

The battle of clean and dirty qubits in the era of partial error correction

Daniel Bultrini^{*1,2}, Samson Wang^{*1,3}, Piotr Czarnik^{1,4}, Max Hunter Gordon^{1,5}, M. Cerezo^{6,7}, Patrick J. Coles^{1,7}, and Lukasz Cincio^{1,7}

¹Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

²Theoretische Chemie, Physikalisch-Chemisches Institut, Universität Heidelberg, INF 229, D-69120 Heidelberg, Germany

³Imperial College London, London, UK

⁴Institute of Theoretical Physics, Jagiellonian University, Krakow, Poland.

⁵Instituto de Física Teórica, UAM/CSIC, Universidad Autónoma de Madrid, Madrid 28049, Spain

⁶Information Sciences, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

⁷Quantum Science Center, Oak Ridge, TN 37931, USA

When error correction becomes possible it will be necessary to dedicate a large number of physical qubits to each logical qubit. Error correction allows for deeper circuits to be run, but each additional physical qubit can potentially contribute an exponential increase in computational space, so there is a trade-off between using qubits for error correction or using them as noisy qubits. In this work, we look at the effects of using noisy qubits in conjunction with noiseless qubits (an idealized model for error-corrected qubits), which we call the “clean and dirty” setup. We employ analytical models and numerical simulations to characterize this setup. Numerically we show the appearance of Noise-Induced Barren Plateaus (NIBPs), i.e., an exponential concentration of observables caused by noise, in an Ising model Hamiltonian variational ansatz circuit. We observe this even if only a single qubit is noisy and given a deep enough circuit, suggesting that NIBPs cannot be fully overcome simply by error-correcting a subset of the qubits. On the positive side, we find that for every noiseless qubit in the circuit, there is an exponential suppression in the concentration of gradient observables, showing the benefit of partial error correction. Finally, our analytical models corroborate these findings by showing that observables concentrate with a scaling in the exponent related to the ratio of dirty-to-total qubits.

^{*}These authors contributed equally to this work.

1 Introduction

Since Feynman’s seminal proposal [1], quantum computing has moved from a purely theoretical exercise to real devices. Yet, quantum computers today are still far from an ideal machine able to truly harness the full power of fault-tolerant computing. Current devices do not have the qubit counts or gate fidelities required to implement large-scale error correction. This means that only small-scale implementations of error-correcting codes are currently feasible [2]. Therefore, implementing Shor’s factoring algorithm [3], the Harrow-Hassidim-Lloyd algorithm [4], and the myriad of possibilities this computational paradigm could bring are still a distant reality.

An important primitive in quantum computing is estimating the expectation values of some operator at the end of a quantum circuit. This appears in the Noisy Intermediate Scale Quantum (NISQ) [5] setting for the case of variational quantum algorithms [6, 7] and quantum machine learning [8], but also in fault-tolerant algorithms [9]. However, it has been shown that noise can have detrimental effects resulting in the exponential concentration of expectation values [10–13]. In the specific case of variational quantum algorithms and quantum machine learning, this implies untrainability of the models, i.e., exponential scaling of resources required for training [13–18]. This scaling phenomenon is known as a Noise-Induced Barren Plateau (NIBP) [13], which is a specific type of barren plateau [14–22]. In addition, it was also shown that error mitigation cannot reverse the effect of noise on expectation value concentration [23]. Therefore, noise presents one of the largest obstacles to the practical utility of quantum computers and to obtaining near-term quantum advantage. This leads to a critical question that we aim to address in this work: *how can we mitigate, or remove, the effect of NIBPs and exponential concentration due to noise?*

To investigate this question we construct a quantum computational model that employs “clean” and “dirty” qubits. Here, we define clean qubits as qubits on which no noise channels are applied, which we envision as the best-case implementation of an error-correcting code. On the other hand, the dirty qubits in our setup are affected by noise. This setup presents a novel, physically motivated paradigm where we can explore the scalability of cost concentration in quantum algorithms. We can motivate this setup as an idealized scenario of the regime where the number of qubits becomes sufficiently large to implement error correction on some, but not all, of the qubits in a quantum computer. The intersection of NISQ and quantum error correction has been considered in other recent works [24–26]. We note our setup is different from earlier clean qubit models [27–29], which work on a system where one or more qubits are in a known state, such as $|0\rangle$, and the rest are in the maximally mixed state².

We find that using clean qubits can mitigate the effect of NIBPs. Specifically, we find that using clean qubits increases the depth accessible to an ansatz before the magnitude of the gradients vanishes. However, NIBPs are not avoided unless all qubits are clean. We analytically explore the effect of clean qubits on the concentration of expectation values and gradient using two toy models. The first is a CNOT ladder acting on an input state followed by depolarizing noise on the dirty qubits. The second explores two initially disjoint subsystems, one or both of which are noisy, that are subsequently connected via CNOTs. Both show that the expectation value concentrates with a scaling in the exponent related to the ratio of dirty to total qubits.

Our numerics arrive at essentially the same conclusions as our analytics. Here, we numerically explore the magnitude of the gradient when optimizing the Hamiltonian Variational Ansatz (HVA) for systems of 4, 6, and 8 qubits under a depolarizing noise model and also a more realistic model of a trapped ion quantum computer. This allows us to investigate how the gradient scales with system size under the effect of noise. We show that using clean qubits leads to gradient scaling similar to the case of reducing error rates of the noise on every qubit. This in turn means that deeper circuits can be employed before NIBPs become an issue. Therefore, our work presents a possible avenue to improve the reach of quantum computing using a combination of error-corrected and dirty qubits.

²Such models have the ability to factor numbers [30], compute partition functions [31] and are hard to simulate classically [32].




Figure 1: **The clean and dirty setup.** A schematic example of the clean and dirty setup. We have n_c noiseless, clean qubits at which no errors happen. For these qubits, we represent the action of a gate by its unitary U . We also have n_d uncorrected, dirty qubits. The action of a gate at these qubits is represented by U followed by a noise channel $\mathcal{N}$. In the case of gates acting at both the dirty and the clean qubits, we represent them by U followed by a noise channel $\mathcal{N}$ acting only at the dirty qubits.

2 The Clean and Dirty Computational Model

In this work, we investigate how the emergence of exponential concentration and NIBPs is affected by introducing clean, noiseless qubits in a quantum computer. Such a setup may provide insights into an early era of quantum error correction when, due to limited numbers of physical qubits available, the number of error-corrected, high-quality logical qubits will be severely restricted [33–35]. In such a case, a question arises as to whether combining a limited number of high-quality logical qubits with uncorrected, noisy qubits could mitigate or perhaps even prevent the exponential concentration of expectation values and gradients. If so, this would highlight the power of post-NISQ quantum computing architectures, relative to their NISQ counterparts.

To provide insight into these questions, we consider an idealized model of such a quantum computer, with n qubits, partitioned into n_c noiseless, clean qubits, and n_d noisy, dirty qubits. We assume that the action of a gate g , defined by a unitary U , on the clean qubits of a state ρ is

$$\rho \rightarrow U\rho U^\dagger. \quad (1)$$

When g acts at the dirty qubits we represent its action by U followed by a noise channel $\mathcal{N}$ acting only at the dirty qubits:

$$\rho \rightarrow \mathcal{N}(U\rho U^\dagger). \quad (2)$$

For a gate g acting between a dirty and clean qubit, we assume no noise acts on the clean qubit. We call this setup the clean and dirty setup, and we illustrate

it schematically in Figure 1. This simple setup can be used to understand the fundamental limitations of a quantum computation involving both high-quality logical, and low-quality, either uncorrected or logical qubits. For $n_d = 0$ the setup is equivalent to a perfect, noiseless quantum computation, and for $n_c = 0$ it becomes equivalent to a noisy quantum device.

We note that this idealized setup neglects error rates of real-world logical qubits due to imperfections of quantum error correction. Furthermore, the model does not account for the larger number of native hardware gates necessary to implement logical non-Clifford gates compared to the Clifford ones [36]. This overhead may result in a higher effective error rate for logical non-Cliffords that the clean-and-dirty setup does not consider. Therefore, we treat this work as an exploration of the potential of quantum algorithms combining qubits with large and small effective error rates rather than a realistic model of future quantum devices.

An alternative approach for early implementations of quantum error correction is to use less robust quantum error correction codes with lower distances to detect and correct a part of the errors. An example of this approach is presented in Ref. [37]. In general, the error rates of logical qubits are expected to decay exponentially with the code distance, with the base of the exponential being proportional to the ratio of the physical qubits' error rate to its threshold value [35]. Therefore, at the beginning of the quantum error correction era, when the physical error rates will be close to the threshold, obtaining high-quality logical qubits would require a large code distance and consequently a large number of physical qubits per logical qubit [35]. Consequently, implementing less powerful codes would require much fewer physical qubits making them more suitable for the first realizations of quantum error correction.

Nevertheless, to obtain a physical qubit count large enough for a quantum advantage, it might be necessary to supplement such logical qubits with noisy, physical ones. Furthermore, a partition of the device qubits into lower-quality error-corrected qubits and noisy ones would increase the number of computational qubits whilst potentially retaining some advantages of quantum error correction. Therefore, a combination of the noisy and quantum-error-corrected qubits modeled by the clean and dirty setup might be the optimal solution when a limited number of device qubits is the primary limitation on the computer's computational power.

While this motivation applies to the first generations of quantum computers utilizing quantum error correction, the clean and dirty setup can also provide insights into later, more advanced architectures combining high and lower-quality logical qubits. Such a

combination might be desirable to optimally utilize qubit counts of quantum hardware when the number of qubits is insufficient to use high-quality error corrections for all logical qubits. Furthermore, in the case of algorithms having qubits with particularly many gates acting upon them, much more errors will occur at these qubits. In such a case, using a code with a higher distance for these qubits might be a natural choice. An example of such an algorithm is quantum phase estimation with a large unitary controlled by an ancilla qubit [38]. Therefore, we expect insights gained with the clean and dirty model to potentially remain relevant beyond the early era of quantum error correction.

3 Analytical Investigation

3.1 Concentration of expectation values

In this section, we present two analytical models in which exponential scaling effects due to noise are observed, even though a proportion of the registers are noise-free. Such scaling effects are present even if there is only one noisy register. The two settings demonstrate two different mechanisms in which noise can spread through a system via entangling gates. In the first setting, we consider a simple circuit where clean and dirty qubits interact via cycles of CNOT gates. Here, we find that the CNOTs allow entropy to spread from noisy to clean registers despite the fact that if the initial state is in the computational basis, no entanglement is actually created between the subsystems. In the second case, we consider a setting where two subsystems, which can in principle be initially disconnected, undergo an entangled measurement. We show that, due to the entangled measurement, local noise in the dirty qubits leads to decoherence and information loss in the joint system of all clean and dirty qubits. In both models, we find that the output distribution of the circuit, as measured in the computational basis, is exponentially indistinguishable from the uniform distribution with increasing circuit depth. However, in both settings, this concentration effect is dampened with increasing number of clean qubits, compared to the number of dirty qubits.

In the first setting, we consider the interaction of local noise with CNOT ladders, as shown in Figure 2. Ladders of 2-qubit gates can be found as a primitive quantum machine learning settings (e.g. in [39]) as well as in variational quantum algorithms settings (e.g. for the unitary coupled cluster ansatz when one implements an exponentiation of a multi-qubit Pauli operator [40]). In order to understand how such ladders interact with noise, we consider a simplified model where layers of CNOT ladders are interleaved




Figure 2: **CNOT ladders.** In our first analytical example we consider an input state ρ acted on by L layers of CNOT ladders. After each layer, we consider an instance of local depolarizing noise $\mathcal{D}_p$ with depolarizing probability p acting on each of the first n_d registers. We compare the output distribution as measured in the computational basis to the uniform distribution.

with instances of local depolarizing noise, where the local noise only acts on the first n_d qubits. In the following proposition, proved in Appendix A, we show how the output distribution converges with circuit depth.

Proposition 1. *Consider the circuit in Figure 2 with L layers, n_d dirty qubits each with depolarizing noise with depolarizing probability p in each layer, and computational basis measurement $O_{z_n} = |z_n\rangle\langle z_n|$; $z_n \in \{0,1\}^n$. Denote the channel that describes the action of the CNOT ladders and noise in the circuit as $\mathcal{W}_L^{(n_d)}$. For any input state ρ , the expectation value concentrates as*

$$\text{Tr}[\mathcal{W}_L^{(n_d)}(\rho)O_{z_n}] - \frac{1}{2^n} \leq (1-p)^{\Gamma_{n,n_d,L}} \sqrt{P(\rho)}, \quad (3)$$

for any computational basis measurement, where $P(\rho) = \text{Tr}[\rho^2]$ denotes the purity of ρ , and we have defined

$$\Gamma_{n,n_d,L} = \begin{cases} \left\lfloor L \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \right\rfloor, & \text{for } n_d = 1, \\ \left\lfloor L \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \right\rfloor, & \text{for } n_d > 1. \end{cases} \quad (4)$$

Proposition 1 shows that the output distribution of the CNOT ladder circuit exponentially concentrates in the number of layers of the circuit. Moreover, compared to the setting where all registers are noisy, the exponent is scaled by a factor approximately equal to n_d/n . Thus, we can consider the scaling to depend on an *effective* noisy depth $L_{\text{eff}} := Ln_d/n$. We note that the bound in Eq. (3) holds even if there is no entanglement in the circuit - that is, it holds even if the input state ρ is a tensor-product state diagonal in the computational basis.

