



HAL
open science

Efficient learning of the structure and parameters of local Pauli noise channels

Cambyse Rouzé, Daniel Stilck Franca

► **To cite this version:**

Cambyse Rouzé, Daniel Stilck Franca. Efficient learning of the structure and parameters of local Pauli noise channels. 2023. hal-04386732

HAL Id: hal-04386732

<https://hal.science/hal-04386732v1>

Preprint submitted on 11 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Efficient learning of the structure and parameters of local Pauli noise channels

Cambyse Rouzé *

Zentrum Mathematik, Technische Universität München, 85748 Garching, Germany

Daniel Stilek França

Univ Lyon, ENS Lyon, UCBL, CNRS, Inria, LIP, F-69342, Lyon Cedex 07, France[†]

The unavoidable presence of noise is a crucial roadblock for the development of large-scale quantum computers and the ability to characterize quantum noise reliably and efficiently with high precision is essential to scale quantum technologies further. Although estimating an arbitrary quantum channel requires exponential resources, it is expected that physically relevant noise has some underlying local structure, for instance that errors across different qubits have a conditional independence structure. Previous works showed how it is possible to estimate Pauli noise channels with an efficient number of samples in a way that is robust to state preparation and measurement errors, albeit departing from a known conditional independence structure.

We present a novel approach for learning Pauli noise channels over n qubits that addresses this shortcoming. Unlike previous works that focused on learning coefficients with a known conditional independence structure, our method learns both the coefficients and the underlying structure. We achieve our results by leveraging a groundbreaking result by Bresler for efficiently learning Gibbs measures and obtain an optimal sample complexity of $\mathcal{O}(\log(n))$ to learn the unknown structure of the noise acting on n qubits. This information can then be leveraged to obtain a description of the channel that is close in diamond distance from $\mathcal{O}(\text{poly}(n))$ samples. Furthermore, our method is efficient both in the number of samples and postprocessing without giving up on other desirable features such as SPAM-robustness, and only requires the implementation of single qubit Cliffords. In light of this, our novel approach enables the large-scale characterization of Pauli noise in quantum devices under minimal experimental requirements and assumptions.

INTRODUCTION

With the increasing size and quality of quantum devices in the last years, it becomes increasingly important to develop tools that can characterize and identify multiple (possibly correlated) error rates. One particular class of channels that has received considerable attention over the recent years is that of Pauli channels [1–5]. This is justified by the fact that randomized compiling techniques can bring the noise affecting a quantum device into this normal form [6] and that it models well logical errors [5]. Furthermore, Pauli channels have much more mathematical structure than arbitrary quantum channels, making it easier to treat them analytically. Unfortunately, even the simpler class of Pauli channels requires a number of samples that scales exponentially in the number of qubits [1] to be characterized globally, i.e. in the diamond norm.

However, it is expected that noise in physical systems will exhibit some locality, i.e. the probability of different Pauli errors exhibits some conditional independence structure, as illustrated and explained further in Fig. 1. In this case, given a hypergraph that models the conditional independence structure of different Pauli errors and assuming that the size of the largest hyperedge is constant, the authors of [2, 7, 8] showed how to learn the quantum channel with a number of samples that scales polynomially in the system’s size. Furthermore, they can

achieve this even in a way that is robust to state preparation and measurement errors (SPAM) and without using auxiliary qubits. However, these works need to assume that the conditional independence structure is known in advance, which limits the application of their methods in practice. In this work, we address and solve this bottleneck in estimating Pauli noise by devising a learning protocol that is both computationally efficient and can learn the conditional independence structure. Furthermore, our method shares the desirable features of previous protocols, namely efficient sample complexity and robustness to SPAM.

To achieve our results, we draw upon recent results on learning Markovian fields [9] and adapt them to our setting. The main technical difficulty we overcome is that these classical works assume access to samples from the underlying distribution. In contrast, without the use of auxiliary qubits it is not possible to sample from the probability distribution of various Pauli errors. To overcome this difficulty, we develop a new method that combines features from randomized benchmarking protocols [10–14] together with shadow process tomography protocols [15–17], and which only requires the implementation of single qubit Cliffords. Akin to process shadows, our protocol allows for the estimation of all marginals of Pauli errors of a given size with a number of samples that scales logarithmically in the number of qubits and exponentially in the size of the marginals. And like random-




FIG. 1. The region B of qubits (in blue) shields regions A (in red) and C (in black) away from each other. A measure μ satisfies conditional independence if $\mu(P_A, P_C | P_B) = \mu(P_A | P_B) \mu(P_C | P_B)$ for any such regions. In terms of Pauli errors, it means that given the errors that occurred in B , the distributions of errors in A and C is independent.

ized benchmarking, this is done in a SPAM-robust way. We then show how to run classical algorithms for learning Markovian fields solely with estimates of the marginals.

As our protocol checks all boxes for a practical and effective Pauli learning algorithm (no auxiliary systems required, SPAM-robustness, only requires single qubit Cliffords, optimal sample and efficient computational complexity and no knowledge of conditional independence structure required), we believe it will find widespread applicability to characterize large scale quantum devices.

PRELIMINARIES

Let $d = 2^n$ be the dimension of an n -qubit system. We use the notation $[d] := \{1, \dots, d\}$. A quantum state is a positive semi-definite Hermitian matrix of trace 1. We denote the identity matrix by $\mathbb{I}_d \in \mathbb{C}^{d \times d}$ and by $\text{id}_d : \mathbb{C}^{d \times d} \rightarrow \mathbb{C}^{d \times d}$ the identity map, and omit the subscript if the dimension is clear from context. A quantum channel is a map $\mathcal{N} : \mathbb{C}^{d \times d} \rightarrow \mathbb{C}^{d \times d}$ of the form $\mathcal{N}(\rho) = \sum_k A_k \rho A_k^\dagger$ where the Kraus operators $\{A_k\}_k$ satisfy $\sum_k A_k^\dagger A_k = \mathbb{I}$. If the quantum channel \mathcal{N} satisfies further $\mathcal{N}(\mathbb{I}) = \mathbb{I}$, it is called *unital*.

