the tools from the previous section to see that it can indeed be checked as a property of the quantum states themselves without explicitly considering the measurements.

Lemma III.2. *The following three statements are equivalent,*

$$\mathcal{A} \text{ is } (\epsilon,\delta)\text{-differentially private} \tag{III.2}$$

$$\Leftrightarrow \quad \sup_{\rho\sim\sigma} E_{e^\epsilon}(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) \leq \delta \tag{III.3}$$

$$\Leftrightarrow \quad \sup_{\rho\sim\sigma} D^\delta_{\max}(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) \leq \epsilon. \tag{III.4}$$

Proof. Note that we can rewrite the condition in Equation (III.1) as

$$\operatorname{Tr}\left(M(\mathcal{A}(\rho) - e^\epsilon\mathcal{A}(\sigma))\right) \leq \delta.$$

Since this has to hold for all measurements and all neighboring input states we can use Lemma II.1 to conclude the first equivalence. The second then follows directly from Lemma II.6. □

Note that Lemma III.2 implies that if all the outputs of an algorithm are diagonal in the same basis, then quantum DP is equivalent to classical DP. This is because for two states that commute, E_γ is the same for the states and the corresponding probability distributions. Note that in the case of $\epsilon = 0$ we have a condition on the trace distance. This is also known as local sensitivity as e.g. defined in [17, Definition 7.1], which in turn is closely related to the stability of an algorithm, see [17, Section 7.3]. For recent quantum generalization of the latter concept see also [4].

The above allows us to immediately conclude some well known properties, however with remarkably simple

proofs solely based on properties of the divergences.

Corollary III.3. *The following properties hold.*

- *(Post-processing) Let $\mathcal{A}$ be (ϵ, δ)-differentially private and $\mathcal{N}$ be an arbitrary quantum channel, then $\mathcal{N} \circ \mathcal{A}$ is also (ϵ, δ)-differentially private.*
- *(Parallel composition) Let $\mathcal{A}_1$ be (ϵ_1, δ_1)-differentially private and $\mathcal{A}_2$ be (ϵ_2, δ_2)-differentially private. Define that $\rho_1 \otimes \rho_2 \sim \sigma_1 \otimes \sigma_2$ if $\rho_1 \sim \sigma_1$ and $\rho_2 \sim \sigma_2$. Then $\mathcal{A}_1 \otimes \mathcal{A}_2$ is $(\epsilon_1 + \epsilon_2, \bar{\delta})$-differentially private on such product states, with $\bar{\delta} = \min\{\delta_1 + e^{\epsilon_1}\delta_2, e^{\epsilon_2}\delta_1 + \delta_2\}$.*

Proof. The post-processing property was first shown in [40, Proposition 1] and composition in [40, Theorem 4] for $\delta_1 = \delta_2 = 0$ and for the general case in [40, Theorem 4], although the latter relied on a erroneous assumption in [40, Lemma 1], see below. Now, post-processing simply follows by data-processing of the hockey-stick divergence. Composition for $\delta_1 = \delta_2 = 0$ is of course implied by the general case, however also follows easily by subadditivity of the hockey-stick divergence which we showed in Lemma II.5. The general case can be seen from the differential privacy formulation in Equation (III.4). Let $\bar{\rho}_1$ and $\bar{\rho}_2$ be the optimizers in the smooth max-relative entropies implied by the differential privacy assumption. One easily sees that

$$\bar{\rho}_1 \otimes \bar{\rho}_2 \leq e^{\epsilon_1+\epsilon_2} \mathcal{A}_1(\sigma_1) \otimes \mathcal{A}_2(\sigma_2) \tag{III.5}$$

and

$$\begin{aligned} &E_1(\mathcal{A}_1(\rho_1) \otimes \mathcal{A}_2(\rho_2) \| \bar{\rho}_1 \otimes \bar{\rho}_2) \\ &\leq \mathrm{Tr}[\bar{\rho}_2] E_1(\mathcal{A}_1(\rho_1) \| \bar{\rho}_1) + \mathrm{Tr}[\mathcal{A}_1(\rho_1)] E_1(\mathcal{A}_2(\rho_2) \| \bar{\rho}_2) \\ &\leq e^{\varepsilon_2}\delta_1 + \delta_2, \end{aligned}$$

which follows from the subadditivity property in Proposition II.5 and similarly,

$$E_1(\mathcal{A}_1(\rho_1) \otimes \mathcal{A}_2(\rho_2) \| \bar{\rho}_1 \otimes \bar{\rho}_2) \leq \delta_1 + e^{\varepsilon_1}\delta_2.$$

This makes $\bar{\rho} = \bar{\rho}_1 \otimes \bar{\rho}_2$ a valid choice to prove the claim. □

We note that the parallel composition property was also claimed in [40, Theorem 5]. However, note that in their analogue of Eq. (II.24) in [40, Lemma 1] their definition of $D^{\epsilon}_{\max}$ requires an optimization over states which are not necessarily positive. As Eq. (III.5) does not hold for operators that are not positive (i.e. $A_1 \leq B_1, A_2 \leq B_2 \not\Rightarrow A_1 \otimes A_2 \leq B_1 \otimes B_2$ in general), the parallel decomposition does not follow. Thus, to the best of our knowledge, Corollary III.3 is the first to establish parallel composition property for quantum differential privacy.

The implementation of many near-term quantum algorithms on noisy devices can be modelled by layers of intended channels $\mathcal{C}_i$, e.g. gates in a circuit or layers of a quantum neural network, directly followed by intermediate noise $\mathcal{N}_i$, i.e.

$$\mathcal{A} = \bigcirc_i^n \mathcal{N}_i \circ \mathcal{C}_i. \tag{III.6}$$

These types of algorithms are predestined to applying an approach based on contraction coefficients. Doing so, we get the following result.

Proposition III.4. *For an algorithm $\mathcal{A}$ of the form in Equation* (III.6) *we have*

$$E_{e^\epsilon}(\mathcal{A}(\rho) \| \mathcal{A}(\sigma)) \leq \left(\prod_i \eta_{e^\epsilon}(\mathcal{N}_i) \right) E_{e^\epsilon}(\rho \| \sigma)\,. \tag{III.7}$$

Proof. The proof follows by alternatingly using the definition of the contraction coefficient to shell off the noise and using data processing to remove the computational layers. □

This directly implies a decay in the δ parameter for differential privacy based on the contraction coefficient of the noise channels. This intuition becomes even more transparent when we consider a special case.

Corollary III.5. *For an algorithm $\mathcal{A}$ of the form in Equation* (III.6) *with all $\mathcal{N}_i = \mathcal{N}$ identical and $\rho \sim \sigma$ if $\frac{1}{2}\|\rho - \sigma\|_1 \leq \kappa$, then $\mathcal{A}$ is (ϵ, δ)-differentially private with*

$$\delta = \left(\eta_{e^\epsilon}(\mathcal{N})\right)^n \kappa. \tag{III.8}$$

Proof. The corollary is a direct application of Proposition III.4. □

This implies in particular an exponential decay of δ with the length of the algorithm, i.e. every such algorithm with $\eta_{e^\epsilon}(\mathcal{N}) < 1$ will eventually become differentially private with vanishing δ for large enough n. Note that generally $\eta_{e^\epsilon}(\mathcal{N})$ is a function of the channel but also ϵ allowing for a certain trade-off between ϵ and δ. Finally, we remark that thanks to Theorem II.2 in the previous section, $\eta_{e^\epsilon}(\mathcal{N})$ is more easily computable and can be controlled analytically or numerically for many noise models. Before moving to more specific models, we give a simple bound on the measured relative entropy, defined as

$$D_M(\rho\|\sigma) = \sup_{\{M_x\}} D(\mathrm{Tr}(M_x\rho) \| \mathrm{Tr}(M_x\sigma)), \tag{III.9}$$

of a differentially private channel. The relative entropy is a common tool when comparing quantum states and therefore the result might be of independent interest.

Lemma III.6. *Let $\mathcal{A}$ be ϵ-differentially private. Then for all $\rho \sim \sigma$,*

$$D_M(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) \leq 2\epsilon E_1(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) \leq 2\epsilon(1-e^{-\epsilon}). \tag{III.10}$$

Proof. Let $\{M_x\}$ be the POVM that achieves the maximum in $D_M(\mathcal{A}(\rho)\|\mathcal{A}(\sigma))$ and let $p_x = \operatorname{Tr} M_x\mathcal{A}(\rho)$ and $q_x = \operatorname{Tr} M_x\mathcal{A}(\sigma)$. Then,

$$\begin{aligned}
D_M(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) &= \sum_x p_x \log\frac{p_x}{q_x} \\
&\leq \sum_x p_x \log\frac{p_x}{q_x} + \sum_x q_x \log\frac{q_x}{p_x} \\
&= \sum_x (p_x - q_x)\log\frac{p_x}{q_x} \\
&\leq \sum_x (p_x - q_x)\epsilon \\
&= \epsilon \sum_x \operatorname{Tr} M_x(\mathcal{A}(\rho) - \mathcal{A}(\sigma)) \\
&\leq 2\epsilon E_1(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) \\
&\leq 2\epsilon\left(1-e^{-\epsilon}\right),
\end{aligned}$$

where the first equality is by definition and the second and third obvious. The first inequality follows because the relative entropy is non-negative, the second by ϵ-DP, the third by a property of the trace distance and the final one by Lemma II.4 and again ϵ-DP. □

This lemma gives a quantum version of several similar relations for classical differential privacy. We remark that it might be possible to find tighter bounds along the lines of the classical [14, Theorem 1]. We leave investigating non-measured quantities for future work, however remark here that because

$$D(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) \leq D_{\max}(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)),$$

ϵ-differential privacy always also implies small relative entropy for neighbouring input states and refer to Section V-B for relaxations involving Rényi relative entropies.

IV. Applications to specific noise models

In the remainder of this section we will discuss particular examples and implications of the above results, including global and local depolarizing noise as well as arbitrary local qubit noise.

A. Global depolarizing noise

A typical noise channel is the depolarizing channel defined as

$$\mathcal{D}_p(\rho) = (1-p)\rho + p\frac{\mathbb{1}}{D}, \tag{IV.1}$$

with $0 \leq p \leq 1$ and D the dimension of the system. In this case we can easily bound the contraction coefficient.