An alternative interpretation of the clean and dirty computational model follows by inspecting Proposition 1 in the low error limit. Namely, considering for simplicity n to be a power of 2, then we have $(1-p)^{\Gamma_{n,n_d,L}} = (1-p)^{L_{\text{eff}}}$ and hence, if $p \ll 1$, then

$$(1-p)^{L_{\text{eff}}} \sim (1-p \frac{n_d}{n} L) \sim (1-p \frac{n_d}{n})^L, \quad (5)$$

up to first order in p . Equation (5) shows that adding clean qubits (left-hand side) is also approximately equivalent to a fully noisy system with an effective noise parameter rescaled by n_d/n . The identification of the right-hand-side can be made precise from Proposition 1 as $(1-p \frac{n_d}{n})^L$ is precisely the scaling coefficient one finds in an L -layered system where all qubits are subject to noise with a probability $p \frac{n_d}{n}$. As such, having fewer dirty qubits has a similar effect to reducing the effective noise in a system where all qubits undergo the same rate of depolarizing noise. This can be thought of as arising due to the fact that the CNOTs propagate the noise from dirty qubits through to the rest of the system.

Thus, for sufficiently deep circuits, the localized contributions of error instances wash into the rest of the circuit, and the decoherence can be broadly characterized by the *total number of error instances*, or equivalently effective error rate, and depth L_{eff} . This interpretation is explored more in our numerical studies in Section 4. We note that our first setting can also be adapted to give similar results for Pauli observables. Specifically, by considering the reverse circuit with Pauli observables we observe the same exponential concentration. We discuss this further in Appendix A.3.1.

In the second setting, we show how local noise impacts the global system due to entangled measurements. This can also equally be considered to demonstrate the role of entanglement when considering the Heisenberg picture (i.e., considering the circuit acting in the reverse order). We consider the circuit in Figure 3(a) and suppose that the registers are split into two subsystems A and B . Both subsystems are acted on by a global unitary W (which can in principle even be separable across the cut $A|B$), before they undergo evolution according to some channel $\mathcal{T}$ which is also separable across the cut $A|B$. One choice for $\mathcal{T}$ is to consider noise-free unitary evolution which we denote as $\mathcal{U} \otimes \mathcal{V}$. In order to characterize the effects of noise, we can also modify these channels to account for hardware noise, as demonstrated in Figure 3(b), in the sense that between each non-parallelizable layer of gates there is an instance of local depolarizing noise acting on each of the noisy qubits. We denote such a noisy modification as $\mathcal{U} \rightarrow \tilde{\mathcal{U}}$ or $\mathcal{V} \rightarrow \tilde{\mathcal{V}}$. Further, we denote the number of such unitary layers as $L_{\mathcal{U}}$ and $L_{\mathcal{V}}$ respectively. Finally, at the end of the circuit there is an entangled measurement across m pairs of registers in the two subsystems, which we denote with the




Figure 3: Entangled registers (a) In our second analytical example we show that entanglement can also cause local noise to spread globally. In this circuit, we consider an arbitrary input state ρ that is acted on by some unitary W . We then partition the system into two subsystems, each of which undergoes its own unitary evolution, denoted as $\mathcal{U}$ and $\mathcal{V}$. In our analytics, we consider a modification where one of the subsystems undergoes noisy evolution whilst the other is noise-free. At the end of the circuit, a subset of the qubits undergoes an entangled measurement across the two subsystems. (b) We define the depth of $\mathcal{U}$ to be the number of non-parallelizable unitary layers $L_{\mathcal{U}}$ in a given hardware implementation. When we consider a noisy modification $\mathcal{U} \rightarrow \tilde{\mathcal{U}}$, we insert layers of local depolarizing noise with depolarizing probability p in between each unitary layer. We define $L_{\mathcal{V}}$ and $\tilde{\mathcal{V}}$ in an analogous manner.

measurement operator $\mathcal{B}_m(|z_{2m}\rangle\langle z_{2m}|)$, where $\mathcal{B}_m$ is the map corresponding to a Hadamard gate followed by a CNOT on m pairs of qubits. We note that a special case of this circuit is the (global) Hilbert-Schmidt test circuit [41]. We find that the following proposition, proved in Appendix A, holds.

Proposition 2. *Consider the circuit in Figure 3(a) partitioned into two subsystems A and B . We have entangling operation $\mathcal{B}_m$ on $m \geq 1$ pairs of qubits across A and B , noisy unitary evolution $\mathcal{T}$ separable across the cut $A|B$, global unitary evolution $\mathcal{W}$, and computational basis measurement $O_{z_n} = |z_n\rangle\langle z_n|$; $z_n \in \{0,1\}^n$. For any input state ρ and any computational basis measurement, the expectation value concentrates as*

$$\text{Tr}[\mathcal{B}_m^\dagger \circ \mathcal{T} \circ \mathcal{W}(\rho) O_{z_n}] - \frac{1}{2^n} \leq (1-p)^{L_{\mathcal{T}}} \sqrt{P(\rho)}, \quad (6)$$

where $P(\rho)$ is the purity of ρ and we have defined

$$L_{\mathcal{T}} = \begin{cases} L_{\mathcal{U}}, & \text{for } \mathcal{T} = \tilde{\mathcal{U}} \otimes \mathcal{V}, \\ L_{\mathcal{V}}, & \text{for } \mathcal{T} = \mathcal{U} \otimes \tilde{\mathcal{V}}, \\ L_{\mathcal{U}} + L_{\mathcal{V}}, & \text{for } \mathcal{T} = \tilde{\mathcal{U}} \otimes \tilde{\mathcal{V}}. \end{cases} \quad (7)$$

Proposition 2 shows the role entanglement (or entangled measurements) can play in spreading noise

from a noisy subsystem to a clean subsystem. Namely, despite noise only occurring in one part of the circuit, the output distribution of the full system exponentially concentrates in the circuit depth. We also note that this occurs even if only one clean qubit is coupled to a dirty qubit ($m = 1$). The proof follows from analysis in the Heisenberg picture and thus implies that entangling registers at the beginning of the computation also leads to an identical effect, which we demonstrate in Appendix A.4.

We note that, when only subsystem A (B) is noisy, there is still exponential scaling, however, compared to the fully noisy setting the exponent is modified by factor $\frac{L_{\mathcal{U}}}{L_{\mathcal{U}}+L_{\mathcal{V}}}$ ($\frac{L_{\mathcal{V}}}{L_{\mathcal{U}}+L_{\mathcal{V}}}$). If we consider the unitaries to have linear depth in the number of qubits and scale proportionally with the same factor (i.e. $L_{\mathcal{U}}/n_A = L_{\mathcal{V}}/n_B$), then the modification to the exponent is n_A/n (n_B/n). That is, the exponent is rescaled by the ratio of the number of dirty qubits to the total number of qubits, which is similar to the effect we find in Proposition 1.

Propositions 1 and 2 together show that despite some of the registers being clean, exponential concentration of expectation values in the circuit depth can still occur, in a similar spirit to that of NIBPs as presented in Ref. [13]. Indeed, when we set $n_d = n$ we recover the same scaling. Moreover, there is also exponential concentration with each increasing dirty qubit (for linear depth circuits, in the case of Proposition 2). We note that following the methods presented in Ref. [13], our results can also be generalized to classes of Pauli noise that have at least two types of Pauli error occurring with non-zero probability.

We remark that, despite the fact that our bounds are one-sided, there is some physical significance to the factor of n_d/n that the exponent is rescaled to when transitioning from n to n_d dirty qubits. This is due to the fact that the exponent in our bounds corresponds to the Pauli string with the smallest amplitude decay. Thus, whilst we would not expect all circuits to exactly exhibit this scaling, we do expect it to be characteristic of circuits in the best case. Further details are given in the proofs of Propositions 1 and 2, which can be found in Appendices A.3 and A.4 respectively.

3.2 Gradient scaling

Our above results can also be readily transported to consider gradient scaling. First, we note that if either of the two above settings is modified to contain a trainable parameter (without breaking initial assumptions) and in settings amenable to the parameter shift rule, then Propositions 1 and 2 automatically imply exponentially vanishing gradients in the depth of the circuit. Namely, this implies barren plateaus for lin-

ear depth circuits if n_d scales linearly with n . This also implies barren plateaus for polynomial depth circuits of degree 2 or greater, for any $n_d > 0$. Further, in more general settings where one cannot apply the parameter shift rule, we now explicitly demonstrate how gradient scaling results can be obtained for modified circuits, and present further details in Appendix A.5.

We consider a trainable unitary of the form

$$Y(\boldsymbol{\theta}) = W_0 \prod_{k=1}^K e^{-i\theta_k H_k} W_k, \quad (8)$$

where $\{W_k\}_{k=0}^K$ are arbitrary fixed unitary operators and $\{H_k\}_{k=1}^K$ are Hermitian operators. We denote the channel that corresponds to this unitary as $\mathcal{Y}(\boldsymbol{\theta})$. In the following propositions, we modify the circuits in Section 3.1 by inserting $Y(\boldsymbol{\theta})$ in the middle of the circuit and constraining the input state to be a computational basis state.

Proposition 3. *Consider a cost function $C(\boldsymbol{\theta}) = \text{Tr}[\mathcal{W}_{L_2}^{(n_d)} \circ \mathcal{Y}(\boldsymbol{\theta}) \circ (\mathcal{W}_{L_1}^{(n_d)})^\dagger(\rho) O_{z_n}]$, where ρ is a computational basis input state, $\mathcal{W}_{L_i}^{(n_d)}$ denotes the channel corresponding to L_i instances of CNOT ladders and local noise on n_d qubits as presented in Figure 2, and with computational basis measurement $O_{z_n} = |z_n\rangle\langle z_n|$; $z_n \in \{0, 1\}^n$. The partial derivative with respect to parameter θ_k is bounded as*

$$|\partial_{\theta_k} C| \leq (1-p)^{\Gamma_{n,n_d,L_1} + \Gamma_{n,n_d,L_2}} \|H_k\|_\infty, \quad (9)$$

where we denote

$$\Gamma_{n,n_d,L_i} = \begin{cases} \lfloor L_i \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \rfloor, & \text{for } n_d = 1, \\ \lfloor L_i \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \rfloor, & \text{for } n_d > 1, \end{cases} \quad (10)$$

for $i \in \{1, 2\}$.

Proposition 4. *Consider a cost function $C(\boldsymbol{\theta}) = \text{Tr}[\mathcal{B}_{m_2}^\dagger \circ \mathcal{T}_2 \circ \mathcal{Y}(\boldsymbol{\theta}) \circ \mathcal{T}_1 \circ \mathcal{B}_{m_1}(\rho) O_{z_n}]$, where ρ is a computational basis input state, $\mathcal{B}_{m_i}$ is an entangling operation between m_i pairs of qubits in subsystems A_i and B_i as considered in Proposition 2, $\mathcal{T}_i$ denotes noisy unitary evolution separable across the cut $A_i|B_i$, and with computational basis measurement $O_{z_n} = |z_n\rangle\langle z_n|$; $z_n \in \{0, 1\}^n$. The partial derivative with respect to parameter θ_k is bounded as*

$$|\partial_{\theta_k} C| \leq (1-p)^{L_{\tau_1} + L_{\tau_2}} \|H_k\|_\infty, \quad (11)$$

where we denote

$$L_{\tau_i} = \begin{cases} L_{\mathcal{U}_i}, & \text{for } \mathcal{T}_i = \tilde{\mathcal{U}}_i \otimes \mathcal{V}_i, \\ L_{\mathcal{V}_i}, & \text{for } \mathcal{T}_i = \mathcal{U}_i \otimes \tilde{\mathcal{V}}_i, \\ L_{\mathcal{U}_i} + L_{\mathcal{V}_i}, & \text{for } \mathcal{T}_i = \tilde{\mathcal{U}}_i \otimes \tilde{\mathcal{V}}_i. \end{cases} \quad (12)$$

for $i \in \{1, 2\}$, and $L_{\mathcal{U}_i}$ and $L_{\mathcal{V}_i}$ are defined in the same manner as $L_{\mathcal{U}}$ and $L_{\mathcal{V}}$ in Figure 3.

Propositions 3 and 4 demonstrate that when modifying the settings considered in Section 3.1 to include a parameterized unitary, the partial derivative of measurement outcomes with respect to any parameter displays similar scaling to the concentration of expectation values. Specifically, if we assume that the largest singular values of the generators $\{H_k\}_{k=1}^K$ of the unitary $Y(\boldsymbol{\theta})$ scale at most polynomially in n , this implies NIBPs for linear depth circuits if n_d also scales linearly in n , or for polynomial-depth circuits of degree 2 or greater for any $n_d > 0$.

4 Numerical Simulations

In this section we numerically investigate the behavior of the gradient of a cost function for the Hamiltonian Variational Ansatz (HVA) using both a single-qubit local depolarizing noise model and a model based on a trapped-ion quantum computer [42, 43]. We start with an introduction of the HVA in Section 4.1, and follow with a presentation of the results in Section 4.

4.1 Hamiltonian Variational Ansatz

We consider an HVA [44, 45] for a task of variational ground state search for a transverse field one-dimensional quantum Ising model. The model is given by a Hamiltonian

$$H = - \sum_{\langle i,j \rangle} X_i X_j - g \sum_{i=1}^n Z_i. \quad (13)$$

Here, g is a constant, and X, Z are Pauli matrices. We choose $g = 1$ and assume Periodic Boundary Conditions (PBC). The ansatz with L layers takes the form

$$|\psi(\boldsymbol{\theta})\rangle = e^{-i\theta_{2L} H_Z} e^{-i\theta_{2L-1} H_{XX}} \dots e^{-i\theta_2 H_Z} e^{-i\theta_1 H_{XX}} |0\rangle, \quad (14)$$

where $H_{XX} = \sum_{\langle i,j \rangle} X_i X_j$, $H_Z = \sum_{i=1}^n Z_i$, and $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_{2L})$ are parameters. We decompose $e^{-i\theta_k H_{XX}}$ and $e^{-i\theta_l H_Z}$ to entangling Molmer-Sørensen gates and single qubit rotations, respectively, as shown in Figure 4. We minimize the energy of the model using a cost function defined as

$$C(\boldsymbol{\theta}) = \langle \psi(\boldsymbol{\theta}) | H | \psi(\boldsymbol{\theta}) \rangle, \quad (15)$$

and compute the gradient

$$\nabla C(\boldsymbol{\theta}) = (\partial_{\theta_1} C(\boldsymbol{\theta}), \partial_{\theta_2} C(\boldsymbol{\theta}), \dots, \partial_{\theta_{2L}} C(\boldsymbol{\theta})), \quad (16)$$

using a parameter shift rule [46]. The gradient (16) is not affected by a barren plateau in the noiseless case [20] making it a convenient setup for numerical investigation of the noise effects on the gradient scaling. For technical details on the gradient computation see Appendix B.