We define the diamond distance between two quantum channels \mathcal{N} and \mathcal{M} as the diamond norm of their difference:

$$\|\mathcal{N} - \mathcal{M}\|_\diamond = \max_{\phi: \text{tr}(\phi) = 1} \|\text{id}_d \otimes (\mathcal{N} - \mathcal{M})(|\phi\rangle\langle\phi|)\|_1$$

where the Schatten 1-norm of a matrix M is defined as $\|M\|_1 = \text{tr}(M^\dagger M)$.

Pauli channels are a special class of quantum channels whose Kraus operators are weighted Pauli operators. Formally, a Pauli quantum channel \mathcal{P} can be written as follows:

$$\mathcal{P}(\rho) = \sum_{P \in \mathbb{P}_n} \mu(P) P \rho P \quad (1)$$

where we recall that the Pauli matrices X, Y, Z take the form

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (2)$$

and where μ is a probability distribution over the set $\mathbb{P}_n = \{\mathbb{I}, X, Y, Z\}^{\otimes n}$ of n -fold tensor products of Pauli matrices (or Pauli strings).

The elements of \mathbb{P}_n either commute or anticommute: let P and Q be two Pauli operators, we have $PQ = (-1)^{P \cdot Q} QP$, where $P \cdot Q := 0$ if $[P, Q] = 0$ and $P \cdot Q := 1$ otherwise. Furthermore, we will denote the number of non-identity elements of a Pauli string P , its weight, by $w(P)$.

We consider the Pauli channel tomography problem which consists of learning a Pauli channel in the diamond norm. Given a precision parameter $\epsilon > 0$, the goal is to construct a Pauli channel $\hat{\mathcal{P}}$ satisfying with at least a probability $2/3$:

$$\|\mathcal{P} - \hat{\mathcal{P}}\|_\diamond \leq \epsilon.$$

An algorithm \mathcal{A} is $1/3$ -correct for this problem if it outputs a Pauli channel ϵ -close to \mathcal{P} with a probability of error at most $1/3$. We choose to learn in the diamond norm because it characterizes the minimal error probability to distinguish between two quantum channels when auxiliary systems are allowed [18]. Since the diamond norm between two Pauli channels is exactly twice the total variation distance between their corresponding probability distributions [13], approximating the Pauli channel \mathcal{P} in diamond norm is equivalent to approximating the probability distribution μ in total variation distance. The latter is defined for two probability distributions p and q on $[d]$ as follows:

$$\|p - q\|_{\text{TV}} = \frac{1}{2} \sum_{i=1}^d |p_i - q_i|.$$

The learner can only extract classical information from the unknown Pauli channel \mathcal{P} by performing a measurement on the output state for a choice of the input state. Throughout the paper, we only consider unentangled or incoherent measurements. That is, the learner can only measure with an n -qubit measurement device, and using auxiliary qubits or measuring multiple copies at once is not allowed. Furthermore, we only require product Pauli

eigenstates as input states and measurements in product Pauli bases. This corresponds to a minimal and feasible experimental setup.

It is possible to identify the Pauli strings modulo phases with the space $(\mathbb{Z}_2 \times \mathbb{Z}_2)^n$ [18]. Through this identification, it is possible to see that for some Pauli string Q

$$\alpha_Q := \frac{1}{2^n} \text{tr}(Q \mathcal{P}(Q)) = \sum_{P \in \mathbb{P}_n} (-1)^{P \cdot Q} \mu(P) \quad (3)$$

is nothing but the Fourier coefficient of the distribution. Thus, it is possible to estimate the marginals of the distribution μ from such Fourier coefficients, as we will explain in more detail below. In the appendix, we provide a simple scheme to learn these coefficients in a way that is robust to both state preparation and measurement (SPAM) errors and is inspired by randomized benchmarking (RB) combined with recent classical process shadows (PC) protocols [15–17]. It combines the best of these two types of protocols: SPAM-robustness (a feature of RB, but not of PC) and parallelization of estimates (a feature of PC, but not of RB). For our protocol, we assume that we can prepare product Pauli eigenstates and measure in Pauli bases, and that the SPAM error is independent of the state and measurement and modeled by quantum channels Φ_1, Φ_2 (see Fig.3). Under these conditions, we have:

Proposition 1 (see Corollary 7). *Let P_1, \dots, P_m be m Pauli strings such that $w(P_i) \leq w$, denote by p_0 the probability that no error occurs, i.e. $p_0 = \mu(\mathbf{I})$, and suppose that for a constant C_{SPAM} we have for all $1 \leq i \leq m$:*

$$2^{-2n} |\text{tr}(\Phi_1^*(P_i)P_i) \text{tr}(\Phi_2^*(P_i)P_i)| \geq C_{\text{SPAM}}. \quad (4)$$

Then $\mathcal{O}(3^w \epsilon^{-2} C_{\text{SPAM}}^{-2} \log(m\delta^{-1}))$ state preparation and measurements in the setting described above suffice to estimate all α_{P_i} up to an error $\epsilon(1 - p_0)$ with probability of success at least $1 - \delta$. Furthermore, the maximal number k of uses of the channel needed in each run satisfies $k = \mathcal{O}(\log(1/(1 - p_0)))$.

From Fourier coefficients, one can easily reconstruct marginals of the distribution μ as follows: denoting by $A \subset [n]$ a certain set of qubits, the marginal μ_A is defined as

$$\mu_A(Q) := \sum_{P \in \mathbb{P}_{\bar{A}}} \mu(Q \otimes P) = \sum_{P' \in \mathbb{P}_{\bar{A}}} (-1)^{Q \cdot P'} \alpha_{P'}, Q \in \mathbb{P}_A$$

where \mathbb{P}_X denotes the set of Pauli strings over a region $X \subset [n]$. For ease of notations, we denote by P_X the Pauli sub-string generated by $P \in \mathbb{P}_n$ on a region $X \subset [n]$.