Lemma IV.1. *For $0 \leq p \leq 1$ and $\gamma \geq 1$ we have*

$$\begin{aligned}
&E_\gamma(\mathcal{D}_p(\rho)\|\mathcal{D}_p(\sigma)) && \text{(IV.2)}\\
&\leq \max\{0, (1-\gamma)\frac{p}{D} + (1-p)E_\gamma(\rho\|\sigma)\} && \text{(IV.3)}
\end{aligned}$$

and

$$\eta_\gamma(\mathcal{D}_p) = \max\{0, (1-\gamma)\frac{p}{D} + (1-p)\}. \tag{IV.4}$$

Proof. Note that

$$\begin{aligned}
&E_\gamma(\mathcal{D}_p(\rho)\|\mathcal{D}_p(\sigma)) \\
&= \operatorname{Tr}((1-\gamma)p\frac{\mathbb{1}}{D} + (1-p)(\rho-\gamma\sigma))^+ \\
&= \operatorname{Tr} P^+((1-\gamma)p\frac{\mathbb{1}}{D} + (1-p)(\rho-\gamma\sigma)),
\end{aligned}$$

where P^+ is the projector onto the positive subspace of $(1-\gamma)p\frac{\mathbb{1}}{D} + (1-p)(\rho-\gamma\sigma)$. Observe that

$$E_\gamma(\mathcal{D}_p(\rho)\|\mathcal{D}_p(\sigma)) > 0 \quad \Rightarrow \quad \operatorname{Tr} P^+ \geq 1.$$

Considering this case we get

$$\begin{aligned}
&E_\gamma(\mathcal{D}_p(\rho)\|\mathcal{D}_p(\sigma)) \\
&= (1-\gamma)\frac{p}{D}\operatorname{Tr} P^+ + (1-p)(\operatorname{Tr} P^+(\rho-\gamma\sigma)) \\
&\leq (1-\gamma)\frac{p}{D} + (1-p)E_\gamma(\rho\|\sigma) \\
&\leq (1-\gamma)\frac{p}{D} + (1-p).
\end{aligned}$$

Note that for sufficiently large γ the upper bound could become negative, but one can easily check that in this case $E_\gamma(\mathcal{D}_p(\rho)\|\mathcal{D}_p(\sigma)) = 0$ implying that we are in the other case. It remains to show that this is optimal but this is easy to see by picking any two orthogonal pure states. □

In Figure 1 we show an example of the contraction coefficient for the depolarizing channel and examples for the exponential decay of δ in Equation (III.8) for different values of ϵ using the contraction coefficient in Equation IV.4. It becomes clear that while iterative algorithms with depolarizing layers lead to strong privacy, using the contraction coefficient bound, there is only little space before the output states also become mostly useless (the case $\epsilon = 0$).

However, if we know more about the structure of the channel, we can give significantly better bounds. We will do this here using Equation (IV.3).

Lemma IV.2. *Say, $\rho \sim \sigma$ if $\frac{1}{2}\|\rho-\sigma\|_1 \leq \kappa$, then $\mathcal{D}_p$ is (ϵ, δ)-differentially private with*

$$\delta = \max\{0, (1-e^\epsilon)\frac{p}{D} + (1-p)\kappa\}. \tag{IV.5}$$

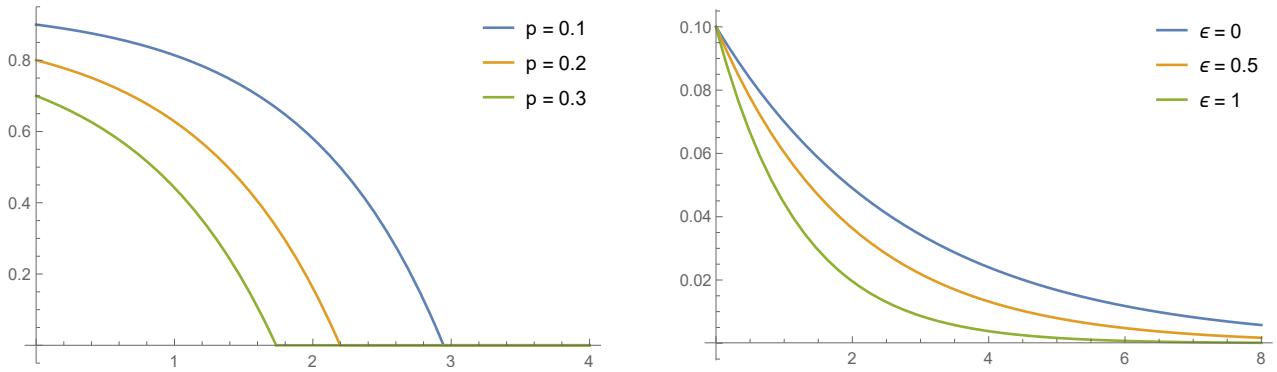

Fig. 1: Left: the contraction coefficient $\eta_{e^\epsilon}(\mathcal{D}_p)$ for $D = 2$ and different values of p plotted against ϵ. Right: the exponential decay given in Equation (III.8) for a depolarizing channel with $p = 0.3$, $\kappa = 0.1$ and different values of ϵ plotted against n, the number of layers with the given depolarizing noise.

For an algorithm $\mathcal{A}$ of the form in Equation (III.6) *with all $\mathcal{N}_i = \mathcal{D}_{p_i}$ and $\rho \sim \sigma$ if $\frac{1}{2}\|\rho - \sigma\|_1 \leq \kappa$, then $\mathcal{A}$ is (ϵ, δ)-differentially private with*

$$\delta = \max\{0, (1 - e^\epsilon)\frac{p_\star}{D} + (1 - p_\star)\kappa\}, \tag{IV.6}$$

where $p_\star = 1 - \prod_i (1 - p_i)$.

Proof. The first statement follows directly from Equation (IV.3). The second statement can be proven by induction. Recall that n is the number of steps in the algorithm, i.e. the number of depolarizing layers. For $n = 1$ the statement follows from (IV.3). Now, assuming the results holds for $n - 1$ layers, for the n-th layer we have

$$\begin{aligned}
&(1-\gamma)\frac{p_\star}{D} + (1 - p_\star)E_\gamma(\mathcal{D}_p(\rho)\|\mathcal{D}_p(\sigma)) \\
&\leq (1-\gamma)\frac{p_\star}{D} \\
&\quad + (1 - p_\star)\left((1-\gamma)\frac{p_n}{D} + (1 - p_n)E_\gamma(\rho\|\sigma)\right) \\
&= (1-\gamma)\frac{1 - (1 - p_\star)(1 - p_n)}{D} \\
&\quad + (1 - p_\star)(1 - p_n)E_\gamma(\rho\|\sigma)
\end{aligned}$$

from which the claim follows directly. □

The above bound is compared to the direct contraction coefficient bound from Equation (III.8) in Figure 2. As can be seen, the improvement is significant. While contraction coefficients give an exponential decay of δ with n, the bound in Equation (IV.6) shows that $\delta = 0$ is sufficient already for $n \geq 3$ if $p = 0.3$ and $n \geq 10$ if $p = 0.1$. We remark that for $\epsilon = 0$ the two bounds give the same result, i.e. in case of the trace distance Lemma IV.2 does not lead to an improvement over the contraction bound. In fact, for the trace distance this simple bound is tight in some cases, because

$$\begin{aligned}
\|\mathcal{D}_p(\rho) - \mathcal{D}_p(\sigma)\|_1 &= \|(1-p)(\rho - \sigma)\|_1 &\text{(IV.7)}\\
&= (1-p)\|\rho - \sigma\|_1\,, &\text{(IV.8)}
\end{aligned}$$

implying that the decrease in trace distance is always exactly $(1 - p)$ for a layer of global depolarizing noise. This easily extends to algorithms $\mathcal{A}$ of the form in Equation (III.6) when all computational layers $\mathcal{C}_i$ are unitary, i.e. don't introduce any noise themselves. These observations lead to a provable separation in n between good privacy and useless algorithm outputs.

Alternatively we can also state a similar result determining a bound on ϵ as a simple corollary.

Corollary IV.3. *Say, $\rho \sim \sigma$ if $\frac{1}{2}\|\rho - \sigma\|_1 \leq \kappa$. For a fixed $\delta \geq 0$, $\mathcal{D}_p$ is (ϵ, δ)-differentially private with*

$$\epsilon \geq \max\{0, \log\left(\frac{D}{p}((1-p)\kappa - \delta) + 1\right)\}. \tag{IV.9}$$

For an algorithm $\mathcal{A}$ of the form in Equation (III.6)*, for a fixed $\delta \geq 0$, $\mathcal{A}$ is (ϵ, δ)-differentially private with*

$$\epsilon \geq \max\{0, \log\left(\frac{D}{p_\star}((1-p_\star)\kappa - \delta) + 1\right)\}. \tag{IV.10}$$

where $p_\star = 1 - \prod_i (1 - p_i)$.

This generalizes some results in the literature, namely, for $\delta = 0$ the first statement reduces to [40, Theorem 3] and the second to [12, Lemma 2]. We present examples of the above bound in Figure 3, showcasing the trade-off between ϵ and δ and the decay of ϵ with growing n.

B. Local depolarizing noise

Previously we have seen how quantum algorithms are affected by global depolarizing noise. However, in a quantum computing device we would rather expect each

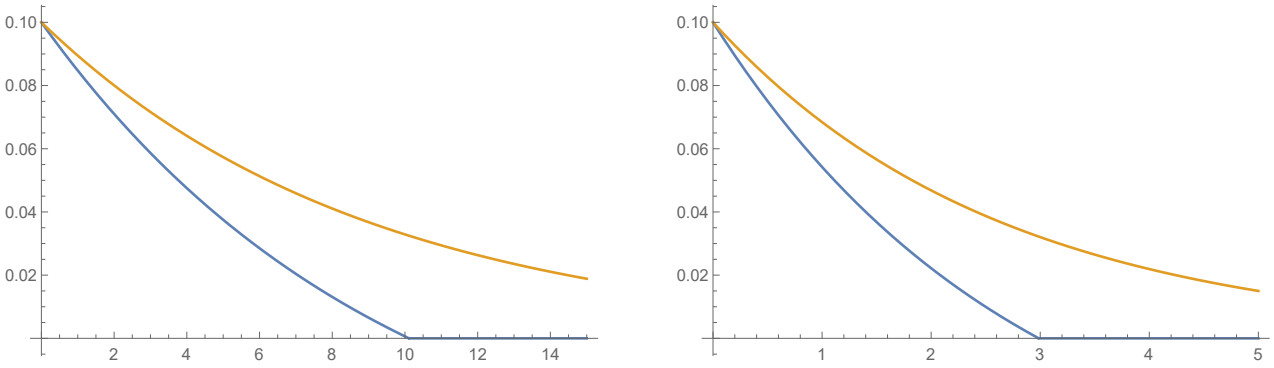

Fig. 2: Comparison of bounds on δ for $\epsilon = 0.1$, $\kappa = 0.1$ and $D = 2$ plotted against n. The top function is the contraction coefficient bound in Equation (III.8) and below the improved bound in Equation (IV.6). Left: $p = 0.1$. Right: $p = 0.3$.