Figure 4: **The HVA ansatz.** A single layer HVA ansatz (14) for a 4-qubit system with periodic boundary conditions. We decompose the ansatz to Molmer-Sørensen $XX(\theta_1) = e^{-i\theta_1 X_j X_{j+1}}$ gates and single qubit rotations $R_Z(\theta_2) = e^{-i(\theta_2/2)Z_k}$ which are native gates of a trapped-ion quantum computer.

4.2 Numerical results

We examine the scaling of the averaged absolute value of the partial derivative of the cost function with respect to a parameter θ_k $|\partial_{\theta_k} C|$ for the cost function of Eq. (15) in a system with $n = 8$ qubits and for an HVA ansatz with the number of layers L increasing from 1 to 600. The chosen range of L allows to showcase the asymptotic scaling of $|\partial_{\theta_k} C|$ with increasing L . The average is uniformly taken over both random sets of parameters θ and the parameters θ_k . We consider $n_d = 0, 1, 2, \dots, n$ in both local depolarizing noise and a realistic trapped-ion noise model [42, 47].

The depolarizing noise model applies a local depolarizing channel after each gate if the qubit is designated to be dirty. In the case of the XX gate, noise is represented by applying a single qubit depolarizing channel to each dirty qubit. The realistic model is based on real-world machines [42] and consists of single-qubit noise channels for single-qubit gates acting on the dirty qubits, and a correlated two-qubit noise channel after XX gates. The XX noise is applied only when both qubits are dirty. A detailed description of the noise models, including figures and parameters, can be found in Appendix C.

We assume that the dirty qubits form a contiguous block as in Figure 1. We gather results in Figure 5. Here we can see that $|\partial_{\theta_k} C|$ decays exponentially with L for $n_d > 0$ (as evidenced by a straight line in the log-linear plot). Furthermore, we can see that the averaged gradient also decreases exponentially with increasing n_d . This can be seen from the fact that for a fixed L , increasing n_d decreases the value of $|\partial_{\theta_k} C|$ by a constant factor in the log-linear plot. Hence, we find in both cases that for $n_d > 0$, $|\partial_{\theta_k} C|$ decays exponentially with L and n_d . Namely, we heuristically observe that $|\partial_{\theta_k} C| \in \mathcal{O}(a^{Ln_d})$ for some $a < 1$.

These results reflect the exponential scaling in L and n_d obtained in our theoretical results from Propositions 1 and 2, thus indicating that the scaling in our propositions might be present in realistic models.

Here we can also verify to what extent Eq. (5) holds in the considered realistic scenario. Namely, we perform numerical simulations for the same HVA ansatz as previously analyzed but now all the qubits are dirty and the local depolarizing probability is rescaled according to Eq. (5). For instance, in the case of depolarizing noise, we consider the error rates

$$p \rightarrow pf, \quad \text{with } f = 1/n, 2/n, \dots, (n-1)/n, 1. \quad (17)$$

The error rates in the case of the trapped-ion noise model are scaled analogously, see Appendix C for details.

Figure 5 shows that the results for $n_d < n$ (solid lines), with the local depolarizing and the trapped-ion noise models, are similar to the case when all qubits are dirty and the noise is rescaled by $f = n_d/n$ and $f = (n_d - 1)/n$, respectively (dashed lines). The difference in f comes from the fact that in the case of depolarizing noise, the number of noise channels per layer of the ansatz is proportional to n_d , while in the case of the trapped-ion noise model we have $n_d - 1$ noisy Molmer-Sørensen gates per layer, as explained in Appendix C. Therefore, our numerics demonstrate that adding clean qubits is approximately equal to rescaling the error rate by a factor n_d/n as in Eq. (5). In Appendix D we additionally analyze the gradient behavior for $n = 4, 6$ qubits, obtaining the same conclusions as here for $n = 8$.

We further explore the behavior of the averaged derivative of the cost function with respect to the parameters in Figure 6 by plotting $|\partial_{\theta_k} C|$ versus the total circuit error rate, which we define as the sum of the error rates of all gates in a circuit. For both the local depolarizing and the trapped-ion noise models we obtain with good accuracy a collapse of $|\partial_{\theta_k} C|$ plotted versus the circuit error rate for all considered values of n_d and f . In Appendix D we conduct a similar analysis for $n = 4, 6$ obtaining the same conclusion. We see small deviations in the depolarizing case which appear to increase with increasing system size. Furthermore, in the same appendix, we also explore the clean and dirty setup for low circuit depths ($L = 1 - 30$) up to $n = 10$ in the realistic noise model, obtaining similar behavior as for the higher depths.

We note that the deviations are within one standard deviation of the mean gradient which may indicate that there are finite sample artifacts. Nevertheless, to elucidate their origin a more systematic investigation would be necessary, which we leave to future work. We find that for the realistic noise model deviations are much smaller resulting in nearly perfect collapse. Finally, we note that in the case of the depolarizing




Figure 5: **HVA gradient scaling versus the number of HVA layers.** We plot the mean absolute value of the partial derivative of the cost function with respect to a parameter θ_k $|\partial_{\theta_k} C|$ for $L = 1, \dots, 600$ and $n = 8$. The average is taken over both 28 random sets of parameters θ and the parameters θ_k . In (a) we show results for local depolarizing noise. In (b) we show results for the trapped-ion realistic noise model. The solid lines are obtained for $n_d = 0, 1, \dots, 8$. In the plots larger n_d is plotted with brighter colors. The dashed lines are obtained for the case of n dirty qubits and error rates of the noisy gates scaled by a factor $f = 1/8, 2/8, \dots, 7/8$ with respect to their noise rates for the aforementioned simulations. We plot the results for n_d and $f = n_d/n$ with the same color.

noise, the total error is proportional to n_d/n , while in the case of the realistic noise, it scales approximately linearly with n_d/n as shown in an inset of Figure 6b. Therefore, the obtained collapses provide further examples of exponential suppression of the gradient by a ratio of n_d/n .

5 Discussion

We have found that partitioning a quantum computer into noiseless “clean” qubits and noisy “dirty” qubits typically has an analogous effect to lowering the overall error rate by the ratio $\frac{n_d}{n}$. Consequently, our model of clean qubits cannot avoid exponential scaling effects due to noise such as barren plateaus, and even a single dirty qubit will ruin the bunch eventually. As it has previously been shown that standard error mitigation techniques also cannot avoid exponential scaling effects due to noise [23, 48], our work provides further evidence that a fully noiseless system may be ultimately required to avoid such scaling. We note that even though we exclusively look at gradient scaling in our numerics, this also implies the concentration of expectation values due to the results of [16]. This means that our results have implications for the scalability of near-term algorithms in general, not just in the realm of training parameterized circuits.

On the other hand, we found numerically that each additional clean qubit exponentially suppresses the rate of gradient decay with respect to depth. Therefore, the inclusion of clean qubits helps to mitigate NIBPs and exponential concentration of expectation values due to noise. Our analytical results point to this being a more general phenomenon in partially

noisy circuits. We note that here, in order to numerically investigate the scaling of trainability at large depths, we assumed that the error rates of the noisy qubits are larger than the ones expected from future machines. Their reduction should enable even more impressive increases in computational reach than the ones shown in our numerics.

We expect this work to lay the groundwork for potential explorations into the post-NISQ era, which is seemingly drawing closer with ever-larger qubit counts and decreasing error rates. Our results indicate that quantum computation with both error-corrected and noisy qubits can mitigate the exponential concentration issues. Therefore, it motivates the search for practically viable realizations of this computational paradigm. The clean and dirty setup is an idealized model of such a quantum computer, as a clean qubit implementation requires the correction of every error occurring at a logical qubit.

A natural follow-up study would be an extension of our model that takes into account that real-world logical qubits are characterized by logical error rates. Such a model would also enable a more comprehensive study of strategies involving combining logical qubits with different logical error rates and their comparison to using the same code distance for all the device qubits. In order to better model real-world logical qubits, it is essential to consider the gate-dependent errors, as non-Clifford gates pose significantly greater implementation challenges than Clifford gates.

Another line of follow-up research would be to consider the application of the clean and dirty setup to algorithms that involve some qubits subject to many more gates than the others, such as the quantum




Figure 6: **HVA gradient scaling versus the total error rate.** The plots show the mean absolute value of the partial derivative of the cost function $|\partial_{\theta_k} C|$ from Figure 5 plotted versus the total error, which in the depolarizing case is proportional to $L_{\text{eff}} = L \frac{n_d}{n}$. The solid lines show the clean and dirty setup, while the dashed ones show the variable error setup, with darker lines representing lower noise levels. Both the depolarizing noise model (a) and the realistic noise (b) results show an exponential decay of the mean absolute value of the partial derivative with the total error rate up to small deviations. The inset in (b) shows the error rate for the two setups and the realistic noise model in relation to n_d demonstrating that it is approximately linear. Consequently, in both cases, we obtain that the mean gradient decays approximately exponentially with n_d .

phase estimation algorithm [38]. For such qubits, one naturally obtains higher effective error rates, and potentially better error suppression when just a fraction of qubits are corrected. Therefore, in such cases, the clean and dirty setup may result in even better error suppression than in the cases investigated here. Another worth-exploring scenario that may result in the clean and dirty setup as an effective model is a coupling of a sensor built from noisy qubits [49] to logical qubits processing information obtained from the sensor [50]. Finally, one can also ask about using the clean and dirty computational paradigm to design custom quantum algorithms that can maximize quantum advantage in the post-NISQ era.

6 Acknowledgements

We thank Andrew Arrasmith, Burak Sahinoglu, Chenfeng Cao, and Hsin-Yuan Huang for their helpful discussions. We also thank Tom O’Leary for producing an open-source Qiskit version of the clean and dirty setup for the repository. This work was supported by the Quantum Science Center (QSC), a National Quantum Information Science Research Center of the U.S. Department of Energy (DOE). DB was supported by the U.S. DOE through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science & Technology Institute and by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 955479. SW was supported by the UKRI EPSRC grant No. EP/T001062/1 the Samsung GRP grant. MHG and LC were initially supported by the U.S.

DOE, Office of Science, Office of Advanced Scientific Computing Research, under the Quantum Computing Application Teams program. Piotr C. acknowledges initial support from the Laboratory Directed Research and Development (LDRD) program of Los Alamos National Laboratory (LANL) under project number 20190659PRD4 with subsequent support by the National Science Centre (NCN), Poland under project 2019/35/B/ST3/01028. MC was supported by the LDRD program of LANL under project number 20210116DR. PJC also acknowledges initial support from the LANL ASC Beyond Moore’s Law Project.

7 Code availability

The scripts and data used to create all the plots and to recreate some of the data can be found in Appendix E.

References

- [1] Richard P. Feynman. “Simulating physics with computers”. *International Journal of Theoretical Physics* **21**, 467–488 (1982).
- [2] Laird Egan, Dripto M Debroy, Crystal Noel, Andrew Risinger, Daiwei Zhu, Debopriyo Biswas, Michael Newman, Muyuan Li, Kenneth R Brown, Marko Cetina, et al. “Fault-tolerant control of an error-corrected qubit”. *Nature* **598**, 281–286 (2021).
- [3] Peter W Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In *Proceedings 35th annual symposium on foun-*