In this paper, we will assume that the underlying distribution μ has full support and has a conditional independence structure modeled by a (potentially unknown)

hypergraph $\mathcal{H} = (V, H)$ with n vertices corresponding to the qubits. In a nutshell, the hypergraph encodes the fact that, if two sets of vertices A, C are shielded away from each other by a third set of vertices B , then the probability that a certain error occurs in vertices A, C is independent conditioned on the errors that happened on B . To illustrate this concept, imagine three qubits arranged on a line such that qubit 1, 2 and 2, 3 interact but not 1, 3. Then we could assume that, once we fixed that a given error happened on qubit 2, the probability of errors happening on 1 and 3 are independent.

The Hammersley-Clifford theorem [19] then implies that such a measure μ corresponds to a Gibbs distribution. This means that

$$\mu(P) := \exp\left(\sum_{\ell=1}^r \sum_{h=(i_1, \dots, i_\ell) \in H} \theta^h(P_h) - C\right). \quad (5)$$

Here $\theta^h : \mathbb{P}_h \rightarrow \mathbb{R}$ is a function that takes as inputs Pauli matrices on qubits $i_1 \dots i_\ell$. These functions are referred to as clique potentials. In the equation above, C is a constant that ensures the distribution is normalized and is called the log-partition function. The parameter r denotes the maximal range of a potential, i.e. how many qubits it acts on. We also assume w.l.o.g. that $\sum_{Q \in \mathbb{P}_h} \theta^h(Q) = 0$, as we can always shift the potential by a constant without changing the distribution.

A hyperedge h is called maximal if no other hyperedge of strictly larger size contains h . A graph $G = (V, E)$ can be obtained by replacing every hyperedge with a clique. In other words, each clique in the graph corresponds to a local interaction in H . Let D be a bound on the maximum degree of the graph. We denote by $\Gamma(u)$ the neighbors of u . We further require the following conditions to ensure that we can learn the hypergraph:

- (a) Every edge $(i, j) \in E$ is contained in some hyperedge $h \in H$ where the corresponding tensor is non-zero.
- (b) For every maximal hyperedge $h \in H$, θ^h takes at least one value lower bounded by some constant $\alpha > 0$ in absolute value.
- (c) All the values θ^h takes are upper bounded by a constant $\beta > 0$ in absolute value.

LEARNING THE CONDITIONAL INDEPENDENCE STRUCTURE

Our first task is to learn the set H of hyperedges. For this, we adapt a greedy approach by Bresler [9]. Given a site u and regions I, S , first introduce the functional

$$\nu_{u, I|S} := \mathbb{E}_{R, G} \mathbb{E}_{P_S \sim \mu_S} \left[|\Delta_{u, I|S}| \right], \quad (6)$$

where $\Delta_{u,I|S} := \mathbb{P}_\mu(P_u = R, P_I = G|P_S) - \mathbb{P}_\mu(P_u = R|P_S)\mathbb{P}_\mu(P_I = G|P_S)$, R is a configuration drawn uniformly at random from \mathbb{P}_u and G is an $|I|$ -tuple drawn uniformly at random from \mathbb{P}_I . Finally, P_S is drawn according to the unknown distribution μ_S . This quantity will play a crucial role in our algorithm. In words, it quantifies the amount by which the site u and region I are correlated, conditioned on the nature of the configuration of Pauli matrices in the region S . In spirit, the larger $\nu_{u,I|S}$, the more likely u and I are to be neighboring, and it will be 0 if S shields u away from I . Since we do not have direct access to this quantity due to its dependence on the unknown distribution $\mu_{\{u\} \cup I \cup S}$, we also define an estimate for it in the standard way:

$$\widehat{\nu}_{u,I|S} := \mathbb{E}_{R,G} \widehat{\mathbb{E}}_{P_S} \left[|\widehat{\Delta}_{u,I|S}| \right]. \quad (7)$$

Here, we assume access to an approximation $\widehat{\mu}_{\{u\} \cup I \cup S}$ of the distribution $\mu_{\{u\} \cup I \cup S}$ which we will later construct from access to the Fourier coefficients introduced in (3). The average $\widehat{\mathbb{E}}_{P_S}$ is therefore defined with respect to the marginal $\widehat{\mu}_S$, whereas $\widehat{\Delta}_{u,I|S} := \mathbb{P}_{\widehat{\mu}_{\{u\} \cup I \cup S}}(P_u = R, P_I = G|P_S) - \mathbb{P}_{\widehat{\mu}_{\{u\} \cup I \cup S}}(P_u = R|P_S)\mathbb{P}_{\widehat{\mu}_{\{u\} \cup I \cup S}}(P_I = G|P_S)$. The following algorithm will succeed in finding the neighborhood of a node u with high probability as long as $\widehat{\nu}_{u,I|S}$ is sufficiently close to $\nu_{u,I|S}$. The algorithm also depends on two parameters $\tau, L > 0$, where τ serves as a thresholding constant, whereas L upper bounds the size of the superset of a neighborhood of u . Both these parameters will be fixed later on.

Algorithm 1: NeighborhoodLearning

Input: Vertex $u \in [n]$, $L, \tau \geq 0$.

Output: Estimated neighborhood S of u .

Initialize $S := \emptyset$;

while $|S| \leq L$ and $\exists I \subset [n] \setminus S$, $|I| \leq r - 1$

s.t. $\widehat{\nu}_{u,I|S} > \tau$ **do**

 | $S \leftarrow S \cup I$;

end

for $i \in S$ **do**

 | if $\widehat{\nu}_{u,i|S \setminus i} < \tau$, $S \leftarrow S \setminus \{i\}$;

end

return S .