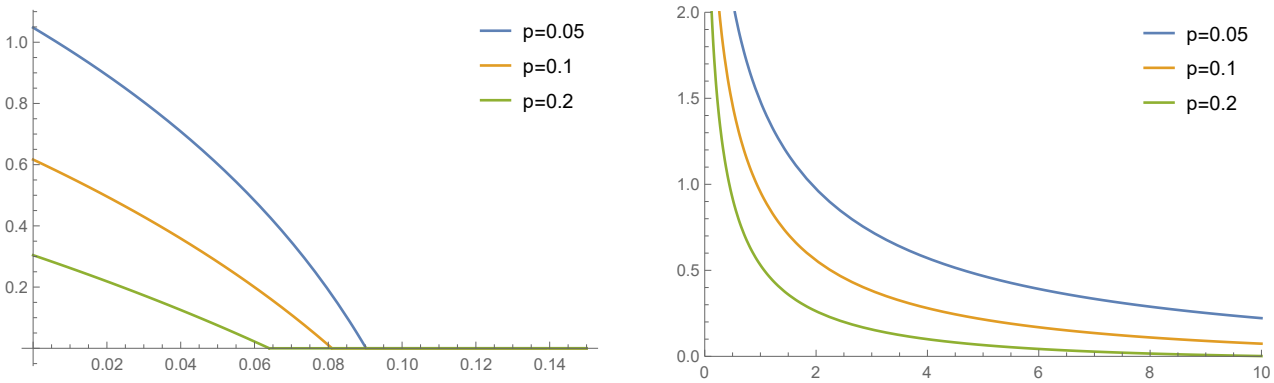

Fig. 3: Bound on ϵ in Equation (IV.10) for $D = 2$ and $\kappa = 0.1$ for different values of $p_i = p$. Left: plotted for $n = 2$ against δ. Right: plotted for $\delta = 0.01$ against n.

qubit to be affected by local noise. In this section we will discus depolarizing noise of the form $\mathcal{D}_p^{\otimes k}$ where k is the number of qubits a quantum algorithm acts on at any given layer. To investigate this setting we first need an equivalent of Lemma IV.1 for local depolarizing noise.

Lemma IV.4. *For $0 \leq p \leq 1$ and $\gamma \geq 1$ we have*

$$E_\gamma(\mathcal{D}_p^{\otimes k}(\rho)\|\mathcal{D}_p^{\otimes k}(\sigma)) \tag{IV.11}$$

$$\leq \max\{0, (1-\gamma)\frac{p^k}{D^k} + (1-p^k)E_\gamma(\rho\|\sigma)\} \tag{IV.12}$$

and

$$\eta_\gamma(\mathcal{D}_p^{\otimes k}) \leq \max\{0, (1-\gamma)\frac{p^k}{D^k} + (1-p^k)\}. \tag{IV.13}$$

Proof. The proof is similar to Lemma IV.1 but requires one more main ingredient. Note that we can always write local depolarizing noise as

$$\mathcal{D}_p^{\otimes k}(\rho) = p^k\frac{\mathbb{1}^{\otimes k}}{D^k} + (1-p^k)\mathcal{M}(\rho),$$

where $\mathcal{M}$ is some CPTP map. This can be checked by direct calculation. With this we have

$$\begin{aligned} &E_\gamma(\mathcal{D}_p(\rho)\|\mathcal{D}_p(\sigma)) \\ &= \mathrm{Tr}((1-\gamma)p^k\frac{\mathbb{1}^{\otimes k}}{D^k} + (1-p^k)\mathcal{M}((\rho-\gamma\sigma)))^+ \\ &= \mathrm{Tr}\, P^+((1-\gamma)p^k\frac{\mathbb{1}^{\otimes k}}{D^k} + (1-p^k)\mathcal{M}((\rho-\gamma\sigma))), \end{aligned}$$

where P^+ is the projector onto the positive subspace of $((1-\gamma)p^k\frac{\mathbb{1}^{\otimes k}}{D^k} + (1-p^k)\mathcal{M}((\rho-\gamma\sigma)))$. Observe that

$$E_\gamma(\mathcal{D}_p^{\otimes k}(\rho)\|\mathcal{D}_p^{\otimes k}(\sigma)) > 0 \quad \Rightarrow \quad \mathrm{Tr}\, P^+ \geq 1.$$

Considering this case we get

$$\begin{aligned}
&E_\gamma(\mathcal{D}_p(\rho)\|\mathcal{D}_p(\sigma)) \\
&= (1-\gamma)\frac{p^k}{D^k}\operatorname{Tr}P^+ + (1-p^k)(\operatorname{Tr}P^+(\mathcal{M}(\rho-\gamma\sigma))) \\
&\leq (1-\gamma)\frac{p^k}{D^k} + (1-p^k)E_\gamma(\mathcal{M}(\rho)\|\mathcal{M}(\sigma)) \\
&\leq (1-\gamma)\frac{p^k}{D^k} + (1-p^k)E_\gamma(\rho\|\sigma) \\
&\leq (1-\gamma)\frac{p^k}{D^k} + (1-p^k).
\end{aligned}$$

Note that for sufficiently large γ the upper bound could become negative, but one can easily check that in this case $E_\gamma(\mathcal{D}_p^{\otimes k}(\rho)\|\mathcal{D}_p^{\otimes k}(\sigma)) = 0$ implying that we are in the other case. □

Note that in this case we provide only an upper bound on the contraction coefficient and determining its exact value remains open. Nevertheless, with the above tool at hand, we can easily generalize Lemma IV.2.

Lemma IV.5. *Say, $\rho \sim \sigma$ if $\frac{1}{2}\|\rho-\sigma\|_1 \leq \kappa$, then $\mathcal{D}_p^{\otimes k}$ is (ϵ,δ)-differentially private with*

$$\delta = \max\{0, (1-e^\epsilon)\frac{p^k}{D^k} + (1-p^k)\kappa\}. \tag{IV.14}$$

For an algorithm $\mathcal{A}$ of the form in Equation (III.6) *with all $\mathcal{N}_i = \mathcal{D}_p^{\otimes k}$ and $\rho \sim \sigma$ if $\frac{1}{2}\|\rho-\sigma\|_1 \leq \kappa$, then $\mathcal{A}$ is (ϵ,δ)-differentially private with*

$$\delta = \max\{0, (1-e^\epsilon)\frac{p_\star}{D^k} + (1-p_\star)\kappa\}, \tag{IV.15}$$

where $p_\star = 1-(1-p^k)^n$.

Proof. The proof is identical to that of Lemma IV.2. □

The bounds on δ are illustrated in Figure 4. On the left we see that for values $\epsilon > 0$, δ is eventually going to reach 0, while for $\epsilon = 0$ i.e. the case of the trace distance, δ decays exponentially but always stays strictly positive. This is the same as for global depolarizing noise. On the right, we compare local and global depolarizing noise. For that we fix a total dimension of 8 at each layer and consider different dimension of the depolarizing noise, from a single global depolarizing channel to 4 local qubit depolarizing channels. From the plot it becomes evident that our bounds guarantee much faster decay of δ for global than for local noise.

Finally, we can also adapt Corollary IV.3 to local noise.

Corollary IV.6. *Say, $\rho \sim \sigma$ if $\frac{1}{2}\|\rho-\sigma\|_1 \leq \kappa$. For a fixed $\delta \geq 0$, $\mathcal{D}_p$ is (ϵ,δ)-differentially private with*

$$\epsilon \geq \max\{0, \log\left(\frac{D^k}{p^k}((1-p^k)\kappa-\delta)+1\right)\}. \tag{IV.16}$$

For an algorithm $\mathcal{A}$ of the form in Equation (III.6) *with all $\mathcal{N}_i = \mathcal{D}_p^k$. For a fixed $\delta \geq 0$, $\mathcal{A}$ is (ϵ,δ)-differentially private with*

$$\epsilon \geq \max\{0, \log\left(\frac{D^k}{p_\star}((1-p_\star)\kappa-\delta)+1\right)\}. \tag{IV.17}$$

where $p_\star = 1-(1-p^k)^n$.

We end this section by providing a lower bound on the trace distance for local depolarizing noise that will demonstrate the substantial difference between decaying computational accuracy and good differential privacy. The bound takes a similar role to the observation around Equation (IV.8) in the previous section.

Proposition IV.7. *For any two states ρ,σ and local depolarizing parameter $0 \leq p \leq 1$ such that $p < \frac{1}{2}$,*

$$\|\rho-\sigma\|_1 \leq \left(\frac{1}{1-2p}\right)^k \|\mathcal{D}_p^{\otimes k}(\rho-\sigma)\|_1 .$$

Therefore, for the noisy circuit $\mathcal{A}$ where all the gates $\mathcal{C}_i$ are chosen to be unitary, and where $\mathcal{N}_i = \mathcal{D}_p^{\otimes k}$, we have that

$$E_1(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) \geq \frac{1}{2}\Big(1-2p\Big)^{kn}\|\rho-\sigma\|_1 . \tag{IV.18}$$

Proof. We start by inverting the depolarizing noise acting on qubit i as:

$$\rho = \frac{\mathcal{D}_p(\rho) - \frac{p}{D}\operatorname{Tr}_i(\rho)\otimes\mathbb{1}_i}{1-p}$$

Hence, we have

$$\begin{aligned}
\|\rho-\sigma\|_1 &= \frac{1}{1-p}\|\mathcal{D}_p(\rho-\sigma) - \frac{p}{D}\operatorname{Tr}_i(\rho-\sigma)\otimes\mathbb{1}_i\|_1 \\
&\leq \frac{1}{1-p}\|\mathcal{D}_p(\rho-\sigma)\|_1 + \frac{p}{(1-p)}\|\rho-\sigma\|_1 .
\end{aligned}$$

Therefore, whenever $p < \frac{1}{2}$, we have that

$$\|\rho-\sigma\|_1 \leq \frac{1}{1-2p}\|\mathcal{D}_p(\rho-\sigma)\|_1 .$$

The result arises from repeating the above step for all the qubits. □

We compare the above lower bound to our previous upper bounds in Figure 5. It can be seen that there is a clear separation between the worst case decay of the trace distance and the depth required to reach differential privacy. It should however be noted that this bound seems to be useful only for small k.