- dations of computer science. *Pages 124–134*. Ieee (1994).
- [4] Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. “Quantum algorithm for linear systems of equations”. *Physical Review Letters* **103**, 150502 (2009).
 - [5] John Preskill. “Quantum computing in the NISQ era and beyond”. *Quantum* **2**, 79 (2018).
 - [6] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. “Variational quantum algorithms”. *Nature Reviews Physics* **3**, 625–644 (2021).
 - [7] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermann Heimonen, Jakob S Kottmann, Tim Menke, et al. “Noisy intermediate-scale quantum algorithms”. *Reviews of Modern Physics* **94**, 015004 (2022).
 - [8] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. “Quantum machine learning”. *Nature* **549**, 195–202 (2017).
 - [9] Michael A. Nielsen and Isaac L. Chuang. “Quantum computation and quantum information”. *Cambridge University Press*. Cambridge (2000).
 - [10] Dorit Aharonov, Michael Ben-Or, Russell Impagliazzo, and Noam Nisan. “Limitations of noisy reversible computation” (1996). url: <https://doi.org/10.48550/arXiv.1106.6189>.
 - [11] Michael Ben-Or, Daniel Gottesman, and Avinatan Hassidim. “Quantum refrigerator” (2013). url: <https://doi.org/10.48550/arXiv.1301.1995>.
 - [12] Daniel Stilck França and Raul Garcia-Patron. “Limitations of optimization algorithms on noisy quantum devices”. *Nature Physics* **17**, 1221–1227 (2021).
 - [13] Samson Wang, Enrico Fontana, M. Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J Coles. “Noise-induced barren plateaus in variational quantum algorithms”. *Nature Communications* **12**, 1–11 (2021).
 - [14] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. “Barren plateaus in quantum neural network training landscapes”. *Nature Communications* **9**, 1–6 (2018).
 - [15] M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. “Cost function dependent barren plateaus in shallow parametrized quantum circuits”. *Nature Communications* **12**, 1–12 (2021).
 - [16] Andrew Arrasmith, Zoë Holmes, Marco Cerezo, and Patrick J Coles. “Equivalence of quantum barren plateaus to cost concentration and narrow gorges”. *Quantum Science and Technology* **7**, 045015 (2022).
 - [17] Andrew Arrasmith, M. Cerezo, Piotr Czarnik, Lukasz Cincio, and Patrick J Coles. “Effect of barren plateaus on gradient-free optimization”. *Quantum* **5**, 558 (2021).
 - [18] M. Cerezo and Patrick J Coles. “Higher order derivatives of quantum neural networks with barren plateaus”. *Quantum Science and Technology* **6**, 035006 (2021).
 - [19] Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe. “Entanglement-induced barren plateaus”. *PRX Quantum* **2**, 040316 (2021).
 - [20] Martin Larocca, Piotr Czarnik, Kunal Sharma, Gopikrishnan Muraleedharan, Patrick J. Coles, and M. Cerezo. “Diagnosing Barren Plateaus with Tools from Quantum Optimal Control”. *Quantum* **6**, 824 (2022).
 - [21] Zoë Holmes, Kunal Sharma, M. Cerezo, and Patrick J Coles. “Connecting ansatz expressibility to gradient magnitudes and barren plateaus”. *PRX Quantum* **3**, 010313 (2022).
 - [22] Supanut Thanasilp, Samson Wang, Nhat A Nghiem, Patrick J. Coles, and M. Cerezo. “Subtleties in the trainability of quantum machine learning models” (2021). url: <https://arxiv.org/abs/2110.14753>.
 - [23] Samson Wang, Piotr Czarnik, Andrew Arrasmith, M. Cerezo, Lukasz Cincio, and Patrick J Coles. “Can error mitigation improve trainability of noisy variational quantum algorithms?” (2021). url: <https://doi.org/10.48550/arXiv.2109.01051>.
 - [24] Ningping Cao, Junan Lin, David Kribs, Yiu-Tung Poon, Bei Zeng, and Raymond Laflamme. “NISQ: Error correction, mitigation, and noise simulation” (2021). url: <https://doi.org/10.48550/arXiv.2111.02345>.
 - [25] Adam Holmes, Mohammad Reza Jokar, Ghasem Pasandi, Yongshan Ding, Massoud Pedram, and Frederic T Chong. “NISQ+: Boosting quantum computing power by approximating quantum error correction”. In 2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA). Pages 556–569. IEEE (2020). url: <https://doi.org/10.1109/ISCA45697.2020.00053>.
 - [26] Yasunari Suzuki, Suguru Endo, Keisuke Fujii, and Yuuki Tokunaga. “Quantum error mitigation as a universal error reduction technique: Applications from the NISQ to the fault-tolerant quantum computing eras”. *PRX Quantum* **3**, 010345 (2022).
 - [27] Emanuel Knill and Raymond Laflamme. “Power of one bit of quantum information”. *Physical Review Letters* **81**, 5672 (1998).
 - [28] Keisuke Fujii, Hirotada Kobayashi, Tomoyuki Morimae, Harumichi Nishimura, Shuhei Tamate,

- and Seiichiro Tani. “Power of Quantum Computation with Few Clean Qubits”. *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)* **55**, 13:1–13:14 (2016).
- [29] Tomoyuki Morimae, Keisuke Fujii, and Harumichi Nishimura. “Power of one nonclean qubit”. *Physical Review A* **95**, 042336 (2017).
- [30] Craig Gidney. “Factoring with $n+2$ clean qubits and $n-1$ dirty qubits” (2017). url: <https://doi.org/10.48550/arXiv.1706.07884>.
- [31] Anirban N. Chowdhury, Rolando D. Somma, and Yi ğit Subaşı. “Computing partition functions in the one-clean-qubit model”. *Physical Review A* **103**, 032422 (2021).
- [32] Keisuke Fujii, Hirotada Kobayashi, Tomoyuki Morimae, Harumichi Nishimura, Shuhei Tamate, and Seiichiro Tani. “Impossibility of classically simulating one-clean-qubit model with multiplicative error”. *Physical Review Letters* **120**, 200502 (2018).
- [33] Raymond Laflamme, Cesar Miquel, Juan Pablo Paz, and Wojciech Hubert Zurek. “Perfect quantum error correcting code”. *Phys. Rev. Lett.* **77**, 198–201 (1996).
- [34] Daniel Gottesman. “An introduction to quantum error correction and fault-tolerant quantum computation”. *Quantum information science and its contributions to mathematics, Proceedings of Symposia in Applied Mathematics* **63**, 13–58 (2010).
- [35] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. “Surface codes: Towards practical large-scale quantum computation”. *Physical Review A* **86**, 032324 (2012).
- [36] A Yu Kitaev. “Quantum computations: algorithms and error correction”. *Russian Mathematical Surveys* **52**, 1191 (1997).
- [37] Chris N Self, Marcello Benedetti, and David Amaro. “Protecting expressive circuits with a quantum error detection code” (2022). url: <https://doi.org/10.48550/arXiv.2211.06703>.
- [38] Rolando D Somma. “Quantum eigenvalue estimation via time series analysis”. *New Journal of Physics* **21**, 123025 (2019).
- [39] Vojtěch Havlíček, Antonio D Córcoles, Kristan Temme, Aram W Harrow, Abhinav Kandala, Jerry M Chow, and Jay M Gambetta. “Supervised learning with quantum-enhanced feature spaces”. *Nature* **567**, 209–212 (2019).
- [40] Andrew G Taube and Rodney J Bartlett. “New perspectives on unitary coupled-cluster theory”. *International journal of quantum chemistry* **106**, 3393–3401 (2006).
- [41] Sumeet Khatri, Ryan LaRose, Alexander Poremba, Lukasz Cincio, Andrew T Sornborger, and Patrick J Coles. “Quantum-assisted quantum compiling”. *Quantum* **3**, 140 (2019).
- [42] Colin J Trout, Muyuan Li, Mauricio Gutiérrez, Yukai Wu, Sheng-Tao Wang, Luming Duan, and Kenneth R Brown. “Simulating the performance of a distance-3 surface code in a linear ion trap”. *New Journal of Physics* **20**, 043038 (2018).
- [43] Lukasz Cincio, Yi ğit Subaşı, Andrew T Sornborger, and Patrick J Coles. “Learning the quantum algorithm for state overlap”. *New Journal of Physics* **20**, 113022 (2018).
- [44] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. “A quantum approximate optimization algorithm” (2014). url: <https://doi.org/10.48550/arXiv.1411.4028>.
- [45] Stuart Hadfield, Zihui Wang, Bryan O’Gorman, Eleanor G Rieffel, Davide Venturelli, and Rupak Biswas. “From the quantum approximate optimization algorithm to a quantum alternating operator ansatz”. *Algorithms* **12**, 34 (2019).
- [46] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. “Evaluating analytic gradients on quantum hardware”. *Physical Review A* **99**, 032331 (2019).
- [47] Lukasz Cincio, Kenneth Rudinger, Mohan Sarovar, and Patrick J. Coles. “Machine learning of noise-resilient quantum circuits”. *PRX Quantum* **2**, 010324 (2021).
- [48] Ryuji Takagi, Suguru Endo, Shintaro Minagawa, and Mile Gu. “Fundamental limits of quantum error mitigation”. *npj Quantum Information* **8**, 114 (2022).
- [49] Sergey Danilin, Nicholas Nugent, and Martin Weides. “Quantum sensing with tuneable superconducting qubits: optimization and speed-up” (2022). url: <https://arxiv.org/abs/2211.08344>.
- [50] Nikolai Lauk, Neil Sinclair, Shabir Barzanjeh, Jacob P Covey, Mark Saffman, Maria Spiropulu, and Christoph Simon. “Perspectives on quantum transduction”. *Quantum Science and Technology* **5**, 020501 (2020).
- [51] Bernhard Baumgartner. “An inequality for the trace of matrix products, using absolute values” (2011). url: <https://doi.org/10.48550/arXiv.1106.6189>.

A Theoretical results

A.1 Preliminaries

Definition 1 (Depolarizing noise). *The action of a single-qubit depolarizing noise channel $\mathcal{D}_p$ with error probability p on a single-qubit quantum state ρ can be written as*

$$\mathcal{D}_p(\rho) = (1-p)\rho + p\frac{\mathbb{1}}{2}, \quad (18)$$

where $\mathbb{1}$ is the 2×2 identity matrix.

Definition 2 (Pauli strings). *We define the set of non-identity Pauli tensor product strings P_n as*

$$P_n = \{\mathbb{1}, X, Y, Z\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}. \quad (19)$$

Action of local depolarizing noise on Pauli strings. We note that for any $\sigma \in P_n$ we have

$$\mathcal{D}_p^{\otimes n}(\sigma) = q_\sigma \sigma, \quad (20)$$

where $|q_\sigma| \leq 1 - p$.

A.2 Useful Lemmas

Lemma 1 (Connecting purity to Hilbert-Schmidt norm). *For any given state ρ , the purity $P(\rho) = \text{Tr}[\rho^2]$ satisfies*

$$P(\rho) - \frac{1}{2^n} = \left\| \rho - \frac{\mathbb{1}}{2^n} \right\|_2^2, \quad (21)$$

where $\|\cdot\|_2$ denotes the Schatten 2-norm.

Proof. We have

$$\text{Tr}[\rho^2] = \text{Tr} \left[\left(\rho - \frac{\mathbb{1}}{2^n} \right)^2 + \frac{2}{2^n} \rho - \frac{1}{2^n} \frac{\mathbb{1}}{2^n} \right] \quad (22)$$

$$= \text{Tr} \left[\left(\rho - \frac{\mathbb{1}}{2^n} \right)^2 \right] + \frac{1}{2^n} \quad (23)$$

$$= \left\| \rho - \frac{\mathbb{1}}{2^n} \right\|_2^2 + \frac{1}{2^n}, \quad (24)$$

where in the first line we have added and subtracted $\frac{2}{2^n} \rho - \frac{1}{2^n} \frac{\mathbb{1}}{2^n}$, in the second line we have explicitly computed the trace for the last two terms, and in the final line we have used the definition of the Schatten 2-norm. $\square$

Lemma 2. (2-norm of commutators.) *Let X and Y be complex matrices. Let $\|\cdot\|_2$ and $\|\cdot\|_\infty$ denote the Schatten 2- and ∞ -norms respectively. We have*

$$\|[X, Y]\|_2 \leq 2 \|X\|_2 \|Y\|_\infty. \quad (25)$$

Proof. Writing out the commutator explicitly, we have

$$\|XY - YX\|_2 \leq \|XY\|_2 + \|YX\|_2 \quad (26)$$

$$\leq \|X\|_2 \|Y\|_\infty + \|Y\|_\infty \|X\|_2, \quad (27)$$

where the first inequality is due to the triangle inequality, and the second inequality is an application of the tracial matrix Hölder inequality [51]. $\square$

A.3 Model 1 - CNOT ladders

In this section we consider a circuit consisting of CNOT ladders as presented in Figure 7(a). After each ladder, a fraction n_d/n of the qubits experiences single-qubit depolarizing noise $\mathcal{D}_p$.

We first provide a brief roadmap of our proof methods. In order to analyze such circuits, we consider the Heisenberg picture and bound the output purity of the reverse circuit (Figure 7(b)). As we only consider computational basis measurements, we only need to consider Pauli tensor product strings of the form $\{1, Z\}^{\otimes n}$. We will adopt the notation $A \dots B$ to represent a tensor product string $A \otimes \dots \otimes B$. Due to Eq. (20), local depolarizing noise on n_d qubits will decrease the amplitude of certain strings in $\{1, Z\}^{\otimes n}$. Specifically, wherever a string has a "Z" on a noisy register, that string's amplitude will decay by a factor $(1 - p)$ for each instance of local depolarizing noise $\mathcal{D}_p$. The goal of our analysis will be to observe patterns in Pauli strings as they are acted on by CNOT ladders and to keep track of the "best case" Pauli string that is least affected by noise. This corresponds to the Pauli string that has the fewest number of "Z"s appearing on noisy registers. This can then be used to establish upper bounds on expectation value concentration.

Definition 3 (CNOT ladder). *We denote the mapping corresponding to one CNOT ladder (consisting of CNOT gates between all pairs of adjacent qubits in a 1D chain of n qubits as in Figure 7(b)) as CNOT_n .*

Definition 4 (CNOT ladder + noise). *We denote the channel corresponding to a CNOT ladder plus depolarizing noise on the first n_d qubits as $\mathcal{W}^{(n_d)\dagger} := \text{CNOT}_n \circ (\mathcal{D}_p^{\otimes n_d} \otimes \mathcal{I}^{n-n_d})$. Further, we denote the channel corresponding to L instances of this channel (i.e. the full circuit in Figure 7(b)) as $\mathcal{W}_L^{(n_d)\dagger} := (\mathcal{W}^{(n_d)\dagger})^L$.*




Figure 7: (a) Our goal is to consider a circuit with a repeating ansatz of L layers of a CNOT ladder. After each layer, the qubits on the first n_d registers go through depolarizing noise channel $\mathcal{D}_p$. At the end of the circuit, we measure in the computational basis. We denote ρ as the input state. (b) In order to analyze such circuits, we consider the Heisenberg picture and bound the output purity of the reverse circuit, with computational basis input states.

As a preliminary example, we consider the mapping of the string $Z \otimes 1 \otimes Z \otimes Z$ under such CNOT ladders:

$$Z1ZZ \xrightarrow{\text{CNOT}_4} ZZ1Z \xrightarrow{\text{CNOT}_4} 1ZZZ \xrightarrow{\text{CNOT}_4} Z11Z \xrightarrow{\text{CNOT}_4} Z1ZZ. \quad (28)$$

We note that after 4 iterations, the $Z1ZZ$ is mapped back to itself. We will refer to this as a string with period 4.

Definition 5 (Periodicity of a string). *We define the periodicity of a string X under a map $\mathcal{M}$ as the smallest value p such that $X \xrightarrow{\mathcal{M}^p} X$, where $\mathcal{M}^p$ denotes p applications of the map $\mathcal{M}$.*

We make two observations about the chain of mappings in (28). First, the string is mapped only to other strings in the set $\{1, Z\}^{\otimes n}$. Second, the mapping is cyclic. Both observations can be noted to be a direct

consequence of the fact that all strings from $\{\mathbb{1}, Z\}^{\otimes n}$ are mapped bijectively to another element of the set under CNOT_n , which we will now formally show.