Next, given $\ell \in \mathbb{N}$ and $\epsilon > 0$, we denote by $A(\ell, \epsilon)$ the event that for all u, I, S with $|I| \leq r - 1$ and $|S| \leq \ell$ simultaneously, $|\nu_{u,I|S} - \widehat{\nu}_{u,I|S}| \leq \epsilon$. In particular, we denote $A \equiv A(L, \tau/2)$. We further choose $\tau := \frac{2\alpha^2 \delta^{r+d-1}}{r^{2r} 4^{r+1} \binom{D}{r-1} \gamma e^{2\gamma}}$ and $L := (8/\tau^2) \ln(4)$, with $\gamma := \sup_{u \in [n]} \sum_{\ell=1}^r \sum_{i_2 < \dots < i_\ell} |\theta^{ui_2 \dots i_\ell}|_{\max}$ and $\eta := \frac{1}{4} e^{-2\gamma}$. Here $|T|_{\max}$ denotes the maximum entry of a tensor T . In [20, Theorem 5.7], it is shown that whenever the event A occurs, Algorithm 1 returns the correct set of neighbors of u for every node u . It remains to estimate the number of samples required for event A to occur with

high probability. This is done in the next Lemma.

Lemma 2. *The event A occurs with probability $1 - \delta$ when given access to*

$$\mathcal{O} \left(\log \left(\frac{1}{1-p_0} \right) \frac{4^{3r+5L}}{(C_{\text{SPAM}} \tau \eta^L)^2} \log \frac{n^{r+L}}{\delta} \right) \quad (8)$$

copies of the channel \mathcal{P} , where $p_0 = \mu(\mathbb{I})$.

Proof. Since we assume that $|I| \leq r - 1$ and $|S| \leq L$, the set $\{u\} \cup I \cup S$ has cardinality at most $r + L$. The number of possible subsets of $[n]$ with at most $r + L$ elements is upper bounded by n^{r+L} , and each such region corresponds to at most 4^{r+L} different configurations. Therefore, to learn all the Fourier coefficients α_P of Pauli strings of weight at most $r + L$ to precision $\tilde{\epsilon}$ with probability $1 - \delta$, it suffices to run the protocol of Cor. 7, which requires access to $\mathcal{O}(-\log(1-p_0) 3^{r+L} \tilde{\epsilon}^{-2} C_{\text{SPAM}}^{-2} \log((r+L)n^{r+L} 4^{r+L} \delta^{-1}))$ copies of the channel \mathcal{P} . Once the coefficients α_P are all known, we compute their inverse Fourier transform, since for any P_A supported in $A \subset [n]$,

$$\mu_A(P_A) = \sum_{Q \in \mathbf{P}_A} (-1)^{P_A \cdot Q} \alpha_Q. \quad (9)$$

Now, knowing the coefficients α_Q to precision $\tilde{\epsilon}$ implies that we can compute the coefficients $\mu_A(P_A)$ to precision at least $4^{r+L} \tilde{\epsilon}$. By [20, Lemma 5.2], this further implies that we can compute $\nu_{u,I|S}$ to precision ϵ as long as $4^{r+L} \tilde{\epsilon} \leq \epsilon 4^{-L} \frac{\eta^L}{5}$. The result follows from choosing $\tilde{\epsilon} = \frac{\tau}{2} 4^{-2L-r} \frac{\eta^L}{5}$. \square

Combining this with [20, Theorem 5.7], we conclude the following.

Theorem 3. *There exists a protocol robust to SPAM errors which uses a number of queries to \mathcal{P} as in Eq. (8) and returns the conditional independence structure G with probability $1 - \delta$. Each query only requires implementing 1-qubit Clifford gates circuits of the form of that of Figure 3 and at most $\mathcal{O}(\log(1/(1-p_0)))$ sequential queries to \mathcal{P} are made per run of a circuit. Moreover, the protocol queries n times the protocol of Algorithm 1, whose run time is of order $\tilde{\mathcal{O}}(Ln^r)$.*

LEARNING THE COEFFICIENTS OF THE INTERACTION

Next, we assume that we have already successfully learned the underlying conditional independence structure, for instance by running Algorithm 1, and we aim at learning the values of the coefficients $\theta^h(P_{\{h\}})$ for any $h \in \mathcal{H}$. Since the distribution μ is assumed to be Gibbs, it is well-known that this task can be achieved

efficiently with access to $\mathcal{O}(\log(n))$ samples drawn from local marginals. The idea behind this fact is that the effective Hamiltonian of the marginal of a Gibbs measure is also local. Thus, learning can be performed locally (see e.g. [21] for a proof in the more general context of a commuting Hamiltonian). The following result therefore follows from the fact that these marginals can be accessed via the Fourier coefficients, cf. Equation 9. In what follows, for any hyperedge $h \in \mathcal{H}$, we denote by $R_h \subseteq [n]$ the smallest set of vertices that contains h in its strict interior (i.e. such that h does not overlap with the inner boundary $\partial_- R_h$ of R_h , that is vertices which are contained in other hyperedges $h' \neq h$). We will also need the orthogonal basis $\{\chi_P := Q \mapsto (-1)^{P \cdot Q}\}_{P \in \mathbb{P}_1}$ of characters over $\mathbb{P}_1 \simeq \mathbb{Z}_2 \times \mathbb{Z}_2$, and more generally define the product basis $\{\chi_P\}_{P \in \mathbb{P}_A}$ over a region A by taking products of the single qubit characters.

Algorithm 2: CoefficientLearning

Input: $h \in H$.

Output: Estimated interactions $\hat{\theta}^h$.

return $\hat{\theta}^h := \frac{1}{4^{|R_h|}} \sum_{P \in \mathbb{P}_h} \chi_P \langle \chi_P, \log(\hat{\mu}_{R_h}) \rangle$.