The previous results suggest that the local noise channels do not contract too fast at small enough noise. Here, we prove that the trace distance, and hence also any hockey stick divergence, will converge exponentially fast in the number of layers as soon as a critical local noise

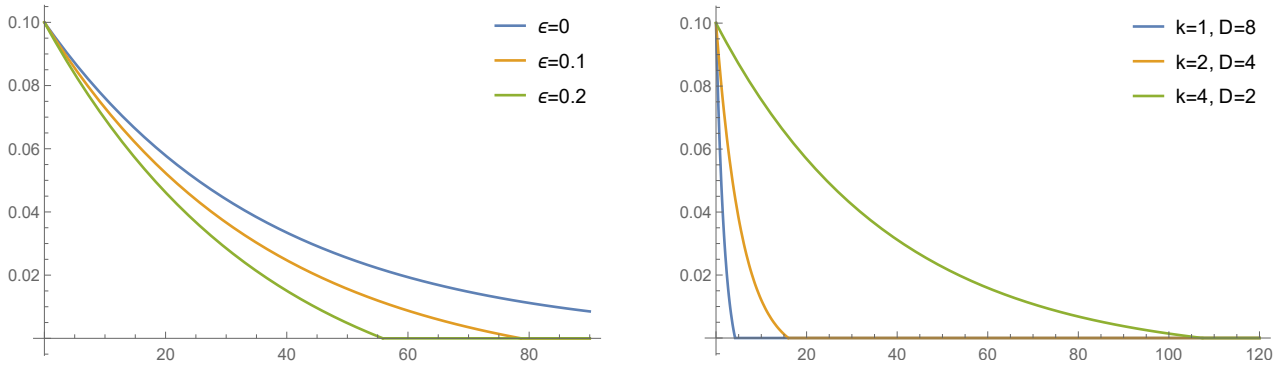

Fig. 4: Left: The bound in Equation (IV.15) on δ for $p = 0.3, k = 3, D = 2, \kappa = 0.1$ plotted against n for different ϵ. Right: The same bound for $p = 0.4, \epsilon = 0.1, \kappa = 0.1$ for different values of k and D, such that the total dimension at each layer is always 8, plotted against n.

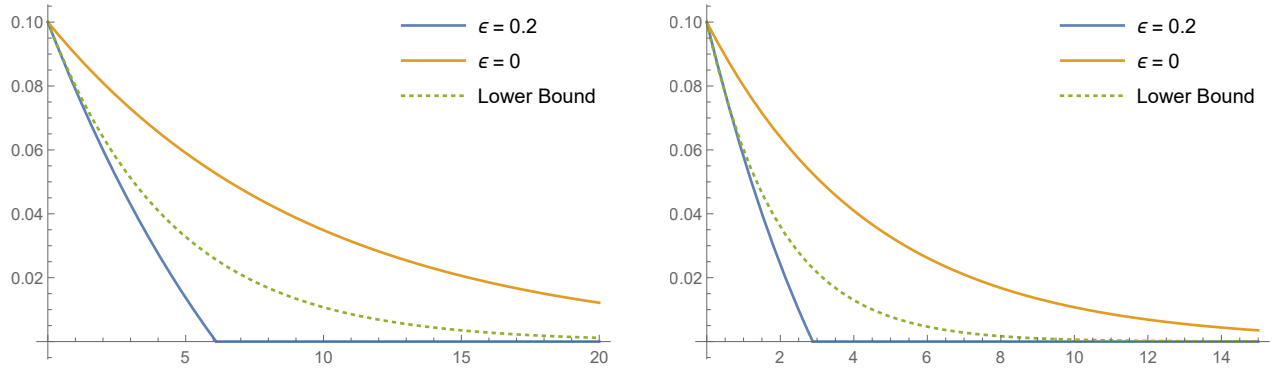

Fig. 5: Comparison of the bound in Equation (IV.15), solid lines for different ϵ, with the lower bound in Equation (IV.18), dotted line. Plotted for $D = 2$, $k = 1$ and $\frac{1}{2}\|\rho - \sigma\|_1 = 0.1$ against n, with $p = 0.1$ on the left and $p = 0.2$ on the right.

is attained. Our result can be compared to earlier upper bounds on the noise threshold for fault-tolerant quantum computing as in [22], [29]. However, as we will see, our bound comes as a simple corollary of basic properties of a recently introduced quantum Wasserstein distance in [11]. Although the argument can be generalized to other Pauli noisy channels (see e.g. [19]), we restrict once again our analysis to the simple depolarizing channel. Here, given a quantum channel Φ acting on k qubits, we define its light-cone as follows: first, for any qubit i, we denote by I_i the minimal subset of qubits such that $\mathrm{Tr}_{I_i}(\Phi(\rho)) = \mathrm{Tr}_{I_i}(\Phi(\sigma))$ for any two k-qubit states ρ and σ such that $\mathrm{Tr}_i(\rho) = \mathrm{Tr}_i(\sigma)$. Then, the light-cone of Φ is defined as

$$|I| := \max_{i\in[k]} |I_i| .$$

Proposition IV.8. *Given the noisy circuit $\mathcal{A}$ in Equation* (III.6) *with n layers, k qubits and local depolarizing noise of parameter $0 \leq p \leq 1$, we assume that each layer of the circuit is a quantum channel of light-cone $|I|$. Then, we have that for any two input states ρ, σ*

$$\|\mathcal{A}(\rho) - \mathcal{A}(\sigma)\|_1 \leq 2k \left(2|I|\,(1-p)\right)^n .$$

In other words, the trace distance between any two output states vanishes in logarithmic depth as soon as p satisfies $2|I|(1-p) < 1$.

Proof. We will use the Wasserstein distance introduced in [11] as follows:

$$W_1(\rho,\sigma) := \sup_{\|O\|_L \leq 1} \mathrm{Tr}(O(\rho - \sigma)) ,$$

where

$$\|H\|_L := 2 \max_{i\in[k]} \min_{H^{(i)}} \|H - H^{(i)}\|_\infty ,$$

where the mininmum is over all self-adjoint operators

$H^{(i)}$ which do not act on qubit i. Next, we have

$$\begin{aligned}\frac{1}{2}\left\|\mathcal{A}(\rho-\sigma)\right\|_1 &\leq W_1(\mathcal{A}(\rho),\mathcal{A}(\sigma))\\ &\leq (2|I|(1-p))^n W_1(\rho,\sigma)\\ &\leq \frac{k}{2}\left(2|I|(1-p)\right)^n\left\|\rho-\sigma\right\|_1\end{aligned}$$

Above, the first and last inequalities follow from [11, Proposition 2], whereas the second inequality comes from an alternating use of [11, Propositions 12 and 13]. The result follows. □

Compared to our earlier results this bounds dependence on k is weaker and we recover the $(1-p)^n$ scaling, previously seen for the global depolarizing noise, whenever p is large enough to overpower the, possibly error-correcting, properties of the computational layers.

C. Arbitrary local qubit noise channels

In this section we will give a bound for the contraction of arbitrary local qubit channels based on recent work in [18]. In particular we will see that any non-unital local qubit noise leads to good differential privacy eventually. The main result is the following lemma that generalizes [18, Lemma 7] by combining their proof with our generalized Fuchs-van-de-Graaf inequality.

Lemma IV.9. *Let $\mathcal{T}=\mathcal{N}^{\otimes k}$ and $\mathcal{N}$ some qubit channel, then*

$$\begin{aligned}&\eta_\gamma(\mathcal{T})\\ &\leq \frac{1}{2}\sqrt{(1+\gamma)^2-4\gamma\left(\frac{\lambda_{\min}(C_{\mathcal{N}^\dagger\circ\mathcal{N}})}{4}\right)^k}+\frac{(1-\gamma)}{2},\end{aligned} \tag{IV.19}$$

where $C_{\mathcal{N}^\dagger\circ\mathcal{N}}$ is the Choi matrix of $\mathcal{N}^\dagger\circ\mathcal{N}$ and $\lambda_{\min}(C_{\mathcal{N}^\dagger\circ\mathcal{N}})$ its smallest eigenvalue. If $\mathcal{N}$ is also non-unital then

$$\lambda_{\min}(C_{\mathcal{N}^\dagger\circ\mathcal{N}})>0. \tag{IV.20}$$

Proof. We begin by applying the generalized Fuchs-van-de-Graaf inequality from Lemma II.3 to the simplified expression of the contraction coefficient in Theorem II.2 as previously stated in Equation (II.12). We then apply a bound on the fidelity shown in [18] that states

$$F(\mathcal{N}^{\otimes k}(\Psi),\mathcal{N}^{\otimes k}(\Phi))\geq\left(\frac{\lambda_{\min}(C_{\mathcal{N}^\dagger\circ\mathcal{N}})}{4}\right)^k,$$

see Appendix A for the details. The final statement about non-unital qubit channels was also argued in [18]. □

This bound applies directly to differential privacy via Corollary III.5. Let $\mathcal{A}$ be of the form in Equation (III.6) with all $\mathcal{N}_i=\mathcal{N}^{\otimes k}$ identical and $\rho\sim\sigma$ if $\frac{1}{2}\|\rho-\sigma\|_1\leq\kappa$. Let $\lambda_{\min}(C_{\mathcal{N}^\dagger\circ\mathcal{N}})=\lambda$, then $\mathcal{A}$ is (ϵ,δ)-differentially private with

$$\delta=\left(\frac{1}{2}\sqrt{(1+e^\epsilon)^2-4e^\epsilon\left(\frac{\lambda}{4}\right)^k}+\frac{(1-e^\epsilon)}{2}\right)^n\kappa. \tag{IV.21}$$

We observe that for growing ϵ we get smaller δ, a natural trade-off between the two security parameters, also implying that δ always decays faster than the trace distance ($\epsilon=0$). Also for any channel with $\lambda>0$, which is the case for every non-unital channel as shown above, good differential privacy will be eventually reached for large enough n. We give some numerical examples in Figure 6.

V. Extensions

A. Local quantum differential privacy

Local differential privacy (LDP) was defined in the classical setting for the scenario in which a database is collected from many clients and each of them demands differential privacy to hold for their individual contribution. In this case the client applies an algorithm $\mathcal{A}$ to mask their contribution to the database and the priority is not to make neighboring states look similar but to hide the general information they are sending. The definition of (ϵ,δ)-LDP therefore coincides with that (ϵ,δ)-DP but with a more general set of possible input states. In the extreme setting one could even consider the set of all possible input states, implying that $\mathcal{A}$ is (ϵ,δ)-LDP if

$$\sup_{\rho,\sigma}E_{e^\epsilon}(\mathcal{A}(\rho)\|\mathcal{A}(\sigma))\leq\delta\,, \tag{V.1}$$

based on Lemma III.2. Clearly this is a much stronger requirement than what is required for (ϵ,δ)-DP. In fact from the properties of E_γ one can see that this condition is restrictive enough to imply a bound on the trace distance contraction coefficient. This generalizes a classical result in [5].