Lemma 3. *The set $\{\mathbb{1}, Z\}^{\otimes n}$ is mapped to itself under CNOT_n bijectively.*

Proof. We can check the mappings for $n = 2$ and a single CNOT gate:

$$\mathbb{1}\mathbb{1} \xrightarrow{\text{CNOT}} \mathbb{1}\mathbb{1}, \quad (29)$$

$$\mathbb{1}Z \xrightarrow{\text{CNOT}} ZZ \left(\xrightarrow{\text{CNOT}} \mathbb{1}Z \right), \quad (30)$$

$$ZZ \xrightarrow{\text{CNOT}} \mathbb{1}Z, \quad (31)$$

$$Z\mathbb{1} \xrightarrow{\text{CNOT}} Z\mathbb{1}. \quad (32)$$

As CNOT_n is composed of $n - 1$ CNOT maps, it follows that strings in $\{\mathbb{1}, Z\}^{\otimes n}$ are mapped to strings in $\{\mathbb{1}, Z\}^{\otimes n}$ by CNOT_n . Finally, the map is unitary and so has an inverse, and thus the map is bijective. $\square$

Returning to our former example, we observe the elements of the cyclic mapping of $Z\mathbb{1}ZZ$ can be written as a slanted list

$$\begin{array}{cccc} Z & \mathbb{1} & Z & Z \\ & Z & \mathbb{1} & Z \\ & & \mathbb{1} & Z & Z & Z \\ & & & Z & \mathbb{1} & \mathbb{1} & Z \\ \hline & & & Z & \mathbb{1} & Z & Z \end{array} \quad \left| \begin{array}{l} r = 0 \\ r = 1 \\ r = 2 \\ r = 3 \\ r = 4 \end{array} \right.$$

where we denote the row of the list with r , and the line marks the end of a cycle. We now make a key remark.

Definition 6 (Inverted Binary Pascal Triangle). *We define the inverted binary Pascal triangle as the triangle of rows of binary strings, where the L elements in the k th row are uniquely generated by $L + 1$ elements of the $(k - 1)$ th row, with each element constructed by summing the element directly above and to the left, with the element directly above and to the right. An inverted binary Pascal triangle with a length n string in the 0th row has in total n rows, ending in a row with a length 1 string.*

Remark. *The list of cyclic mappings of any element $L \in \{\mathbb{1}, Z\}^{\otimes n}$ under CNOT_n has a one-to-one correspondence with the inverted binary Pascal triangle whose string in row 0 is generated by taking $L \otimes \mathbb{1}^{\otimes n-1}$ and applying the map $Z \leftrightarrow 1, \mathbb{1} \leftrightarrow 0$.*

As an example for the above remark, we can observe this correspondence for the string $Z\mathbb{1}ZZ$ (we truncate the triangle at $r = 3$).

$$\begin{array}{cccccc} Z & \mathbb{1} & Z & Z & \mathbb{1} & \mathbb{1} & \mathbb{1} \\ & Z & Z & \mathbb{1} & Z & \mathbb{1} & \mathbb{1} \\ & & \mathbb{1} & Z & Z & Z & \mathbb{1} \\ & & & Z & \mathbb{1} & \mathbb{1} & Z \end{array} \rightarrow \begin{array}{cccccc} 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ & 1 & 1 & 0 & 1 & 0 & 0 \\ & & 0 & 1 & 1 & 1 & 0 \\ & & & 1 & 0 & 0 & 1 \end{array}$$

We now prove some general statements about inverted binary Pascal triangles.

Lemma 4. *Given an inverted binary Pascal triangle with length $n = 2^x + 1$; $x \in \mathbb{Z}$ of the form $A \dots B$, the single entry in row $n - 1$ has value $A \oplus B$. That is, all triangles of this size take the form*

$$\begin{array}{ccc} A & \dots & B \\ & \dots & \dots \\ & & \dots \\ & & A \oplus B \end{array} \quad \left| \begin{array}{l} r = 0 \\ \dots \\ \dots \\ r = n - 1 \end{array} \right.$$

Proof. We prove the lemma by induction. The primal case of $n = 2$ can be checked by explicitly writing the inverted Pascal triangles as

$$\begin{array}{cccc} 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ & 0 & & 1 & & 1 & & 0 \end{array} .$$

Now suppose that strings of some longer length k also always satisfy this property. For instance, for length k string $A \dots B$, we have

$$\begin{array}{cccc|l} A & \dots & B & & r = 0 \\ & \dots & \dots & & \dots \\ & & \dots & & \dots \\ & & A \oplus B & & r = k - 1 \end{array}$$

Further, we consider another length k string $B \dots C$ that satisfies this property. This implies that the merged length $2k - 1$ string $A \dots B \dots C$ has an inverted Pascal triangle of the form

$$\begin{array}{cccccc|l} A & \dots & B & \dots & C & & r = 0 \\ & \dots & \dots & \dots & \dots & & \dots \\ & & \dots & & \dots & & \dots \\ & & A \oplus B & \dots & B \oplus C & & r = k - 1 \\ & & & \dots & \dots & & \dots \\ & & & & \dots & & \dots \\ & & & & A \oplus C & & r = 2k - 2 \end{array}$$

By merging strings together inductively, we see that strings of a particular length satisfy this property. The lengths satisfy a sequence with values

$$a(m) = 2a(m-1) - 1 \quad ; \quad m \in \mathbb{Z}, \quad a(0) = 2. \quad (33)$$

It can be verified that this sequence has values $a(m) = 2^m + 1$, which completes the proof. $\square$

Corollary 1 (Bit-wise addition of strings). *Given a string composed of adjacently joining two strings of length $n \in 2^x$; $x \in \mathbb{Z}$, starting in row 0, the first n entries of the string in row n is the binary bit-wise addition of both strings.*

Proof. Using Lemma 4, this can be seen by considering two strings $A \dots B$, $C \dots D$, each of length n , and by drawing out the triangle corresponding to the string $A \dots B \ C \dots D$:

$$\begin{array}{cccccc|l} A & \dots & BC & \dots & D & & r = 0 \\ & \dots & \dots & \dots & \dots & & \dots \\ & & \dots & & \dots & & \dots \\ & & A \oplus C & \dots & B \oplus D & & r = n \end{array}$$

$\square$

Lemma 5 (Maximum cycle period). *Under the CNOT ladder CNOT_n , any length n string has a cycle of periodicity $p = 2^x$; $x \in \mathbb{Z}$, where*

$$x \leq \lceil \log_2 n \rceil. \quad (34)$$

That is, the periodicity is an integer power of 2, and at most the closest power of 2 to n that is greater than n .

Proof. Suppose we start with a length n string $A \dots B$. Now we can append arbitrary elements to the start of the string such that the total length is $2^{\lceil \log_2 n \rceil}$, and we denote this new string as $A' \dots A \dots B$. Due to Corollary 1, we have the following triangle

$$\begin{array}{cccccc|l} A' \dots A \dots B & 0 & \dots & 0 & & & r = 0 \\ & \dots & \dots & \dots & \dots & & \dots \\ & & \dots & & 0 & & \dots \\ & & A' \dots A \dots B & & & & r = 2^{\lceil \log_2 n \rceil} \end{array}$$

where $A' \dots A \dots B \dots 0 \dots 0$ denotes the string with $2 \cdot 2^{\lceil \log_2 n \rceil}$ entries where we have appended $2^{\lceil \log_2 n \rceil}$ elements in "0". Thus, $A \dots B$ reappears after $2^{\lceil \log_2 n \rceil}$ steps, implying it has a period of at most $2^{\lceil \log_2 n \rceil}$. As after $2^{\lceil \log_2 n \rceil}$ steps $A \dots B$ must be mapped back to itself, the period must be a factor of $2^{\lceil \log_2 n \rceil}$, which only has prime factor 2. Therefore, the period of $A \dots B$ must itself be a power of 2. $\square$

We now have the tools to understand how purity decays with circuit depth under CNOT_n with local depolarizing noise. We will consider the circuit in Figure 7.

Lemma 6. *For any computational basis state $|z_n\rangle\langle z_n|$ with $z_n \in \{0,1\}^n$, we have*

$$P\left(\mathcal{W}_L^{(n_d)\dagger}(|z_n\rangle\langle z_n|)\right) - \frac{1}{2^n} \leq (1-p)^{2\Gamma_{n,n_d,L}} \left(P(|z\rangle\langle z|) - \frac{1}{2^n}\right) = (1-p)^{2\Gamma_{n,n_d,L}}, \quad (35)$$

with

$$\Gamma_{n,n_d,L} = \begin{cases} \lfloor L \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \rfloor, & \text{for } n_d = 1, \\ \lceil L \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \rceil, & \text{for } n_d > 1, \end{cases} \quad (36)$$

where $P(\cdot)$ denotes the purity function, and $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ denote the integer floor and ceiling functions respectively.

Remark. *If n is a power of 2, then the exponent in Eqs. (35) is given*

$$\Gamma_{n,n_d,L} = \left\lfloor L \cdot \frac{n_d}{n} \right\rfloor. \quad (37)$$

Thus, we see that the exponent in Eqs. (35) is approximately equal to $L \cdot \frac{n_d}{n}$.

Proof. We start with the proof for number of dirty qubits $n_d = 1$, before showing how a looser bound can be constructed for arbitrary n_d . From Lemma 5 we know that the period of a string of length n is at most $2^{\lceil \log_2 n \rceil}$. We now show that there always exists a binary string with period $2^{\lceil \log_2 n \rceil}$ that corresponds to the Pauli string that is least affected by noise in the circuit in Fig. 7.

Consider a string of length n made up of all "0"s, with "1" on the $2^{\lceil \log_2 n \rceil}$ th entry. By inspection, we see that the "1" propagates out to the left and the right in a lightcone structure. In addition, due to Lemma 5, the string must repeat after $2^{\lceil \log_2 n \rceil}$ steps as

$$\begin{array}{cccccccc} 0 & \dots & \dots & 0 & 1 & 0 & \dots & \dots & 0 \\ 0 & \dots & 0 & 1 & 1 & 0 & \dots & \dots & 0 \\ & \searrow & \swarrow & \swarrow & \dots & \searrow & \searrow & & \searrow \\ & & 0 & 1 & \dots & \dots & 1 & 0 & \dots & 0 \\ & & & 1 & \dots & & 1 & 0 & \dots & \dots & 0 \\ & & & & 0 & \dots & \dots & 0 & 1 & 0 & \dots & \dots & 0 \end{array} \quad \left| \begin{array}{l} r = 0 \\ r = 1 \\ \dots \\ r = 2^{\lceil \log_2 n \rceil} - 2 \\ r = 2^{\lceil \log_2 n \rceil} - 1 \\ r = 2^{\lceil \log_2 n \rceil} \end{array} \right.$$

Due to the lightcone structure, it is clear there is only one string per cycle with "1" as its first entry. Thus, every $2^{\lceil \log_2 n \rceil}$ steps there is one string with a "1" in the first entry. Now we show that no other string that produces a cycle with proportionally fewer "1"s in the first entry every L steps.

First, we remark that no nontrivial string with period smaller than $2^{\lceil \log_2 n \rceil}$ can produce such a cycle, as due to the lightcone structure of "1"s propagating left, there must always be at least one string with a "1" in the first entry. Second, we note that if n is a power of 2, then $2^{\lceil \log_2 n \rceil} = 2^{\log_2 n}$ and thus with the above we have already found the cycle with the proportionally fewest "1"s in the first entry.

Thus, it only remains to consider strings with length n not a power of 2, and with period $2^{\lceil \log_2 n \rceil}$ (as this is the maximum possible period, as set by Lemma 5). We remark that for such strings, there must be at least one string with "1" as the first entry within each cycle, as any "1" propagates left in a lightcone. Without loss of generality, we consider the first string of each cycle to be such a string. As we are considering non-trivial strings and cycles, the second string in the cycle must also have a "1" at some position. If the left-most "1" is the first entry, then we have trivially found a cycle with two adjacent strings with "1" in the first entry. If the left-most "1" is not the first entry, the "1" will propagate left with each step in a light cone. As $2^{\lceil \log_2 n \rceil} > n$ if n is not a power of 2, it is assured that the light cone of "1"s reaches the first entry before the end of the cycle. Thus, any string with with period $2^{\lceil \log_2 n \rceil}$ has at least two strings per cycle with "1" in the first entry.

We can summarize the above result as

$$\text{number of "1"s in first entry after } L \text{ steps} \geq \left\lfloor L/2^{\lceil \log_2 n \rceil} \right\rfloor, \quad (38)$$

for any string in $\{\mathbb{1}, Z\}^{\otimes n}$, where the bound is achievable with the string of all "0"s, with "1" on the $2^{\lceil \log_2 n \rceil} \text{th}$ entry.