Theorem 4. *Assuming that the hypergraph $\mathcal{H} = (V, H)$ is known, there exists a protocol robust to SPAM errors which uses a number of queries to \mathcal{P} of order*

$$\mathcal{O} \left(\log \left(\frac{1}{1-p_0} \right) \frac{e^{\beta r D^{r+1}} 4^{6r D^{r+1}} \log n^{r D^{r+1}}}{(\epsilon C_{\text{SPAM}})^2} \log \frac{n^{r D^{r+1}}}{\delta} \right). \quad (10)$$

It utilizes at most $\mathcal{O}(\log(1/(1-p_0)))$ sequential queries to \mathcal{P} and circuits consisting of single-qubit Clifford gates like in Figure 3, and returns, for each $h \in H$, an estimator $\hat{\theta}^h$ of θ^h such that $\|\hat{\theta}^h - \theta^h\|_\infty \leq \epsilon$ with probability $1 - \delta$. Furthermore, the algorithm queries $\mathcal{O}(n^{r D^{r+1}})$ times the protocol of Algorithm 2, whose run time is of order $\tilde{\mathcal{O}}(1)$.

Proof. By a simple adaptation of the proof in [21], it can be shown that the estimator $\hat{\theta}^h$ from Algorithm 2 satisfies $\|\hat{\theta}^h - \theta^h\|_\infty \leq \epsilon$ as soon as $\|\mu_{R_h} - \hat{\mu}_{R_h}\|_{\text{TV}} \leq \epsilon e^{-|R_h| |\theta^h|} 4^{-|R_h| - 2r}$. Since $|R_h| \leq r D^{r+1}$ and $|\theta^h| \leq \beta$ by assumption, the task then reduces to that of estimating all the marginals μ_{R_h} up to an error of $\epsilon e^{-\beta r D^{r+1}} 4^{-r D^{r+1} - 2r}$ in total variation. This can be done in a SPAM error robust manner as in the proof of Lemma 2: since the number of possible subsets of $[n]$ with at most $r D^{r+1}$ elements is upper bounded by $r D^{r+1} n^{r D^{r+1}}$, and since each such region corresponds to at most $4^{r D^{r+1}}$ different configurations, to learn all the corresponding Fourier coefficients to precision ϵ' with probability $1 - \delta$, it suffices to run the protocol of Corollary 7, which requires access to $\mathcal{O}(-\log(1 - p_0) 3^{r D^{r+1}} \epsilon'^{-2} C_{\text{SPAM}}^{-2} \log(r D^{r+1} n^{r D^{r+1}} 4^{r D^{r+1}} \delta^{-1}))$ copies of the channel \mathcal{P} . Moreover, by Equation 9,

knowing these Fourier coefficients to precision ϵ' implies that we can approximate the marginals in sup norm to precision $4^{r D^{r+1}} \epsilon'$, and therefore in total variation to precision $4^{2r D^{r+1}} \epsilon'$. The result follows after setting $4^{2r D^{r+1}} \epsilon' = \epsilon e^{-\beta r D^{r+1}} 4^{-r D^{r+1} - 2r}$. \square

Note that the fact that we can estimate multiple Fourier coefficients in parallel is crucial to ensure the $\log(n)$ scaling of the sample complexity, something previous approaches to learning the coefficients [4, 7, 11, 22] did not achieve.

At this stage, we turn our attention to our original problem of learning local Pauli channels in diamond norm. For this, we introduce some few more notations: Given $Q_h \in \mathbb{P}_h$, we denote $e_{Q_h} := \frac{1}{2^n} \langle \mu, \chi_{Q_h} \rangle$ and $\lambda_{Q_h} := \frac{1}{2^n} \langle \chi_{Q_h}, \log \mu \rangle$, similarly, we denote by $\hat{e}_{Q_h} := \frac{1}{2^n} \langle \hat{\mu}, \chi_{Q_h} \rangle$ and $\hat{\lambda}_{Q_h} := \frac{1}{2^n} \langle \chi_{Q_h}, \log \hat{\mu} \rangle$, where the inner products are with respect to the ℓ_2 norm on \mathbb{P}_n . By Plancherel identity, it is clear that $\|\lambda - \hat{\lambda}\|_{\ell_2}^2 = \sum_{|h| \leq r} \sum_{P_h \in \mathbb{P}_h} |\theta^h(P_h) - \hat{\theta}^h(P_h)|^2$. Moreover, for classical Gibbs measures, the following strong convexity property holds (see e.g. [23, 24]):

$$\|e - \hat{e}\|_{\ell_2} \leq \Gamma \|\lambda - \hat{\lambda}\|_{\ell_2},$$

for some constant $\Gamma = \mathcal{O}(1)$ that depends exponentially on the degree D but is independent of n . Here the euclidean space ℓ_2 is with respect to the set of parameters Q_h for $h \in H$, i.e.

$$\|x\|_{\ell_2}^2 := \sum_{h \in H} \sum_{Q_h \in \mathbb{P}_h} |x_{Q_h}|^2.$$

Corollary 5. *The learning schemes described in Theorems 3 and 4 with re-scaled required precision $\mathcal{O}(\epsilon n^{-r})$ provides an estimator $\hat{\mu}$ of the unknown measure μ whose corresponding Pauli channel $\hat{\mathcal{P}}$, defined by $\hat{\mathcal{P}}(\rho) := \sum_{P \in \mathbb{P}_n} \hat{\mu}(P) P \rho P$, satisfies*

$$\|\mathcal{P} - \hat{\mathcal{P}}\|_\diamond \leq \epsilon.$$

In particular, this strategy provides a classical description of the channel $\hat{\mathcal{P}}$ with access to $\tilde{\mathcal{O}}(n^r \epsilon^{-2})$ uses of the channel \mathcal{P} and a run-time of $\tilde{\mathcal{O}}(n^{r(2+D^{r+1})} \epsilon^{-2})$.