Corollary V.1. *Let $\mathcal{A}$ be (ϵ,δ)-LDP and $\varphi(\epsilon,\delta)=1-e^{-\epsilon}(1-\delta)$. Then*

$$\eta_1(\mathcal{A})\leq\varphi(\epsilon,\delta). \tag{V.2}$$

Proof. Clearly, by definition of LDP and Theorem II.2 we have

$$\eta_{e^\epsilon}(\mathcal{A})\leq\delta.$$

Now the corollary follows because from Lemma II.4 we get

$$\eta_1(\mathcal{A})\leq 1-\frac{1-\eta_\gamma(\mathcal{A})}{\gamma},$$

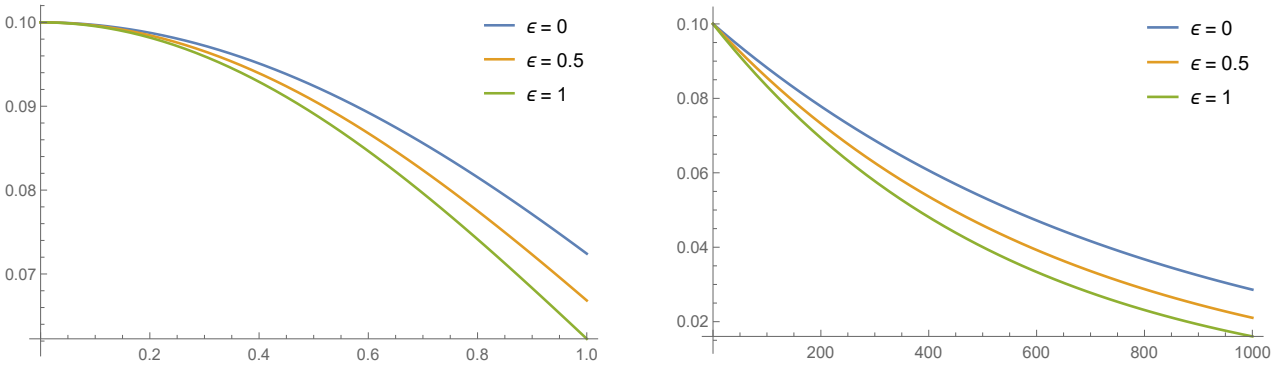

Fig. 6: Guaranteed δ for some local noise channel based on Equation (IV.21) with $\kappa = 0.1$, $k = 2$ and different values of ϵ. Left: plotted for $n = 10$ against λ. Right: plotted for $\lambda = 0.2$ against n.

which together with the definition of (ϵ, δ)-LDP concludes the proof. □

In particular, this implies that requiring $\mathcal{A}$ to be (ϵ, δ)-LDP can strongly limit the usefulness of the output states. In particular applying $\mathcal{A}$ iteratively will lead to very strong privacy guarantees but also make the output states indistinguishable and therefore useless for further computations, compare Proposition III.4.

Note that the above argument also works the same with somewhat less restrictive definitions. In particular, we get the following equivalence, which is similar to an observation in [1].

Corollary V.2. *$\mathcal{A}$ being (ϵ, δ)-LDP with respect to the set of all states is equivalent to $\mathcal{A}$ being (ϵ, δ)-LDP with respect to the set of all pure states.*

Proof. This follows directly from the convexity of E_γ. □

In principle, Corollary V.1 also holds for (ϵ, δ)-LDP with respect to any set of states that includes all orthogonal pure states. For the minimal such set, (ϵ, δ)-LDP would not just imply but indeed be equivalent to $\eta_{e^\epsilon}(\mathcal{A}) \leq \delta$ as observed in the classical setting [5].

B. Rényi quantum differential privacy

In many practical settings ϵ-differential privacy can be too strong of a criteria. On the other hand, (ϵ, δ)-differential privacy allows for rare events that leak a significant amount of information and its composition theorem requires adding the δ's of each algorithm. As an intermediate privacy requirement Mironov proposed Rényi differential privacy [26] and proved that it has several desirable properties. As we have seen in the previous sections, ϵ-differential privacy can be cast in terms of the max-relative entropy. Essentially, (ϵ, α)-Rényi differential privacy is defined by replacing the max-relative entropy by the general α-Rényi relative entropy. In this section we will propose a quantum extension of this concept.

The obvious generalization of the classical definition is to go a similar route to the original quantum differential privacy definition and define measurement outcomes of a POVM $\{M_x\}$ as $p(x) = \operatorname{Tr} M_x \mathcal{A}(\rho)$ and $q(x) = \operatorname{Tr} M_x \mathcal{A}(\sigma)$ and require the classical definition to hold. This can be cast in terms of the measured Rényi relative entropy,

$$\sup_{\rho \sim \sigma} D_{M,\alpha}(\mathcal{A}(\rho) \| \mathcal{A}(\sigma)) = \sup_{\rho \sim \sigma} \sup_{\{M_x\}} D_\alpha(p \| q) \leq \epsilon. \tag{V.3}$$

This coincides with ϵ-differential privacy for $\alpha \to \infty$, however for other values of α there is no known closed (measurement independent) formula for the measured Rényi relative entropy and the definition remains classical at its heart. This also comes with a concrete disadvantage when considering composition bounds, namely that the measured Rényi relative entropy is not generally subadditive[1]. That means there exist examples where

$$\begin{aligned} &D_{M,\alpha}(\mathcal{A}_1 \otimes \mathcal{A}_2(\rho_1 \otimes \rho_2) \| \mathcal{A}_1 \otimes \mathcal{A}_2(\sigma_1 \otimes \sigma_2)) \\ &> D_{M,\alpha}(\mathcal{A}_1(\rho_1) \| \mathcal{A}_1(\sigma_1)) + D_{M,\alpha}(\mathcal{A}_2(\rho_2) \| \mathcal{A}_2(\sigma_2)) \end{aligned} \tag{V.4}$$

implying that, using this quantity, (ϵ, α)-Rényi differential privacy of $\mathcal{A}_1$ and $\mathcal{A}_2$ generally does not imply $(2\epsilon, \alpha)$-Rényi differential privacy for $\mathcal{A}_1 \otimes \mathcal{A}_2$. In this section, we want to propose a fully quantum definition of (ϵ, α)-Rényi differential privacy that avoids this problem.

While in the classical setting there is a uniquely defined

[1]This follows e.g. because the measured Rényi relative entropy is strictly smaller than the sandwiched Rényi relative entropy, see e.g. [8, Theorem 6], however they are equal under regularization, see Equation (V.8)

Rényi relative entropy, in the quantum setting, due to its non-commutative nature, there is an arbitrary number of generalizations. Here, we will not fix a particular definition but simply consider an arbitrary family of Rényi relative entropies $\mathbb{D}_\alpha$ as defined in [34] based on a number of requirements that the quantity needs to fulfill. For completeness a list of the properties can be found in Appendix A and we refer to [34] for details. Here we only note that the list includes several typical properties such as data-processing, additivity and unitary invariance. We further note that the commonly used quantum generalizations such as the Petz-Rényi divergence D_α [28], the sandwiched Rényi divergence $\tilde{D}_\alpha$ [27], [38] and the geometric Rényi divergence $\hat{D}$ [25] are special instances of $\mathbb{D}_\alpha$ for the range of α in which the mentioned properties hold. Formally, we define quantum Rényi differential privacy as follows.

Definition V.3. *We call a quantum channel $\mathcal{A}$ (ϵ, α)-Rényi differentially private if*

$$\sup_{\rho \sim \sigma} \mathbb{D}_\alpha(\mathcal{A}(\rho) \| \mathcal{A}(\sigma)) \leq \epsilon. \tag{V.5}$$

This generally leaves us with a lot of freedom to pick our favourite Rényi relative entropy. Note however that those which include the limit $\alpha \to \infty$ all have the particular feature that in this limit the above definition includes ϵ-differential privacy as a special case. The sandwiched Rényi relative entropy is an example of such a family. We now state the first property of our definition.

Lemma V.4. *If $\mathcal{A}$ is ϵ-differentially private then it is (ϵ, α)-Rényi differentially private.*

Proof. The claim is a direct consequence of

$$\mathbb{D}_\alpha(\mathcal{A}(\rho) \| \mathcal{A}(\sigma)) \leq D_{\max}(\mathcal{A}(\rho) \| \mathcal{A}(\sigma)),$$

which can be seen as follows. We have the following chain of arguments,

$$\mathbb{D}_\alpha(\rho \| \sigma) \leq \hat{D}_\alpha(\rho \| \sigma) \leq \hat{D}_\infty(\rho \| \sigma) = D_{\max}(\rho \| \sigma),$$

where the first inequality is [34, Equation (4.34)], the second is monotonicity of $\hat{D}_\alpha$ in α and the equality is [34, Equation (4.36)]. □

The classical Rényi differential privacy is furthermore known to be stronger than (ϵ, δ)-differential privacy. We now show that the same holds for our quantum definition.

Lemma V.5. *If $\mathcal{A}$ is (ϵ, α)-Rényi differentially private then it is $(\epsilon + \frac{g(\delta)}{\alpha - 1}, \delta)$-differentially private with $g(\epsilon) = -\log(1 - \sqrt{1 - \epsilon^2})$.*

Proof. The main ingredients are [34, Proposition 6.22] as stated in the appendix and an auxiliary lemma relating $D^\delta_{\max}$ with a similar quantity more commonly found in the literature stated in Lemma A.3. From it we have

$$D^\delta_{\max}(\rho \| \sigma) \leq \mathbb{D}_\alpha(\rho \| \sigma) + \frac{g(\delta)}{\alpha - 1} \leq \epsilon + \frac{g(\delta)}{\alpha - 1},$$

where the last inequality is by assumption. The proof follows from Lemma III.2. □

This can be relaxed by noting that $g(\delta) \leq \log \frac{2}{\delta^2}$, which brings it closer to the classical equivalent [26, Proposition 3].

Additionally, we can easily show several desirable properties.