We now return to the quantum problem, recalling that Z corresponds to "1" and $\mathbb{1}$ corresponds to "0". From Eq. (38) we can conclude that through the L channel instances that make up $\mathcal{W}_L^{(1)}$, any input string in $\{Z, \mathbb{1}\}^{\otimes n}$ will produce at least $\lfloor L/2^{\lceil \log_2 n \rceil} \rfloor$ strings with Z in the first register. Thus, for any $\sigma \in \{Z, \mathbb{1}\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}$ we have

$$\mathcal{W}_L^{(1)\dagger}(\sigma) = q_{\sigma,L}^{(1)} \sigma', \quad (39)$$

where $\sigma' \in \{Z, \mathbb{1}\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}$ and for some $\{q_{\sigma,L}^{(1)}\}_\sigma$ satisfying

$$\max_{\sigma \in \{Z, \mathbb{1}\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}} |q_{\sigma,L}^{(1)}| = (1-p)^{\lfloor L/2^{\lceil \log_2 n \rceil} \rfloor}. \quad (40)$$

Moreover, the map in Eq. 39 bijectively maps Pauli strings to Pauli strings. We note that computational basis states can be written as $|z_n\rangle\langle z_n| = \frac{1}{2^n} \sum_{\sigma \in \{\mathbb{1}, Z\}^{\otimes n}} s_{z_n, \sigma} \sigma$ where $s_{z_n, \sigma} \in \{+1, -1\}$ and $s_{z_n, \mathbb{1}} = 1$. We can write

$$P\left(\mathcal{W}_L^{(1)\dagger}(|z_n\rangle\langle z_n|)\right) - \frac{1}{2^n} = \left\| \frac{1}{2^n} \sum_{\sigma \in \{\mathbb{1}, Z\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}} s_{z_n, \sigma} \mathcal{W}_L^{(1)\dagger}(\sigma) \right\|_2^2 \quad (41)$$

$$= \left\| \frac{1}{2^n} \sum_{\sigma \in \{\mathbb{1}, Z\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}} s_{z_n, \sigma} q_{\sigma,L}^{(1)} \sigma' \right\|_2^2 \quad (42)$$

$$= \frac{1}{2^{2n}} \sum_{\sigma \in \{\mathbb{1}, Z\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}} |s_{z_n, \sigma} q_{\sigma,L}^{(1)}|^2 \text{Tr}[\mathbb{1}] \quad (43)$$

$$= \frac{1}{2^n} \sum_{\sigma \in \{\mathbb{1}, Z\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}} |q_{\sigma,L}^{(1)}|^2 \quad (44)$$

$$\leq (1-p)^{2\lfloor L/2^{\lceil \log_2 n \rceil} \rfloor} \left(\frac{2^n - 1}{2^n} \right) \quad (45)$$

$$= (1-p)^{2\lfloor L/2^{\lceil \log_2 n \rceil} \rfloor} \left(P(|z_n\rangle\langle z_n|) - \frac{1}{2^n} \right), \quad (46)$$

where in the first equality we have used Lemma 1, in the second equality we have used Eq. (39), in the third equality we have used the definition of the Schatten 2-norm and the orthogonality of the Pauli matrices under the Hilbert-Schmidt inner product. The inequality comes from Eq. 40.

In order to generalize the results from $\mathcal{W}_L^{(1)\dagger}$ to $\mathcal{W}_L^{(n_d)\dagger}$ (that is, one dirty qubit to n_d dirty qubits), we argue that for each extra dirty qubit, any string $\{Z, \mathbb{1}\}^{\otimes n}$ generates a cycle that has at least one extra Z per cycle in the first n_d qubits. We now relax our previous result and consider cycles of maximal length $2^{\lceil \log_2 n \rceil}$. As argued above, any non-trivial string generates a cycle such that at least one string with "1" as the first entry. We suppose that such a string appears in row k . From the rules of a Pascal triangle, this implies there must be a "1" in row $k-1$ either as the first entry or the second, i.e. we have either

$$\begin{array}{cc|c} 0 & 1 & 1 & 0 \\ & 1 & & 1 \end{array} \quad \left| \quad \begin{array}{l} r = k-1 \\ r = k. \end{array} \right.$$

We recall that this corresponds to a Pauli string with a Z in the first or second qubit. Thus, if the first and second qubits are now dirty, this guarantees for any input string an additional factor of $1-p$ per cycle. In general we can consider $n_d \leq n$ dirty qubits and the light cone of influence up from the first entry of row k up to the first n_d entries of row $k - n_d - 1$. In this triangular light cone, there will exist at least n_d "1"s. Thus, we have

$$\text{number of "1"s in first } n_d \text{ entries after } L \text{ steps} \geq \left\lfloor L n_d / 2^{\lceil \log_2 n \rceil} \right\rfloor. \quad (47)$$

This implies that

$$\mathcal{W}_L^{(n_d)\dagger}(\sigma) = q_{\sigma,L}^{(n_d)} \sigma', \quad (48)$$

where $\sigma' \in \{Z, \mathbb{1}\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}$ and for some $\{q_{\sigma,L}^{(1)}\}_\sigma$ satisfying

$$\max_{\sigma \in \{Z, \mathbb{1}\}^{\otimes n} / \{\mathbb{1}^{\otimes n}\}} |q_{\sigma,L}^{(n_d)}| = (1-p) \lfloor n_d L / 2^{\lceil \log_2 n \rceil} \rfloor. \quad (49)$$

Proceeding with the same steps as the above, for n_d dirty qubits we obtain the desired result with exponent $\lfloor L \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \rfloor$. $\square$

We now restate and prove Proposition 1, which shows how the output distribution of Figure 7(a) exponentially concentrates.

Proposition 1. *Consider the circuit in Figure 7(a) with n_d dirty qubits and computational basis measurement $O_{z_n} = |z_n\rangle\langle z_n|$; $z_n \in \{0,1\}^n$. Denote the channel that describes the action of the gates and noise in the circuit as $\mathcal{W}_L^{(n_d)}$. For any input state ρ , the expectation value concentrates as*

$$\left| \text{Tr}[\mathcal{W}_L^{(n_d)}(\rho) O_{z_n}] - \frac{1}{2^n} \right| \leq (1-p)^{\Gamma_{n,n_d,L}} \sqrt{P(\rho)}, \quad (50)$$

for any computational basis measurement, where $P(\rho)$ denotes the purity of ρ and we have defined

$$\Gamma_{n,n_d,L} = \begin{cases} \lfloor L \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \rfloor, & \text{for } n_d = 1, \\ \lfloor L \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \rfloor, & \text{for } n_d > 1. \end{cases} \quad (51)$$

Proof. We have

$$\left| \text{Tr}[\mathcal{W}_L^{(n_d)}(\rho) O_{z_n}] - \frac{1}{2^n} \right| = \left| \text{Tr}[\mathcal{W}_L^{(n_d)}(\rho) O_{z_n}] - \text{Tr}[\mathcal{W}_L^{(n_d)}(\rho) \frac{\mathbb{1}}{2^n}] \right| \quad (52)$$

$$= \left| \text{Tr} \left[\mathcal{W}_L^{(n_d)}(\rho) \left(O_{z_n} - \frac{\mathbb{1}}{2^n} \right) \right] \right| \quad (53)$$

$$\leq \|\rho\|_2 \left\| \left(\mathcal{W}_L^{(n_d)} \right)^\dagger \left(O_{z_n} - \frac{\mathbb{1}}{2^n} \right) \right\|_2 \quad (54)$$

$$= \sqrt{P(\rho)} \sqrt{P \left(\left(\mathcal{W}_L^{(n_d)} \right)^\dagger (|z_n\rangle\langle z_n|) \right) - \frac{1}{2^n}} \quad (55)$$

$$\leq \sqrt{P(\rho)} (1-p)^{\Gamma_{n,n_d,L}}, \quad (56)$$

where in the first equality we have used the fact that $\text{Tr}[\mathcal{W}_L^{(n_d)}(\rho)] = 1$, the second equality is simply a grouping of terms, the first inequality is due to the tracial matrix Hölder's inequality [51], the third equality is an application of Lemma 1, and the final inequality is an application of Lemma 6. $\square$

A.3.1 Pauli observables

We remark that Eq. 48 also implies the concentration of Pauli observables with ladder circuits. We make this precise with the following proposition.

Supplemental Proposition 1 (Pauli observables). *Consider the reverse ladder circuit $\mathcal{W}_L^{(n_d)\dagger}$ (pictured in Fig. 7b) with computational basis input state $|z_n\rangle\langle z_n|$; $z_n \in \{0,1\}^n$ and measurement observable O . Suppose further that O can be decomposed in the Pauli basis as $O = \sum_{\sigma \in W} \omega_\sigma \sigma$ with $W \subset \{\mathbb{1}, X, Y, Z\}^{\otimes n}$. Then, the output of this circuit concentrates as*

$$\left| \text{Tr} \left[\mathcal{W}_L^{(n_d)\dagger} (|z_n\rangle\langle z_n|) O \right] - \frac{1}{2^n} \text{Tr} [O] \right| \leq |W| \|\omega\|_\infty (1-p)^{\lfloor n_d L / 2^{\lceil \log_2 n \rceil} \rfloor}, \quad (57)$$

where we denote $\|\omega\|_\infty := \max_\sigma \{\omega_\sigma\}_\sigma$ as the Pauli coefficient of largest magnitude. Thus, the output of the circuit concentrates to the maximally mixed value exponentially quickly with increasing circuit depth.




Figure 8: (a) We consider a class of circuits with entangled measurements, that includes as a special case the Hilbert-Schmidt Test circuit. (b) In order to analyze such circuits, we bound the output purity of the reverse circuit with computational basis input states.

Proof. Starting from Eq. 48 we can directly evaluate

$$\left| \text{Tr} \left[\mathcal{W}_L^{(n_d)\dagger} (|z_n\rangle\langle z_n|) O \right] - \frac{1}{2^n} \text{Tr} [O] \right| = \frac{1}{2^n} \sum_{\sigma \in \{1, Z\}^{\otimes n} / \{1^{\otimes n}\}} \sum_{\sigma'' \in W} \left| \text{Tr} \left[q_{\sigma, L}^{(n_d)} \sigma' \omega_{\sigma''} \sigma'' \right] \right| \quad (58)$$

$$= \frac{1}{2^n} \sum_{\sigma \in \{1, Z\}^{\otimes n} / \{1^{\otimes n}\}} \sum_{\sigma'' \in W} \left| q_{\sigma, L}^{(n_d)} \omega_{\sigma''} \right| \delta_{\sigma' \sigma''} 2^n \quad (59)$$

$$\leq \frac{1}{2^n} \sum_{\sigma \in \{1, Z\}^{\otimes n} / \{1^{\otimes n}\}} \sum_{\sigma'' \in W} (1-p)^{\lfloor n_d L / 2^{\lfloor \log_2 n \rfloor} \rfloor} \|\omega\|_{\infty} \delta_{\sigma' \sigma''} 2^n \quad (60)$$

$$\leq |W| \|\omega\|_{\infty} (1-p)^{\lfloor n_d L / 2^{\lfloor \log_2 n \rfloor} \rfloor}, \quad (61)$$

where in the first line we have decomposed $|z_n\rangle\langle z_n|$ and O into the Pauli basis, the second line comes from the identity $\text{Tr}[\sigma' \sigma''] = \delta_{\sigma' \sigma''} 2^n$ where $\delta_{\sigma' \sigma''}$ is the Dirac delta function, the third line comes from taking the maximal magnitudes of $q_{\sigma, L}^{(n_d)}$ and $\omega_{\sigma''}$ separately, and the final line comes by noting there can only be at most $|W|$ non-zero terms in the sum. $\square$

We remark that this bound is particularly relevant if the measurement operator is composed of a small number of Pauli operators with bounded coefficients, compared to the circuit depth. For instance, if $\|\omega\|_{\infty} := \max_{\sigma} \{\omega_{\sigma}\}_{\sigma} \in \mathcal{O}(\text{poly}(n))$ and $N_O \in \mathcal{O}(\text{poly}(n))$ but $L = \Omega(n)$, then Supplemental Proposition 1 implies exponential concentration with increasing system size.

A.4 Model 2 - entangled registers

In this section we consider circuits of the form in Figure 8(a). In order to analyze such circuits we bound the output purity of the reverse circuit with computational basis input states (Figure 8(b)). This will allow us to consider the circuit in Figure 8(a) using the Heisenberg picture.

We consider two distinct subsystems A and B , of size n_A qubits and n_B qubits respectively. The input to the circuit in Figure 8(b) is a computational basis state $|z_n\rangle = |z_{n_A}\rangle |z_{n_B}\rangle$; $z_n \in \{0, 1\}^n$. We first implement m entangling operations with Hadamard gates and CNOTs on a subset of $2m$ registers between A and B , with an operation we denote as $\mathcal{B}_m$. Note that we allow for any generic m satisfying $m > 0$. The second step of the circuit is to act on A with unitary evolution $\mathcal{U}$, and act on B with unitary $\mathcal{V}$. Finally, we allow the action of a global unitary $\mathcal{W}$ on the joint system AB . We remark that this circuit includes as a special case the Hilbert-Schmidt Test circuit [41].

We consider three modifications of this circuit where we now consider the effects of noise. In the first modification, we consider subsystem A to be noisy, and thus we modify the channel $\mathcal{U} \rightarrow \tilde{\mathcal{U}}$, where $\tilde{\mathcal{U}}$ is the channel corresponding to the hardware implementation of the unitary channel $\mathcal{U}$ with an instance of local depolarizing noise $\mathcal{D}_p^{\otimes n_A}$ in between each layer of gates. This follows the noise model as considered in Ref. [13],




Figure 9: We consider the depth of $\mathcal{U}$ to be the number of non-parallelizable unitary layers in a given hardware implementation. When we consider a noisy modification $\mathcal{U} \rightarrow \tilde{\mathcal{U}}$, we insert layers of local depolarizing noise with depolarizing probability p in between each unitary layer.

and we present a schematic of this modification in Figure 9. In the second modification, we consider $\mathcal{V} \rightarrow \tilde{\mathcal{V}}$, where $\tilde{\mathcal{V}}$ is the noisy version of the unitary channel $\mathcal{V}$. We define the depth of the unitaries as the number of unitary layers in such an implementation and denote the depth of $\tilde{\mathcal{U}}(\tilde{\mathcal{V}})$ as $L_{\mathcal{U}}(L_{\mathcal{V}})$. Finally, we also consider the scenario when both A and B are noisy, and we have $\mathcal{U} \rightarrow \tilde{\mathcal{U}}, \mathcal{V} \rightarrow \tilde{\mathcal{V}}$.