Proof. We first recall that [13]:

$$\begin{aligned} \|\mathcal{P} - \hat{\mathcal{P}}\|_\diamond^2 &= (2 \|\mu - \hat{\mu}\|_{\text{TV}})^2 \\ &\leq D(\mu \|\hat{\mu}) + D(\hat{\mu} \|\mu), \end{aligned}$$

where in the second line, we use Pinsker's inequality which relates the total variation distance between two probability measures to their relative entropies:

$$D(\mu_1 \|\mu_2) := \sum_{P \in \mathbb{P}_n} \mu_1(P) (\log(\mu_1(P)) - \log(\mu_2(P))).$$

By using the definitions for $\hat{\mu}$ and μ , we further can show that

$$\begin{aligned}
D(\mu\|\hat{\mu}) + D(\hat{\mu}\|\mu) &= \sum_{|h|\leq r} \sum_{Q_h \in \mathbb{P}_h} (e_{Q_h} - \hat{e}_{Q_h})(\lambda_{Q_h} - \hat{\lambda}_{Q_h}) \\
&\leq \|e - \hat{e}\|_{\ell_2} \|\lambda - \hat{\lambda}\|_{\ell_2} \\
&\leq \Gamma \|\lambda - \hat{\lambda}\|_{\ell_2}^2 \\
&= \Gamma \sum_{|h|\leq r} \sum_{P_h \in \mathbb{P}_h} |\theta^h(P_h) - \hat{\theta}^h(P_h)|^2 \\
&= \mathcal{O}(n^r) \max_h \|\theta^h - \hat{\theta}^h\|_{\infty}^2.
\end{aligned}$$

The result follows after re-scaling ϵ to ϵn^{-r} in Theorems 3 and 4 in order to get the required precision. \square

Acknowledgements. CR acknowledges financial support from the ANR project QTraj (ANR-20-CE40-0024-01) of the French National Research Agency (ANR), as well as from the Humboldt Foundation. This work is part of HQI initiative (www.hqi.fr) and is supported by France 2030 under the French National Research Agency award number ‘‘ANR-22-PNCQ-0002’’. We thank Matthias Caro, Steve Flammia and Aadil Oufkir for fruitful discussions and comments on the draft.

* Cambyse Rouz e cambyse.rouze@tum.de

† Daniel Stilck Fran a daniel.stilck_franca@ens-lyon.fr

- [1] O. Fawzi, A. Oufkir, and D. S. Fran a, Lower bounds on learning pauli channels (2023), arXiv:2301.09192 [quant-ph].
- [2] S. T. Flammia and J. J. Wallman, ACM Transactions on Quantum Computing **1**, 1 (2020).
- [3] S. Chen, S. Zhou, A. Seif, and L. Jiang, Physical Review A **105**, 032435 (2022).
- [4] S. T. Flammia and R. O’Donnell, Quantum **5**, 549 (2021).
- [5] T. Wagner, H. Kampermann, D. Bru , and M. Kliesch, Phys. Rev. Lett. **130**, 200601 (2023).
- [6] J. J. Wallman and J. Emerson, Phys. Rev. A **94**, 052325 (2016).
- [7] R. Harper, S. T. Flammia, and J. J. Wallman, Nature Physics **16**, 1184 (2020).
- [8] R. Harper and S. T. Flammia, Learning correlated noise in a 39-qubit quantum processor (2023), arXiv:2303.00780 [quant-ph].
- [9] G. Bresler, in *Proceedings of the forty-seventh annual ACM symposium on Theory of computing* (2015) pp. 771–782.
- [10] D. S. Fran a and A. Hashagen, Journal of Physics A: Mathematical and Theoretical **51**, 395302 (2018).
- [11] J. Helsen, I. Roth, E. Onorati, A. H. Werner, and J. Eisert, PRX Quantum **3**, 020357 (2022).
- [12] J. Helsen, X. Xue, L. M. Vandersypen, and S. Wehner, npj Quantum Information **5**, 71 (2019).
- [13] E. Magesan, J. M. Gambetta, and J. Emerson, Physical Review A **85**, 042311 (2012).

- [14] M. Heinrich, M. Kliesch, and I. Roth, Randomized benchmarking with random quantum circuits (2023), arXiv:2212.06181 [quant-ph].
- [15] D. S. Fran a, L. A. Markovich, V. V. Dobrovitski, A. H. Werner, and J. Borregaard, Efficient and robust estimation of many-qubit hamiltonians (2022), arXiv:2205.09567 [quant-ph].
- [16] J. Kunjummen, M. C. Tran, D. Carney, and J. M. Taylor, Physical Review A **107**, 042403 (2023).
- [17] R. Levy, D. Luo, and B. K. Clark, Classical shadows for quantum process tomography on near-term quantum computers (2021), arXiv:2110.02965 [quant-ph].
- [18] J. Watrous, *The theory of quantum information* (Cambridge university press, 2018).
- [19] J. Besag, Journal of the Royal Statistical Society: Series B (Methodological) **36**, 192 (1974).
- [20] L. Hamilton, F. Koehler, and A. Moitra, Advances in Neural Information Processing Systems **30** (2017).
- [21] A. Anshu, S. Arunachalam, T. Kuwahara, and M. Soleimanifar, Efficient learning of commuting Hamiltonians on lattices, unpublished notes available at Anurag Anshu’s website, (link to note).
- [22] S. T. Flammia, Averaged circuit eigenvalue sampling (2021), arXiv:2108.05803 [quant-ph].
- [23] A. Montanari, Electronic Journal of Statistics **9**, 10.1214/15-ejs1059 (2015).
- [24] M. Vuffray, S. Misra, A. Lokhov, and M. Chertkov, Advances in neural information processing systems **29** (2016).
- [25] L. Devroye, M. Lerasle, G. Lugosi, and R. I. Oliveira, The Annals of Statistics **44**, 10.1214/16-aos1440 (2016).
- [26] R. Harper, I. Hincks, C. Ferrie, S. T. Flammia, and J. J. Wallman, Phys. Rev. A **99**, 052350 (2019).