Corollary V.6. *The following properties hold.*

- *(Post-processing) Let $\mathcal{A}$ be (ϵ, α)-Rényi differentially private and $\mathcal{N}$ be an arbitrary quantum channel, then $\mathcal{N} \circ \mathcal{A}$ is also (ϵ, α)-Rényi differentially private.*
- *(Parallel composition) Let $\mathcal{A}_1$ be (ϵ_1, α)-Rényi differentially private and $\mathcal{A}_2$ be (ϵ_2, α)-Rényi differentially private. Define that $\rho_1 \otimes \rho_2 \sim \sigma_1 \otimes \sigma_2$ if $\rho_1 \sim \sigma_1$ and $\rho_2 \sim \sigma_2$. Then $\mathcal{A}_1 \otimes \mathcal{A}_2$ is $(\epsilon_1 + \epsilon_2, \alpha)$-Rényi differentially private on such product states.*

Proof. The first result follows by data-processing and the second by additivity of $\mathbb{D}_\alpha$. □

Finally, we remark that our general quantum definition of Rényi differential privacy implies Rényi differential privacy in the semi-classical definition in Equation (V.3) because

$$D_{M,\alpha}(\mathcal{A}(\rho) \| \mathcal{A}(\sigma)) \leq \mathbb{D}_\alpha(\mathcal{A}(\rho) \| \mathcal{A}(\sigma)) \tag{V.6}$$

which is a simple consequence of data-processing.

We have defined Rényi differential privacy based on the general Rényi relative entropy $\mathbb{D}_\alpha$ and while at this point any choice of a particular Rényi relative entropy would seem justified, we will present some brief arguments that might hint that the sandwiched Rényi relative entropy $\tilde{D}_\alpha$ should be the quantity of choice. First, $\tilde{D}_\alpha$ obeys data-processing for $\alpha \geq \frac{1}{2}$, which makes it a valid choice for this whole range of α, and it is equal to $D_{\max}$ in the limit $\alpha \to \infty$. Furthermore, $\tilde{D}_\alpha$ is the minimal Rényi relative entropy, which means that

$$\mathbb{D}_\alpha(\rho \| \sigma) \geq \tilde{D}_\alpha(\rho \| \sigma)\,, \tag{V.7}$$

see e.g. [34]. This implies that choosing $\tilde{D}_\alpha$ is the least restrictive choice and for a fixed α this Rényi differential privacy would be implied by any other choice. Lastly, while we have seen that the measured Rényi relative entropy has some undesirable properties, the sandwiched

Rényi relative entropy equals the regularized measured Rényi relative entropy [34],

$$\tilde{D}_\alpha(\rho\|\sigma) = \lim_{n\to\infty} \frac{1}{n} D_{M,\alpha}(\rho^{\otimes n}\|\sigma^{\otimes n}), \tag{V.8}$$

giving it a close resemblance to the classical Rényi differential privacy.

C. Hypothesis testing

At its core, differential privacy is a requirement on the probabilities associated to determining a used input based on the output of some information processing. To gain a better intuition of the implications of the imposed restrictions, classical differential privacy can as often be reinterpreted in terms of hypothesis testing. This was first considered in [35] for ϵ-differential privacy and then extended to (ϵ, δ)-differrential privacy in [21]. Here we will present and discuss a quantum generalization of this analogy. Besides the intuitive formulation, we will see that it allows for a convenient graphical representation of differential privacy and simple proofs of some additional properties. Before we start, we remark that also Rényi differential privacy has recently been discussed in terms of hypothesis testing [7] but we leave its quantum generalization for future work.

The basic setup we will discuss is binary hypothesis testing between a state $\mathcal{A}(\rho)$, the null hypothesis, and a state $\mathcal{A}(\sigma)$, the alternative hypothesis. If we were able to discriminate between the two states, we could infer which input state was used. We are therefore interested in the corresponding probabilities of error, which are the Type I error $\alpha = \mathrm{Tr}(I-M)\mathcal{A}(\rho)$ of falsely rejecting the null hypothesis and the Type II error $\beta = \mathrm{Tr}\, M\mathcal{A}(\sigma)$ of falsely accepting it.

As differential privacy has to hold for all neighbouring input states and all measurements we get the following set of restrictions on the Type I end Type II errors,

$$1-\alpha \leq e^\epsilon \beta + \delta, \tag{V.9}$$
$$1-\beta \leq e^\epsilon \alpha + \delta \tag{V.10}$$
$$\beta \leq e^\epsilon (1-\alpha) + \delta \tag{V.11}$$
$$\alpha \leq e^\epsilon (1-\beta) + \delta, \tag{V.12}$$

which follow by exchanging $\rho \leftrightarrow \sigma$ and $M \leftrightarrow (I-M)$ in the definition of differential privacy. Based on these inequalities we can define the privacy region of (ϵ, δ)-differential privacy as

$$\mathcal{R}(\epsilon,\delta) = \{(\alpha,\beta) \mid \text{Equations (V.9)-(V.12) hold}\}. \tag{V.13}$$

Next we define the privacy region of a quantum algorithm $\mathcal{A}$ as

$$\mathcal{R}(\mathcal{A}) = \{(\mathrm{Tr}(I-M)\mathcal{A}(\rho), \mathrm{Tr}\, M\mathcal{A}(\sigma)) \mid 0 \leq M \leq I \text{ and } \rho \sim \sigma\}. \tag{V.14}$$

This allows us to state differential privacy in terms of privacy regions.

Theorem V.7. *A quantum channel $\mathcal{A}$ is (ϵ, δ)-differentially private if and only if*

$$\mathcal{R}(\mathcal{A}) \subseteq \mathcal{R}(\epsilon, \delta). \tag{V.15}$$

Proof. The proof is a direct consequence of the above definitions. Note that we could have equivalently defined $\mathcal{R}(\epsilon,\delta)$ with any subset of the Equations (V.9)-(V.12), for example only picking the first one, however this representation will be beneficial later on. □

We continue by stating some properties of the risk regions.

Lemma V.8. *The following holds.*

- *(Concatenation) For arbitrary quantum algorithms $\mathcal{A}$ and $\mathcal{N}$ we have*

$$\mathcal{R}(\mathcal{N}\circ\mathcal{A}) \subseteq \mathcal{R}(\mathcal{A}). \tag{V.16}$$

- *(Symmetry) It holds that $\mathcal{R}(\epsilon,\delta)$ is symmetric with respect to the line $\alpha+\beta=1$.*

Proof. The first statement follows by noting that

$$\mathcal{R}(\mathcal{N}\circ\mathcal{A}) = \{(\mathrm{Tr}(I-N)\mathcal{A}(\rho), \mathrm{Tr}\, N\mathcal{A}(\sigma)) \mid N = \mathcal{N}^\dagger(M) \text{ and } 0 \leq M \leq I \text{ and } \rho\sim\sigma\} \tag{V.17}$$
$$\subseteq \mathcal{R}(\mathcal{A}), \tag{V.18}$$

which follows because $\mathcal{N}^\dagger$ is completely positive and unital. The second statement follows directly from examining Equations (V.9)-(V.12). □

Let us have a look at the graphical representation implied by the above definitions. In Figure 7 we give several examples of $\mathcal{R}(\epsilon,\delta)$ that illustrate the risk region of differential privacy. We also want to give a concrete numerical example of how the risk region of a channel contracts. Let's consider a simple example where only two qubit input states ρ and σ are available which are considered neighbouring. The states are chosen such that $E_1(\rho,\sigma) \leq \frac{1}{3}$. For simplicity we consider a trivial algorithm to which we now want to add depolarizing noise such that the outputs become $(0.2, 0.01)$-differentially private. From Equation (IV.3) we can estimate that this should be the case if we choose $p \approx 0.72$. To verify our observation numerically we simulate $\mathcal{R}(\mathcal{D}_p)$ by drawing random POVMs and compare the resulting pairs (α,β) to the desired privacy region. We can observe in Figure 8

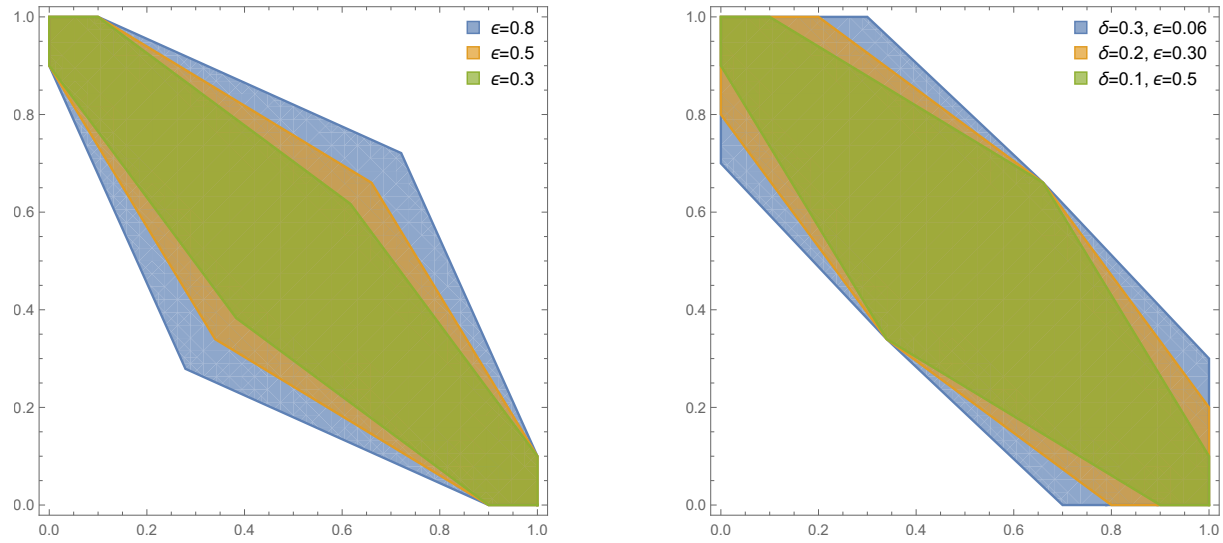

Fig. 7: Examples for $\mathcal{R}(\epsilon, \delta)$. Left: plotted for $\delta = 0.1$ and different values of ϵ. Right: plotted for different values of δ and ϵ, with ϵ chosen according to Lemma V.9.

that this is indeed consistent with what we expected, namely that for $p = 0.72$ all drawn points are within $\mathcal{R}(0.2, 0.01)$ while for smaller values of p the noise is clearly not sufficient for $(0.2, 0.01)$-differential privacy.

Finally, we will see that phrasing differential privacy in terms of hypothesis testing does not only have advantages in terms of intuition, but also allows us to easily prove some useful results.