Lemma 7. Consider the circuit in Figure 8(b) with entangling operation $\mathcal{B}_m$, separable noisy unitary evolution $\mathcal{T}$ and global unitary evolution $\mathcal{W}$. The purity of the circuit concentrates as

$$P(\mathcal{W} \circ \mathcal{T} \circ \mathcal{B}_m(|z_n\rangle\langle z_n|)) - \frac{1}{2^n} \leq (1-p)^{L_{\mathcal{T}}} \left(P(|z_n\rangle\langle z_n|) - \frac{1}{2^n} \right), \quad (62)$$

for any computational basis state input $|z_n\rangle\langle z_n|$; $z_n \in \{0,1\}^n$, where we have defined

$$L_{\mathcal{T}} = \begin{cases} L_{\mathcal{U}}, & \text{for } \mathcal{T} = \tilde{\mathcal{U}} \otimes \mathcal{V}, \\ L_{\mathcal{V}}, & \text{for } \mathcal{T} = \mathcal{U} \otimes \tilde{\mathcal{V}}, \\ L_{\mathcal{U}} + L_{\mathcal{V}}, & \text{for } \mathcal{T} = \tilde{\mathcal{U}} \otimes \tilde{\mathcal{V}}. \end{cases} \quad (63)$$

Thus, when only subsystem A (B) is noisy, there is still exponential scaling, however, compared to the fully noisy setting the exponent is modified by factor $\frac{L_{\mathcal{U}}}{L_{\mathcal{U}}+L_{\mathcal{V}}}$ ($\frac{L_{\mathcal{V}}}{L_{\mathcal{U}}+L_{\mathcal{V}}}$). If we consider the unitaries to have linear depth in the number of qubits and scaling with the same proportionally factor (i.e. $L_{\mathcal{U}}/n_A = L_{\mathcal{V}}/n_B$), then the modification to the exponent is $\frac{n_A}{n_A+n_B}$ ($\frac{n_B}{n_A+n_B}$).

Proof. Consider the mapping of strings in $\{\mathbb{1}, Z\}^{\otimes 2}$ under a Hadamard gate followed by a CNOT. For map $\mathcal{B}(\cdot) = (H \otimes \mathbb{1})\text{CNOT}(\cdot)\text{CNOT}^\dagger(H \otimes \mathbb{1})$ we have

$$\mathbb{1}\mathbb{1} \xrightarrow{\mathcal{B}} \mathbb{1}\mathbb{1}, \quad (64)$$

$$\mathbb{1}Z \xrightarrow{\mathcal{B}} ZZ, \quad (65)$$

$$Z\mathbb{1} \xrightarrow{\mathcal{B}} XX, \quad (66)$$

$$ZZ \xrightarrow{\mathcal{B}} -YY. \quad (67)$$

This implies that after the CNOTs in the circuit in Figure 8(b), the strings are (up to a phase factor) exclusively of the form $S_A \otimes S_B \in \text{Bell}$, where $S_A = \{\mathbb{1}, Z\}^{\otimes n_A-m} \otimes A \cup \{\mathbb{1}^{\otimes n_A}\}$ and $S_B \in A \otimes \{\mathbb{1}, Z\}^{\otimes n_B-m} \cup \{\mathbb{1}^{\otimes n_B}\}$, with $A \in \{X, Y, Z\}^{\otimes m}$. In particular, we make a key observation that $S_A = \mathbb{1}^{\otimes n_A}$ if and only if $S_B = \mathbb{1}^{\otimes n_B}$. This implies that we can import over the results of Ref. [13] and treat subsystems A and B individually in our analysis of the evolution under $\mathcal{T}$. Namely, we have for any $\sigma_A \in P_{n_A}$, with P_{n_A} the group of Pauli operators acting on subsystem A , we have

$$\mathcal{D}_p^{\otimes n_A}(\sigma_A) = q_{\sigma_A} \sigma_A, \quad (68)$$

$$\mathcal{D}_p^{\otimes n_B}(\sigma_B) = q_{\sigma_B} \sigma_B, \quad (69)$$

where $|q_{\sigma_B}| \leq 1-p$. A consequence of this is that for any unitary channel $\mathcal{Y}$ and for input operator

$\sum_{\sigma_A \in P_{n_A}, \sigma_B \in P_{n_B}} \lambda(\sigma_A, \sigma_B) \sigma_A \otimes \sigma_B$ where $\lambda(\sigma_A, \sigma_B) \in \mathbb{C}$ are arbitrary coefficients, we have

$$\left\| \mathcal{Y} \circ (\mathcal{D}_p^{\otimes n_A} \otimes \mathcal{I}) \left(\sum_{\sigma_A \in P_{n_A}, \sigma_B \in P_{n_B}} \lambda(\sigma_A, \sigma_B) \sigma_A \otimes \sigma_B \right) \right\|_2 = \left\| \mathcal{Y} \left(\sum_{\sigma_A \in P_{n_A}, \sigma_B \in P_{n_B}} q_{\sigma_A} \lambda(\sigma_A, \sigma_B) \sigma_A \otimes \sigma_B \right) \right\|_2 \quad (70)$$

$$= \left\| \sum_{\sigma_A \in P_{n_A}, \sigma_B \in P_{n_B}} q_{\sigma_A} \lambda(\sigma_A, \sigma_B) \sigma_A \otimes \sigma_B \right\|_2 \quad (71)$$

$$= \left(\sum_{\sigma_A \in P_{n_A}, \sigma_B \in P_{n_B}} |q_{\sigma_A} \lambda(\sigma_A, \sigma_B)|^2 \text{Tr}[(\sigma_A \otimes \sigma_B)^2] \right)^{1/2} \quad (72)$$

$$\leq \left(\sum_{\sigma_A \in P_{n_A}, \sigma_B \in P_{n_B}} (1-p)^2 |\lambda(\sigma_A, \sigma_B)|^2 \text{Tr}[(\sigma_A \otimes \sigma_B)^2] \right)^{1/2} \quad (73)$$

$$= (1-p) \left\| \sum_{\sigma_A \in P_{n_A}, \sigma_B \in P_{n_B}} \lambda(\sigma_A, \sigma_B) \sigma_A \otimes \sigma_B \right\|_2, \quad (74)$$

where in the first equality we have used Eq. (68), in the second equality we have used the unitary invariance of Schatten norms, the third equality comes from the definition of the Schatten 2-norm, and the inequality comes from the bound $|q_{\sigma_A}| \leq 1-p$. The final equality again comes from the definition of the Schatten 2-norm. We note that the output of $S_A \otimes S_B \in \text{Bell}/\{\mathbb{1}^{\otimes n}\}$ under unital channels separable across the cut $A|B$ can generically be written in the form as the argument in the left-hand side of Eq. (70). As we only consider such unitary channels in our analysis of $\mathcal{T}$, we thus can iteratively apply the above to write the concentration inequality

$$\left\| \tilde{\mathcal{U}} \otimes \mathcal{I} \left(\sum_{S_A \otimes S_B \in \text{Bell}/\{\mathbb{1}^{\otimes n}\}} \lambda(S_A, S_B) S_A \otimes S_B \right) \right\|_2 = \quad (75)$$

$$= \left\| \mathcal{D}_p^{\otimes n_A} \circ \mathcal{U}_L \circ \mathcal{D}_p^{\otimes n_A} \circ \dots \circ \mathcal{D}_p^{\otimes n_A} \circ \mathcal{U}_1 \otimes \mathcal{I} \left(\sum_{S_A \otimes S_B \in \text{Bell}/\{\mathbb{1}^{\otimes n}\}} \lambda(S_A, S_B) S_A \otimes S_B \right) \right\|_2 \quad (76)$$

$$\leq (1-p)^{L_U} \left\| \sum_{S_A \otimes S_B \in \text{Bell}/\{\mathbb{1}^{\otimes n}\}} \lambda(S_A, S_B) S_A \otimes S_B \right\|_2, \quad (77)$$

where $\lambda(S_A, S_B) \in \mathbb{C}$ are arbitrary coefficients. By repeating analogous steps to Eqs. (70)-(77) we can also obtain

$$\left\| \mathcal{I} \otimes \tilde{\mathcal{V}} \left(\sum_{S_A \otimes S_B \in \text{Bell}/\{\mathbb{1}^{\otimes n}\}} \lambda(S_A, S_B) S_A \otimes S_B \right) \right\|_2 \leq (1-p)^{L_V} \left\| \sum_{S_A \otimes S_B \in \text{Bell}/\{\mathbb{1}^{\otimes n}\}} \lambda(S_A, S_B) S_A \otimes S_B \right\|_2. \quad (78)$$

We now have the tools to prove the lemma. We recall that $|z_n\rangle\langle z_n| = \frac{1}{2^n} \sum_{\sigma \in \{\mathbb{1}, Z\}^{\otimes n}} s_{z_n, \sigma} \sigma$ where $s_{z_n, \sigma} \in$

$\{+1, -1\}$ and $s_{z_n, 1} = 1$. Then, we can write

$$P(\mathcal{W} \circ \mathcal{T} \circ \mathcal{B}_m(|z_n\rangle\langle z_n|)) - \frac{1}{2^n} = \left\| \mathcal{W} \circ \mathcal{T} \circ \mathcal{B}_m \left(|z_n\rangle\langle z_n| - \frac{\mathbb{1}^{\otimes n}}{2^n} \right) \right\|_2^2 \quad (79)$$

$$= \left\| \mathcal{W} \circ \mathcal{T} \left(\frac{1}{2^n} \sum_{S_A \otimes S_B \in \text{Bell}/\{\mathbb{1}^{\otimes n}\}} (-1)^{f(z_n, S_A, S_B)} S_A \otimes S_B \right) \right\|_2^2 \quad (80)$$

$$\leq (1-p)^{L_{\mathcal{T}}} \left\| \frac{1}{2^n} \sum_{S_A \otimes S_B \in \text{Bell}/\{\mathbb{1}^{\otimes n}\}} (-1)^{f(z_n, S_A, S_B)} S_A \otimes S_B \right\|_2^2 \quad (81)$$

$$= (1-p)^{L_{\mathcal{T}}} \left\| \mathcal{B}_m \left(|z_n\rangle\langle z_n| - \frac{\mathbb{1}^{\otimes n}}{2^n} \right) \right\|_2^2 \quad (82)$$

$$= (1-p)^{L_{\mathcal{T}}} \left\| |z_n\rangle\langle z_n| - \frac{\mathbb{1}^{\otimes n}}{2^n} \right\|_2^2 \quad (83)$$

$$= (1-p)^{L_{\mathcal{T}}} \left(P(|z_n\rangle\langle z_n|) - \frac{1}{2^n} \right), \quad (84)$$

where $L_{\mathcal{T}}$ is defined in Eq. (62) and $f(z_n, S_A, S_B)$ is a function that generates the appropriate phase. In the first equality we have used Lemma 1, in the second equality we have used Eqs. (64)-(67). The inequality comes from Eqs. (77) and (78), and the unitary invariance of the Schatten norms. In the last three equalities, we again use Eqs. (64)-(67), the unitary invariance of the Schatten norms, and Lemma 1. $\square$

Using Lemma 7 we can now prove our proposition on the concentration of output distributions in circuits with entangled measurements.

Proposition 2. *Consider the circuit in Figure 8(a) with computational basis measurement $O_{z_n} = |z_n\rangle\langle z_n|$; $z_n \in \{0, 1\}^n$. For any input state ρ and choice of measurement string z_n , the expectation value concentrates as*

$$\text{Tr}[\mathcal{B}_m^\dagger \circ \mathcal{T} \circ \mathcal{W}(\rho) O_{z_n}] - \frac{1}{2^n} \leq (1-p)^{L_{\mathcal{T}}} \sqrt{P(\rho)}, \quad (85)$$

where $P(\rho)$ is the purity of ρ and

$$L_{\mathcal{T}} = \begin{cases} L_{\mathcal{U}}, & \text{for } \mathcal{T} = \tilde{\mathcal{U}} \otimes \mathcal{V}, \\ L_{\mathcal{V}}, & \text{for } \mathcal{T} = \mathcal{U} \otimes \tilde{\mathcal{V}}, \\ L_{\mathcal{U}} + L_{\mathcal{V}}, & \text{for } \mathcal{T} = \tilde{\mathcal{U}} \otimes \tilde{\mathcal{V}}. \end{cases} \quad (86)$$

Proof. In similar steps to the proof of Proposition 1

$$\text{Tr}[\mathcal{B}_m^\dagger \circ \mathcal{T} \circ \mathcal{W}(\rho) O_{z_n}] - \frac{1}{2^n} = \text{Tr}[\mathcal{B}_m^\dagger \circ \mathcal{T} \circ \mathcal{W}(\rho) O_{z_n}] - \text{Tr}[\mathcal{B}_m^\dagger \circ \mathcal{T} \circ \mathcal{W}(\rho) \frac{\mathbb{1}}{2^n}] \quad (87)$$

$$= \text{Tr}[\mathcal{B}_m^\dagger \circ \mathcal{T} \circ \mathcal{W}(\rho) \left(O_{z_n} - \frac{\mathbb{1}}{2^n} \right)] \quad (88)$$

$$\leq \|\rho\|_2 \left\| \mathcal{W}^\dagger \circ \mathcal{T}^\dagger \circ \mathcal{B}_m \left(O_{z_n} - \frac{\mathbb{1}}{2^n} \right) \right\|_2 \quad (89)$$

$$= \sqrt{P(\rho)} \sqrt{P(\mathcal{W}^\dagger \circ \mathcal{T}^\dagger \circ \mathcal{B}_m(|z_n\rangle\langle z_n|)) - \frac{1}{2^n}} \quad (90)$$

$$\leq \sqrt{P(\rho)} (1-p)^{L_{\mathcal{T}}}, \quad (91)$$

where in the first equality we have used the fact that $\text{Tr}[\mathcal{B}_m^\dagger \circ \mathcal{T} \circ \mathcal{W}(\rho)] = 1$, the second equality is simply a grouping of terms, the first inequality is due to the tracial matrix Hölder's inequality [51], the third equality is an application of Lemma 1, and the final inequality is an application of Lemma 7. $\square$

A.5 Gradient scaling

In this section, we show how to transport our results to consider gradient scaling. We modify the settings studied in Section 3 to include a trainable unitary in the middle of the circuit. In doing so, we constrain ourselves to consider only computational basis input states.




Figure 10: Modification to the circuit in Figure 2 to include trainable unitary $Y(\theta)$.