SPAM robust parallel estimation of Pauli eigenvalues

Our protocol requires us to estimate the eigenvalues of the Pauli channel for Paulis with small support, as these correspond to low-weight Fourier coefficients. The estimation of eigenvalues of a Pauli channel has been considered in several works in the literature [2, 4, 22]. In particular, using tools from randomized benchmarking techniques, it is possible to estimate these coefficients in a way that is robust against both state preparation and measurement errors (SPAM). This is a highly desirable feature for process tomography protocols since in current devices these errors are not negligible. On the other hand, recent schemes inspired by the classical shadows framework for learning states [15–17] showed how to infer various parameters of a quantum channel from a single class of randomized experiments. Indeed, such schemes are able to estimate m Pauli diagonal elements of the form

$$2^{-n} \text{tr}(\mathcal{P}\mathcal{P}(P)) \quad (11)$$

for Paulis of weight at most w up to an error ϵ with probability of success at least $1 - \delta$ from $\mathcal{O}(9^w \epsilon^{-2} \log(m\delta^{-1}))$




FIG. 2. Class of circuits we use for learning the Pauli eigenvalues on three qubits in the absence of SPAM. Here, C_i are independent uniformly distributed random one-qubit Clifford unitaries, P_i, Q_i are uniform independent random Pauli matrices.

randomized experiments. Thus, we see that this protocol is highly parallelizable. However, it is not as robust to SPAM errors as randomized benchmarking. We will now show how to obtain the best of the two worlds.

We will denote by Φ^* the Hilbert Schmidt dual (i.e. evolution in Heisenberg picture) to the map Φ , and by \mathcal{Q} the action of unitary conjugation by Q , i.e. $\mathcal{Q}(\rho) = Q\rho Q^\dagger$.

The class of circuits we consider are illustrated in Fig. 2. We first prepare the system in the $|0^n\rangle$ state, and proceed to apply one layer of one-qubit random Cliffords and one layer of random Pauli matrices. After that, we apply the target channel k times. This is followed by another layer of random Paulis and 1-qubit Cliffords, which are drawn independently from the previous ones.

The post-processing of this data will be inspired by both the shadow process of [15] and the character randomized benchmarking of [10, 12]. For a target Pauli string P , we define the function $\chi_P : \mathbb{P}_n \rightarrow \{-1, 1\}$ given by

$$\chi_P(Q) = (-1)^{P \cdot Q}. \quad (12)$$

We also define the P -Twirling \mathcal{T}_P as

$$\mathcal{T}_P(X) := \frac{1}{|\mathbb{P}_n|} \sum_{Q \in \mathbb{P}_n} \chi_P(Q) Q X Q. \quad (13)$$

It follows from standard results of representation theory that $\mathcal{T}_P(X) = 2^{-n} \text{tr}(PX)P$. Furthermore, given a Pauli string P , an n qubit Clifford unitary channel $\mathcal{C} = \otimes_{i=1}^n C_i$ and bitstring $x \in \{0, 1\}^n$ we define the function ω_P as

$$\omega_P(\mathcal{C}, x) := 3^{w(P)} \langle x | C^\dagger P C | x \rangle \langle 0^n | C^\dagger P C | 0^n \rangle, \quad (14)$$

Given a sample generated from the circuit as in Fig. 2 where we used the channel k times with the corresponding random Cliffords and Paulis and a reference Pauli string P , we consider the random variable

$$\Omega_k := \omega_P(\mathcal{C}, X_k) \chi_P(Q^{\text{in}}) \chi_P(Q^{\text{out}}), \quad (15)$$




FIG. 3. Structure of our noisy circuits for 3 qubits in the presence of SPAM.

where X_k corresponds to the output to the circuit in Figure 2. Let $\alpha_P^k := 2^{-n} \text{tr}(P\mathcal{P}^k(P))$ and $\alpha_P = \alpha_P^1$. We will now show that $\mathbb{E}[\Omega_k] = C\alpha_P^k$ for some constant C that only depends on the amount of SPAM errors. Furthermore, we have that the variance of Ω_k is bounded as $\mathbb{V}(\Omega_k) \leq 3^{w(P)}$. Thus, by using a median of means estimator, we can ensure that we obtain estimates of $C\alpha_P^k$ for different choices of P . By then fitting the curve $k \mapsto C\alpha_P^k$ for various values of k , as we will describe in more detail below, it is possible to obtain an estimate of α_P that is robust to SPAM errors.

Before we proceed to prove the claims made in the previous paragraph, let us first make some assumptions about how the noise acts. We will model the SPAM noise and the noise incurred by the random Paulis by assuming that between those two steps, there are channels Φ_1, Φ_2 acting on the systems. This is illustrated in Fig. 3. Note that we assume that the noise is independent of the gates and independent of k . The latter assumption is justified by the fact that the circuit we need to implement for the measurement and preparations is independent of k . The former is a standard assumption in RB, but recent works [11] have relaxed it to weakly dependent noise. Generalizing it to this setting should be possible here as well, but would go beyond the scope of this work. We now have:

Proposition 6. *Let Ω_k be the random variable defined in Equation 15, where we draw $\mathcal{C}, Q^{\text{in}}$ and Q^{out} uniformly and independently at random. Then*

$$\mathbb{E}[\Omega_k] = C_P \alpha_P^k \quad (16)$$

and

$$\mathbb{V}(\Omega_k) \leq 3^{w(P)}, \quad (17)$$

where C_P is given by

$$C_P := 2^{-2n} \text{tr}(\Phi_1^*(P)P) \text{tr}(\Phi_2^*(P)P). \quad (18)$$

Proof. We will first fix the values of Q^{in} and Q^{out} and consider only the expectation values w.r.t. \mathcal{C} . Note that $\omega_P(\mathcal{C}, X_k)$ is only nonzero if \mathcal{C} maps the computational basis to the eigenbasis of P , a fact we will use repeatedly. As the local Cliffords are uniformly random, this

happens with probability $3^{-w(P)}$. Let P_+, P_- be the projectors onto the positive and negative eigenvalues of P . Again by the uniformity of the Cliffords, conditioned on the fact that we mapped to the eigenbasis of P under the Cliffords, we prepare the initial state $\frac{P_+}{2^{n-1}}$ with probability $1/2$ and $\frac{P_-}{2^{n-1}}$ with probability $1/2$. We then evolve this state by $\Phi_2 Q^{\text{out}} \mathcal{P}^k Q^{\text{in}} \Phi_1$ and measure the POVM $\{P_+, P_-\}$. By the definition of the random variable $\omega_P \equiv \omega_P(\mathcal{C}, X_k)$ we have that its expectation value conditioned on \mathcal{C} mapping onto the eigenbasis of P is given by

$$\begin{aligned} \mathbb{E}[\omega_P | \mathcal{C}] &= 3^{w(P)} 2^{-n} \left(\text{tr} (P_+ \Phi_2 Q^{\text{out}} \mathcal{P}^k Q^{\text{in}} \Phi_1 (P_+)) \right. \\ &\quad + \text{tr} (P_- \Phi_2 Q^{\text{out}} \mathcal{P}^k Q^{\text{in}} \Phi_1 (P_-)) \\ &\quad - \text{tr} (P_- \Phi_2 Q^{\text{out}} \mathcal{P}^k Q^{\text{in}} \Phi_1 (P_+)) \\ &\quad \left. - \text{tr} (P_+ \Phi_2 Q^{\text{out}} \mathcal{P}^k Q^{\text{in}} \Phi_1 (P_-)) \right) \\ &= 3^{w(P)} 2^{-n} \text{tr} (P \Phi_2 Q^{\text{out}} \mathcal{P}^k Q^{\text{in}} \Phi_1 (P)) \end{aligned}$$

As the probability of hitting the right basis is $3^{-w(P)}$, it will cancel out the $3^{w(P)}$. Let us now analyze the effect of the additional Pauli twirling. By definition of the Twirling (see Eq. (13)), if we take the expectation value over Q^{in} and Q^{out} now, we obtain:

$$\mathbb{E}[\Omega_k] = 2^{-n} \text{tr} (\mathcal{T}_P^*(\Phi_2^*(P)) \mathcal{P}^k \circ \mathcal{T}_P(\Phi_1(P))). \quad (19)$$

Note that $\mathcal{T}_P^* = \mathcal{T}_P$ and, thus

$$\mathcal{T}_P^*(\Phi_2^*(P)) = 2^{-n} \text{tr} (\Phi_2^*(P) P) P \quad (20)$$

and similarly

$$\mathcal{T}_P(\Phi_1(P)) = 2^{-n} \text{tr} (\Phi_1(P) P) P = 2^{-n} \text{tr} (\Phi_1^*(P) P) P, \quad (21)$$

which proves our claim about the expected value. The estimate of the variance follows again from noting that the random variable Ω_k is nonzero only if \mathcal{C} maps the computational basis to an eigenbasis of P , which happens with probability $3^{-w(P)}$. In that case, the value of

the random variable squared is bounded by $9^{w(P)}$, which gives that the expected value of Ω_k^2 is bounded by $3^{w(P)}$. Our bound on the variance follows. \square

Some remarks are in order: first, note that as P is an eigenvector of \mathcal{P} , Q^{in} has no effect in the protocol. However, this extra step ensures that small deviations from \mathcal{P} being a Pauli channel will be suppressed. Furthermore, it should be straightforward to obtain a similar protocol for the case where input and output Pauli matrices are not the same at the expense of a higher sample complexity in terms of the locality.

The following corollary follows directly from the joint use of Proposition 6 and a median of means argument:

Corollary 7. *In the same setting as Prop. 6, let P_1, \dots, P_m be m Pauli strings such that $w(P_i) \leq w$, denote by p_0 the probability that no error occurs, i.e. $p_0 = \mu(\mathbb{I})$, and*

$$2^{-2n} |\text{tr} (\Phi_1^*(P_i) P_i) \text{tr} (\Phi_2^*(P_i) P_i)| \geq C_{\text{SPAM}}. \quad (22)$$

Then $\mathcal{O}(3^w \epsilon^{-2} C_{\text{SPAM}}^{-2} \log(m\delta^{-1}))$ runs of circuits of the form of Fig. 3 suffice to estimate all α_{P_i} up to an error $\epsilon(1-p_0)$ with probability of success at least $1-\delta$. Furthermore, the maximal value of k needed in each run satisfies $k = \mathcal{O}(\log(1/(1-p_0)))$.

Proof. From Prop. 6 it follows that using a median of means estimator [25] we can estimate the mean of all random variables with expectation value $C_{P_i} \alpha_{P_i}^k$ for a given k up to $C_{\text{SPAM}} \epsilon$ and probability of success at least $1-\delta$ with the advertised sample complexity. We can then use the results of [26] to fit these values to an exponential curve and obtain an estimate that satisfies the error bound $\epsilon(1-p_0)$ by using the channel $k = \mathcal{O}(-\log(1-p_0))$ times. \square

Thus, by combining the randomized experiments with an exponential fitting, it is possible to recover all Pauli coefficients in a SPAM-robust way. Note that C_{SPAM} will usually also scale exponentially in the locality of the Pauli string for local noise models. E.g., for local depolarizing noise with depolarizing probability q , $C_{\text{SPAM}} \geq (1-q)^{2w}$.