Lemma V.9. *If $\mathcal{A}$ is (ϵ, δ)-differentially private, then it is also $(\tilde{\epsilon}, \tilde{\delta})$-differentially private with $\tilde{\delta} \geq \delta$ and*

$$e^{\tilde{\epsilon}} \geq \frac{(1-\tilde{\delta})}{(1-\delta)}(1+e^{\epsilon}) - 1. \tag{V.19}$$

Proof. This follows easily from the graphical representation. We want to prove that

$$\mathcal{R}(\epsilon, \delta) \subseteq \mathcal{R}(\tilde{\epsilon}, \tilde{\delta}).$$

Let us consider the lower bounds on the risk region. It can easily be checked that they coincide at the point

$$(\alpha^*, \beta^*) = (\frac{1-\delta}{1+e^{\epsilon}}, \frac{1-\delta}{1+e^{\epsilon}}),$$

which gives a corner point of the region. Since we require $\tilde{\delta} \geq \delta$, it suffices if

$$\frac{1-\tilde{\delta}}{1+e^{\tilde{\epsilon}}} \leq \frac{1-\delta}{1+e^{\epsilon}}.$$

This gives the claimed bound on $\tilde{\epsilon}$. □

This result allows us to observe a certain trade-off between ϵ and δ, in particular, by raising δ one can get differential privacy with a better value of ϵ. This is also illustrated in the right part of Figure 7.

VI. Conclusions

In this work we gave a new approach to exploring quantum differential privacy via information theoretic tools. In particular, we used the quantum hockey-stick divergence to give a simple framework in which we can bound differential privacy parameters for practically relevant noise models such as quantum circuits and quantum neural networks implemented on near-term quantum devices. This includes comparing local and global depolarizing noise and contrasting achieving differential privacy with an undesirable decay in trace distance. On the way we showed several new properties of the said divergence and gave it a new operational interpretation.

Given that our approach promises to be simpler and more powerful than the previously used ones we expect it to play a crucial role going forward when investigating differential privacy.

Naturally we are left with some open problems. In Lemma III.6 we showed a bound on the measured relative entropy of the outputs of a differentially private algorithm. Classically several such results are known, including bounds on other entropic quantities such as the mutual information. An interesting open problem going forward is to find bounds on fully quantum quantities such as the quantum relative entropy or quantum mutual information. This would allow investigating connections to other privacy related quantities such as the quantum privacy funnel [10] generalizing work from the classical setting [31].

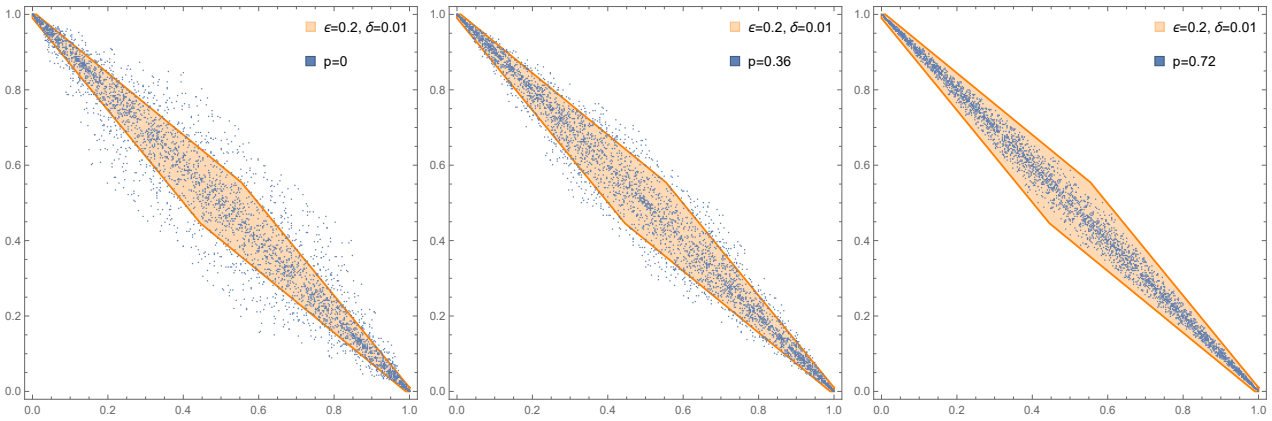

Fig. 8: Numerical values for points in $\mathcal{R}(\mathcal{D}_p)$ for different values of p. Points are based on two qubit states with $E_1(\rho,\sigma) \leq \frac{1}{3}$ and 1000 randomly drawn POVMs. The orange region in the background corresponds to $\mathcal{R}(0.2, 0.01)$.

ACKNOWLEDGMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie Grant Agreement No. H2020-MSCA-IF-2020-101025848. DSF acknowledges financial support from the VILLUM FONDEN via the QMATH Centre of Excellence (Grant no. 10059) the QuantERA ERA-NET Cofund in Quantum Technologies implemented within the European Union's Horizon 2020 Programme (QuantAlgo project) via the Innovation Fund Denmark and from the European Research Council (grant agreement no. 81876). CR is partially supported by a Junior Researcher START Fellowship from the MCQST. CR acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC-2111 390814868.

REFERENCES

[1] S. Aaronson and G. N. Rothblum. Gentle measurement of quantum states and differential privacy. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 322–333, 2019.

[2] A. Angrisani, M. Doosti, and E. Kashefi. Differential privacy amplification in quantum and quantum-inspired algorithms, 2022.

[3] A. Angrisani and E. Kashefi. Quantum local differential privacy and quantum statistical query model, 2022.

[4] S. Arunachalam, Y. Quek, and J. Smolin. Private learning implies quantum stability. *arXiv preprint arXiv:2102.07171*, 2021.

[5] S. Asoodeh, M. Aliakbarpour, and F. P. Calmon. Local differential privacy is equivalent to contraction of e_γ-divergence. *arXiv preprint arXiv:2102.01258*, 2021.

[6] S. Asoodeh, M. Diaz, and F. P. Calmon. Privacy analysis of online learning algorithms via contraction coefficients. *arXiv preprint arXiv:2012.11035*, 2020.

[7] B. Balle, G. Barthe, M. Gaboardi, J. Hsu, and T. Sato. Hypothesis testing interpretations and renyi differential privacy. In *International Conference on Artificial Intelligence and Statistics*, pages 2496–2506. PMLR, 2020.

[8] M. Berta, O. Fawzi, and M. Tomamichel. On variational expressions for quantum relative entropies. *Letters in Mathematical Physics*, 107(12):2239–2265, 2017.

[9] P. J. Coles, J. Kaniewski, and S. Wehner. Equivalence of wave–particle duality to entropic uncertainty. *Nature Communications*, 5(5814), December 2014. arXiv:1403.4687.

[10] N. Datta, C. Hirche, and A. Winter. Convexity and operational interpretation of the quantum information bottleneck function. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 1157–1161. IEEE, 2019.

[11] G. De Palma, M. Marvian, D. Trevisan, and S. Lloyd. The quantum wasserstein distance of order 1. *IEEE Transactions on Information Theory*, 67(10):6627–6643, 2021.

[12] Y. Du, M.-H. Hsieh, T. Liu, D. Tao, and N. Liu. Quantum noise protects quantum classifiers against adversaries. *Physical Review Research*, 3(2):023153, 2021.

[13] Y. Du, M.-H. Hsieh, T. Liu, S. You, and D. Tao. Quantum differentially private sparse regression learning. *IEEE Transactions on Information Theory*, 68(8):5217–5233, 2022.

[14] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy, data processing inequalities, and minimax rates. *arXiv preprint arXiv:1302.3203*, 2013.

[15] C. Dwork. Differential privacy. In *International Colloquium on Automata, Languages, and Programming*, pages 1–12. Springer, 2006.

[16] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.

[17] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3-4):211–407, 2013.

[18] O. Fawzi, A. Müller-Hermes, and A. Shayeghi. A lower bound on the space overhead of fault-tolerant quantum computation, 2022. arXiv:2202.00119v1.

[19] L. Gao and C. Rouzé. Ricci curvature of quantum channels on non-commutative transportation metric spaces. *arXiv preprint arXiv:2108.10609*, 2021.

[20] C. Hirche, C. Rouzé, and D. S. França. On contraction coefficients, partial orders and approximation of capacities for quantum channels. *arXiv preprint arXiv:2011.05949*, 2020.

[21] P. Kairouz, S. Oh, and P. Viswanath. The composition theorem for differential privacy. In *International conference on machine learning*, pages 1376–1385. PMLR, 2015.

[22] J. Kempe, O. Regev, F. Unger, and R. De Wolf. Upper bounds on the noise threshold for fault-tolerant quantum computing. In *International Colloquium on Automata, Languages, and Programming*, pages 845–856. Springer, 2008.

[23] W. Li, S. Lu, and D.-L. Deng. Quantum federated learning through blind quantum computing. *Science China Physics, Mechanics & Astronomy*, 64(10):1–8, 2021.

[24] J. Liu, P. Cuff, and S. Verdú. E_γ-resolvability. *IEEE Transactions on Information Theory*, 63(5):2629–2658, 2016.

[25] K. Matsumoto. A new quantum version of f-divergence. In *Nagoya Winter Workshop: Reality and Measurement in Algebraic Quantum Theory*, pages 229–273. Springer, 2015.

[26] I. Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.

[27] M. Muller-Lennert, F. Dupuis, O. Szehr, S. Fehr, and M. Tomamichel. On quantum Rényi entropies: A new generalization and some properties. *Journal of Mathematical Physics*, 54(12):122203, jun 2013.

[28] D. Petz. Quasi-entropies for finite quantum systems. *Reports on mathematical physics*, 23(1):57–65, 1986.

[29] A. A. Razborov. An upper bound on the threshold quantum decoherence rate. *arXiv preprint quant-ph/0310136*, 2003.

[30] M. B. Ruskai. Beyond strong subadditivity? improved bounds on the contraction of generalized relative entropy. *Reviews in Mathematical Physics*, 6(05a):1147–1161, 1994.

[31] S. Salamatian, F. P. Calmon, N. Fawaz, A. Makhdoumi, and M. Médard. Privacy-utility tradeoff and privacy funnel. 2020.

[32] M. Senekane, M. Mafu, and B. M. Taele. Privacy-preserving quantum machine learning using differential privacy. In *2017 IEEE AFRICON*, pages 1432–1435. IEEE, 2017.

[33] N. Sharma and N. A. Warsi. On the strong converses for the quantum channel capacity theorems. *arXiv preprint arXiv:1205.1712*, 2012.

[34] M. Tomamichel. *Quantum information processing with finite resources: mathematical foundations*, volume 5. Springer, 2015.

[35] L. Wasserman and S. Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010.

[36] W. M. Watkins, S. Y.-C. Chen, and S. Yoo. Quantum machine learning with differential privacy. *arXiv preprint arXiv:2103.06232*, 2021.

[37] M. M. Wilde, M. Berta, C. Hirche, and E. Kaur. Amortized channel divergence for asymptotic quantum channel discrimination. *Letters in Mathematical Physics*, 110(8):2277–2336, 2020.

[38] M. M. Wilde, A. Winter, and D. Yang. Strong converse for the classical capacity of entanglement-breaking and hadamard channels via a sandwiched rényi relative entropy. *Communications in Mathematical Physics*, 331(2):593–622, July 2014.

[39] Y. Yoshida and M. Hayashi. Classical mechanism is optimal in classical-quantum differentially private mechanisms. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 1973–1977, 2020.

[40] L. Zhou and M. Ying. Differential privacy in quantum computation. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 249–262. IEEE, 2017.

APPENDIX

The following is a generalization of the Fuchs-van-de-Graaf inequality to general positive semi-definite operators proven in [9, Supplementary Lemma 3], see also [37, Appendix B] for an alternative proof.

Lemma A.1 ([9])**.** *For positive semi-definite, trace class operators A and B acting on a separable Hilbert space, we have that*

$$\|A-B\|_1^2 + 4\left\|A^{1/2}B^{1/2}\right\|_1^2 \leq \left(\operatorname{Tr}[A+B]\right)^2. \quad \text{(A.1)}$$

Recently it was proven in [18] that for any quantum channel $\mathcal{T}$ with d its input and output dimension,

$$\eta_1(\mathcal{T}) \leq \sqrt{1 - \frac{\lambda_{\min}(\mathcal{T}^\dagger \circ \mathcal{T})}{d^2}}, \quad \text{(A.2)}$$

where $\lambda_{\min}(\mathcal{T}^\dagger \circ \mathcal{T}) = \min_{\Psi,\Phi} \langle\Psi|(\mathcal{T}^\dagger \circ \mathcal{T})(|\Phi\rangle\langle\Phi|)|\Psi\rangle$. As we would like to generalize this result to the hockey-stick divergence, we extract the following lemma from their proof.

Lemma A.2 ([18])**.** *For any quantum channel $\mathcal{T}$ with d its input and output dimension and pure input states Ψ and Φ,*

$$F(\mathcal{T}(\Psi), \mathcal{T}(\Phi)) \geq \frac{\lambda_{\min}(\mathcal{T}^\dagger \circ \mathcal{T})}{d^2}. \quad \text{(A.3)}$$

It was furthermore noted that for $\mathcal{T} = \mathcal{N}^{\otimes k}$ with $\mathcal{N}$ a qubit channel, one has

$$F(\mathcal{T}(\Psi), \mathcal{T}(\Phi)) \geq \left(\frac{\lambda_{\min}(C_{\mathcal{N}^\dagger \circ \mathcal{N}})}{4}\right)^k, \quad \text{(A.4)}$$

where $C_{\mathcal{N}^\dagger \circ \mathcal{N}}$ is the Choi matrix of $\mathcal{N}^\dagger \circ \mathcal{N}$, and if $\mathcal{N}$ is also non-unital one gets

$$\lambda_{\min}(C_{\mathcal{N}^\dagger \circ \mathcal{N}}) > 0. \quad \text{(A.5)}$$

We also need bounds on $D^\epsilon_{\max}(\rho\|\sigma)$. As remarked earlier, our definition of $D^\epsilon_{\max}$ is a bit different than the one usually used in the quantum information literature, as it uses a different distance measure. However, to apply a known result we need to compare our definition to the usual one. The standard smooth $\max$-relative entropy, $D_{\max,s}$, is defined as

$$D^\epsilon_{\max,s}(\rho\|\sigma) = \inf_{\bar{\rho} \in B^\epsilon_s(\rho)} D_{\max}(\bar{\rho}\|\sigma) \quad \text{(A.6)}$$

and

$$B^\epsilon_s(\rho) = \{\bar{\rho} : \bar{\rho} \in \mathcal{S}_\leq(\mathcal{H}) \wedge P(\rho, \bar{\rho}) \leq \epsilon\}, \quad \text{(A.7)}$$

where P is the purified distance, i.e. the minimal trace distance between purifications of the states. See e.g. [34, Definition 3.15] for a discussion of this quantity. We prove the following auxiliary lemma.

Lemma A.3. *Let $\rho, \sigma \in \mathcal{S}_=(\mathcal{H})$, then*

$$D^\epsilon_{\max}(\rho\|\sigma) \leq D^\epsilon_{\max,s}(\rho\|\sigma) \quad \text{(A.8)}$$

Proof. The claim follows immediately by showing that $B^\epsilon_s(\rho) \subseteq B^\epsilon(\rho)$. To that end observe,

$$\begin{aligned} B^\epsilon_s(\rho) &= \{\bar{\rho} : \bar{\rho} \in \mathcal{S}_\leq(\mathcal{H}) \wedge P(\rho, \bar{\rho}) \leq \epsilon\} && \text{(A.9)} \\ &\subseteq \{\bar{\rho} : \bar{\rho} \in \mathcal{S}_\leq(\mathcal{H}) \wedge E_1(\rho, \bar{\rho}) \leq \epsilon\} && \text{(A.10)} \\ &\subseteq \{\bar{\rho} : \bar{\rho} \in \mathcal{P}(\mathcal{H}) \wedge E_1(\rho, \bar{\rho}) \leq \epsilon\} && \text{(A.11)} \\ &= B^\epsilon(\rho), && \text{(A.12)} \end{aligned}$$

where the second inclusion is clear, but the first needs some justification. Note that

$$E_1(\rho,\bar{\rho}) \leq \max\{E_1(\rho,\bar{\rho}), E_1(\bar{\rho},\rho)\} \tag{A.13}$$
$$= \Delta(\rho,\bar{\rho}) \tag{A.14}$$
$$\leq P(\rho,\bar{\rho}), \tag{A.15}$$

where Δ is the generalized trace distance and the equality holds by its definition. The first inequality is immediate and the second is a generalized Fuchs-van-de-Graaf type inequality, see e.g. [34, Lemma 3.17]. This concludes the proof. □

This enables us to use the following result.

Lemma A.4. *Let* $0 \leq \epsilon \leq 1$ *and* $\alpha \in (1,\infty)$*, then*

$$D^{\epsilon}_{\max}(\rho\|\sigma) \leq \mathbb{D}_{\alpha}(\rho\|\sigma) + \frac{g(\epsilon)}{\alpha-1}, \tag{A.16}$$

where $g(\epsilon) = -\log(1-\sqrt{1-\epsilon^2})$ *and* $\mathbb{D}_{\alpha}$ *any quantum Rényi divergence.*

Proof. The result immediately follows from [34, Proposition 6.22], which shows the same result for $D^{\epsilon}_{\max,s}(\rho\|\sigma)$, and the fact that $D^{\epsilon}_{\max}(\rho\|\sigma) \leq D^{\epsilon}_{\max,s}(\rho\|\sigma)$ proved in Lemma A.3. □

In Section V-B, we introduced the new concept of Rényi quantum differential privacy and based it on a very general framework of Rényi relative entropies that solely have to fulfil certain properties as listed in [34]. For completeness we list those properties here.

A quantum Rényi divergence is a quantity $\mathbb{D}(\cdot\|\cdot)$ that fulfills the following properties:

1) **Continuity:** $\mathbb{D}(\rho\|\sigma)$ is continuous in ρ and σ, wherever $\rho \neq 0$ and $\sigma >> \rho$.
2) **Unitary invariance:** $\mathbb{D}(\rho\|\sigma) = \mathbb{D}(U\rho U^{\dagger}\|U\sigma U^{\dagger})$ for any unitary U.
3) **Normalization:** $\mathbb{D}(1\|\frac{1}{2}) = \log 2$.
4) **Order:** If $\rho \geq \sigma$, then $\mathbb{D}(\rho\|\sigma) \geq 0$ and if $\rho \leq \sigma$ then $\mathbb{D}(\rho\|\sigma) \leq 0$.
5) **Additivity:** $\mathbb{D}(\rho\otimes\tau\|\sigma\otimes\omega) = \mathbb{D}(\rho\|\sigma) + D(\tau\|\omega)$.
6) **General mean:** There exists a continuous and strictly monotonic function g such that $Q := g(\mathbb{D})$ satisfies,
$$Q(\rho\oplus\tau\|\sigma\oplus\omega) = \frac{\mathrm{Tr}(\rho)}{\mathrm{Tr}(\rho+\tau)}Q(\rho\|\sigma) + \frac{\mathrm{Tr}(\tau)}{\mathrm{Tr}(\rho+\tau)}Q(\tau\|\omega).$$
7) **Positive Definiteness:** $\mathbb{D}(\rho\|\sigma) \geq 0$ with equality iff $\rho = \sigma$.
8) **Data-processing:** $\mathbb{D}(\rho\|\sigma) \geq \mathbb{D}(\mathcal{N}(\rho)\|\mathcal{N}(\sigma))$.
9) Either **joint convexity** or **joint concavity** of Q.
10) **Dominance:** For $\sigma \leq \sigma'$, one has $\mathbb{D}(\rho\|\sigma) \geq \mathbb{D}(\rho\|\sigma')$.

In the classical case, properties 1-6 uniquely define the Rényi relative entropies. This is not the case in the quantum setting where one additionally requires the operationally motivated properties 7-10.

Finally, a family of quantum Rényi relative entropies is a one-parameter family $\alpha \to \mathbb{D}_{\alpha}(\cdot\|\cdot)$ of quantum Rényi relative entropies such that for some open interval containing 1, the family is monotonically increasing in α.

Christoph Hirche is a Marie Skłodowska-Curie Global Fellow at the Technical University Munich, Germany, and the National University of Singapore. His current research focuses on quantum information theory. Previously, he was a Postdoctoral researcher at QMATH, University of Copenhagen (Denmark). He obtained his PhD from the Universitat Autonoma de Barcelona (Spain). Prior to his PhD studies, he obtained Bachelor and Master degrees in Physics from the University of Hannover (Germany).

Cambyse Rouzé Cambyse Rouzé received the bachelor's degree from the École Centrale Paris in 2012, the master's degree in theoretical physics from the École Normale Supérieure Paris in 2013, and the master of advanced study degree in mathematics and the Ph.D. degree from the University of Cambridge, in 2014 and 2018, respectively. From 2019 to 2022, he was a TUFF and START Fellow at the Department of Mathematics, Technical University of Munich (TUM), and the Munich Center for Quantum Science and Technology (MCQST). He is currently a Humboldt Fellow at the TUM. His research interests include quantum information theory, continuous variables quantum systems, quantum functional inequalities, and quantum many-body systems

Daniel Stilck França holds an Inria Starting Faculty position hosted by the École Normale Supérieure de Lyon (ENS Lyon). He works in quantum information and computation with a focus on mathematical aspects of these fields. Before joining ENS, Daniel was a postdoc at the Qmath group at the University of Copenhagen (Denmark), and before that he did his PhD at the Technical University of Munich (Germany).