We consider a trainable unitary of the form

$$Y(\theta) = W_0 \prod_{k=1}^K e^{-i\theta_k H_k} W_k, \quad (92)$$

where $\{W_k\}_{k=0}^K$ are arbitrary fixed unitary operators and $\{H_k\}_{k=0}^K$ are Hermitian operators. We denote the channel that corresponds to this unitary as $\mathcal{Y}(\theta)$. We modify the circuit in Figure 2 to include this trainable unitary, which we present in Figure 10. In the following proposition, we show how partial derivatives of parameters in this unitary exponentially vanish in the circuit depth under a CNOT ladder circuit.

Proposition 3. *Consider a cost function corresponding to the circuit in Figure 10 which we denote as $C(\theta) = \text{Tr}[\mathcal{W}_{L_2}^{(n_d)} \circ \mathcal{Y}(\theta) \circ (\mathcal{W}_{L_1}^{(n_d)})^\dagger(\rho) O_{z_n}]$, where ρ is a computational basis input state, $\mathcal{W}_{L_i}^{(n_d)}$ denotes the channel corresponding to L_i instances of CNOT ladders and local noise on n_d qubits as defined in Section A.3, and with computational basis measurement $O_{z_n} = |z_n\rangle\langle z_n|$; $z_n \in \{0, 1\}^n$. The partial derivative with respect to parameter θ_k is bounded as*

$$|\partial_{\theta_k} C| \leq (1-p)^{\Gamma_{n,n_d,L_1} + \Gamma_{n,n_d,L_2}} \|H_k\|_\infty, \quad (93)$$

where we denote

$$\Gamma_{n,n_d,L_i} = \begin{cases} \lfloor L_i \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \rfloor, & \text{for } n_d = 1, \\ \lfloor L_i \cdot \frac{n_d}{2^{\lceil \log_2 n \rceil}} \rfloor, & \text{for } n_d > 1, \end{cases} \quad (94)$$

for $i \in \{1, 2\}$.

Proof. We have

$$|\partial_{\theta_k} C| = \left| \text{Tr} \left[\mathcal{W}_{L_2}^{(n_d)} \circ \mathcal{Y}_-(\theta) \partial_{\theta_k} \mathcal{Y}_+(\theta) \circ (\mathcal{W}_{L_1}^{(n_d)})^\dagger(\rho) O_{z_n} \right] \right| \quad (95)$$

$$\leq \left\| \mathcal{Y}_{k-}^\dagger(\theta) \circ (\mathcal{W}_{L_2}^{(n_d)})^\dagger(O_{z_n}) \right\|_2 \left\| \partial_{\theta_k} \mathcal{Y}_{k+}(\theta) \circ (\mathcal{W}_{L_1}^{(n_d)})^\dagger(\rho) \right\|_2, \quad (96)$$

where the inequality is due to Hölder's tracial matrix inequality [51], and we denote $\mathcal{Y}_{k-}^\dagger(\theta)$ and $\mathcal{Y}_{k+}^\dagger(\theta)$ as the channels corresponding to unitary operators

$$Y_{k-}(\theta) = W_0 \prod_{j=1}^k e^{-i\theta_j H_j} W_j, \quad Y_{k+}(\theta) = \prod_{j=k}^K e^{-i\theta_j H_j} W_j, \quad (97)$$

respectively. The first term in Eq. (95) can be bounded by considering the unitary invariance of Schatten norms and by using Lemma 6 to obtain

$$\left\| \mathcal{Y}_{k-}^\dagger(\theta) \circ (\mathcal{W}_{L_2}^{(n_d)})^\dagger(O_{z_n}) \right\|_2 \leq (1-p)^{\Gamma_{n,n_d,L_2}} \sqrt{P(O_{z_n})} \quad (98)$$

$$= (1-p)^{\Gamma_{n,n_d,L_2}}, \quad (99)$$




Figure 11: Modification to the circuit in Figure 3 to include trainable unitary $Y(\theta)$.

where in the last line we have used the fact that the purity of pure states is equal to 1. The second term in Eq. (95) can be bounded as

$$\left\| \partial_{\theta_k} \mathcal{Y}_{k+}(\theta) \circ (\mathcal{W}_{L_1}^{(n_d)})^\dagger(\rho) \right\|_2 = \left\| \left[H_k, \mathcal{Y}_{k+}(\theta) \circ (\mathcal{W}_{L_1}^{(n_d)})^\dagger(\rho) \right] \right\|_2 \quad (100)$$

$$\leq \|H_k\|_\infty \left\| \mathcal{Y}_{k+}(\theta) \circ (\mathcal{W}_{L_1}^{(n_d)})^\dagger(\rho) \right\|_2 \quad (101)$$

$$\leq \|H_k\|_\infty (1-p)^{\Gamma_{n,n_d,L_1}}, \quad (102)$$

where the equality comes by directly evaluating the partial derivative, the first inequality comes from the application of Lemma 2, and the second inequality is an application of Lemma 6. Substituting Eqs. (99) and (102) into Eq. (95) we obtain the gradient scaling result Eq. (93) as desired. $\square$

Proposition 3 shows that the partial derivatives with respect to any parameter θ_k exponentially vanishes in the circuit depth. This implies that so long as $\|H_k\|_\infty$ grows at most polynomially in the number of qubits n , the circuit in Figure 10 has an NIBP for linear circuit depths.

We can similarly modify the circuit in Figure 3 to include the trainable unitary $Y(\theta)$ and entangling gates at the start of the circuit. We present this circuit in Figure 11. In the following proposition, we show that partial derivatives of parameters in this unitary also exponentially vanish with increasing depth.

Proposition 4. Consider a cost function corresponding to the circuit in Figure 11 which we denote as $C(\theta) = \text{Tr}[\mathcal{B}_m^\dagger \circ \mathcal{T}_2 \circ \mathcal{Y}(\theta) \circ \mathcal{T}_1 \circ \mathcal{B}_m(\rho) O_{z_n}]$, where ρ is a computational basis input state, $\mathcal{B}_{m_i}$ is an entangling operation between m_i pairs of qubits in subsystems A_i and B_i as considered in Proposition 2, $\mathcal{T}_i$ denotes noisy unitary evolution separable across the cut $A_i|B_i$, and with computational basis measurement $O_{z_n} = |z_n\rangle\langle z_n|$; $z_n \in \{0,1\}^n$. The partial derivative with respect to parameter θ_k is bounded as

$$|\partial_{\theta_k} C| \leq (1-p)^{L_{\mathcal{T}_1} + L_{\mathcal{T}_2}} \|H_k\|_\infty, \quad (103)$$

where we denote

$$L_{\mathcal{T}_i} = \begin{cases} L_{\mathcal{U}_i}, & \text{for } \mathcal{T}_i = \tilde{\mathcal{U}}_i \otimes \mathcal{V}_i, \\ L_{\mathcal{V}_i}, & \text{for } \mathcal{T}_i = \mathcal{U}_i \otimes \tilde{\mathcal{V}}_i, \\ L_{\mathcal{U}_i} + L_{\mathcal{V}_i}, & \text{for } \mathcal{T}_i = \tilde{\mathcal{U}}_i \otimes \tilde{\mathcal{V}}_i, \end{cases} \quad (104)$$

for $i \in \{1, 2\}$, and $L_{\mathcal{U}_i}$ and $L_{\mathcal{V}_i}$ are defined in the same manner as $L_{\mathcal{U}}$ and $L_{\mathcal{V}}$ in Section A.4.

Proof. Similar to the proof of Proposition 3, we have

$$|\partial_{\theta_k} C| = \left| \text{Tr} \left[\mathcal{B}_m^\dagger \circ \mathcal{T}_2 \circ \mathcal{Y}_-(\theta) \partial_{\theta_k} \mathcal{Y}_+ \circ \mathcal{T}_1 \circ \mathcal{B}_m(\rho) O_{z_n} \right] \right| \quad (105)$$

$$\leq \left\| \mathcal{Y}_{k-}^\dagger(\theta) \circ \mathcal{T}_2^\dagger \circ \mathcal{B}_m(O_{z_n}) \right\|_2 \left\| \partial_{\theta_k} \mathcal{Y}_{k+}(\theta) \circ \mathcal{T}_1 \circ \mathcal{B}_m(\rho) \right\|_2, \quad (106)$$

where the inequality is due to Hölder's tracial matrix inequality [51]. The first term is bounded as

$$\left\| \mathcal{Y}_{k-}^\dagger(\boldsymbol{\theta}) \circ (\mathcal{W}_{L_2}^{(n_d)})^\dagger(O_{z_n}) \right\|_2 \leq (1-p)^{\mathcal{T}_2}, \quad (107)$$

where we have used Lemma 7, and the fact that the lemma gives the same result if we consider the adjoint map of $\mathcal{T}_2$. As in the proof of Proposition 3, the second term in Eq. 105 is bounded as

$$\left\| \partial_{\theta_k} \mathcal{Y}_{k+}(\boldsymbol{\theta}) \circ \mathcal{T}_1 \circ \mathcal{B}_m(\rho) \right\|_2 = \left\| [H_k, \mathcal{Y}_{k+}(\boldsymbol{\theta}) \circ \mathcal{T}_1 \circ \mathcal{B}_m(\rho)] \right\|_2 \quad (108)$$

$$\leq \|H_k\|_\infty \left\| \mathcal{Y}_{k+}(\boldsymbol{\theta}) \circ \mathcal{T}_1 \circ \mathcal{B}_m(\rho) \right\|_2 \quad (109)$$

$$\leq \|H_k\|_\infty (1-p)^{\mathcal{T}_1}, \quad (110)$$

where the equality comes by directly evaluating the partial derivative, the first inequality comes from the application of Lemma 2, and the second inequality is an application of Lemma 7. Substituting Eqs. (107) and (110) into Eq. (105) we obtain the gradient scaling result Eq. (103) as required. $\square$

B Numerical gradient computation for the HVA ansatz

To compute the gradient with respect to θ_k corresponding to $e^{-i\theta_k H_{XX}}$ in the ansatz (14) we introduce "dummy" variables $\tilde{\theta}_{k,1}, \tilde{\theta}_{k,2}, \dots, \tilde{\theta}_{k,n}$ such that

$$e^{-i\theta_k H_{XX}} = e^{-i\tilde{\theta}_{k,1} X_1 X_2} e^{-i\tilde{\theta}_{k,2} X_2 X_3} \dots e^{-i\tilde{\theta}_{k,n} X_n X_1} \big|_{\tilde{\theta}_{k,1}=\tilde{\theta}_{k,2}=\dots=\tilde{\theta}_{k,n}=\theta_k}. \quad (111)$$

We have

$$\partial_{\theta_k} C(\theta_1, \dots, \theta_{k-1}, \theta_k, \theta_{k+1}, \dots, \theta_{2L}) = \sum_{i=1, \dots, n} \partial_{\tilde{\theta}_{k,i}} C(\theta_1, \dots, \theta_{k-1}, \tilde{\theta}_{k,1}, \tilde{\theta}_{k,2}, \dots, \tilde{\theta}_{k,n}, \theta_{k+1}, \dots, \theta_{2L}) \big|_{\tilde{\theta}_{k,1}=\tilde{\theta}_{k,2}=\dots=\tilde{\theta}_{k,n}=\theta_k}. \quad (112)$$

Using a parameter shift rule for Pauli strings [46] leads to

$$\begin{aligned} & \partial_{\tilde{\theta}_{k,1}} C(\theta_1, \dots, \theta_{k-1}, \tilde{\theta}_{k,1}, \tilde{\theta}_{k,2}, \dots, \tilde{\theta}_{k,n}, \theta_{k+1}, \dots, \theta_{2L}) \big|_{\tilde{\theta}_{k,1}=\tilde{\theta}_{k,2}=\dots=\tilde{\theta}_{k,n}=\theta_k} = \\ & C(\theta_1, \dots, \theta_{k-1}, \theta_k + \pi/4, \theta_k, \dots, \theta_k, \theta_{k+1}, \dots, \theta_{2L}) - C(\theta_1, \dots, \theta_{k-1}, \theta_k - \pi/4, \theta_k, \dots, \theta_k, \theta_{k+1}, \dots, \theta_{2L}). \end{aligned} \quad (113)$$

Analogously we compute $\partial_{\tilde{\theta}_{k,2}}, \dots, \partial_{\tilde{\theta}_{k,n}}$. We treat similarly $\partial_{\theta_k} C$ corresponding to $e^{-i\theta_k H_Z}$ in the ansatz (14).

C Noise models

In our numerical studies, we use the local depolarizing noise model and a noise model of a near-term trapped-ion quantum computer developed by Trout *et al.* [42]. We assume single-qubit rotations $R_X(\theta) = e^{-i\theta/2X}$, $R_Y(\theta) = e^{-i\theta/2Y}$, $R_Z(\theta) = e^{-i\theta/2Z}$, and the two-qubit Molmer-Sørensen gate $XX(\theta) \equiv e^{-i\theta X_i \otimes X_j}$ as native gates of the quantum computer. Here, X , Y , and Z are Pauli matrices.

C.1 Local depolarizing noise

The local depolarizing noise channel with the error rate p is

$$\mathcal{D}_p : \rho \mapsto (1-p)\rho + \frac{p}{3}(X^\dagger \rho X + Y^\dagger \rho Y + Z^\dagger \rho Z). \quad (114)$$

This is consistent with the definition in 18 and is included here for the reader's convenience. We choose $p = 2.425 \times 10^{-3}$ for simulations of the clean and dirty setup and $p = 2.425 \cdot 10^{-3} f$ where $f = 1/n, 2/n, \dots, (n-1)/n$ for the simulations with all dirty qubits and variable error rates. The error channels are applied to the dirty qubits after each layer of non-parallelizable gates as shown in Figure 12.




Figure 12: **Local depolarizing noise model.** A 4-qubit HVA circuit with the local depolarizing noise and 2 dirty qubits. The depolarizing noise channel $\mathcal{D}_p$ is applied to all the dirty qubits after each layer of non-parallelizable gates.

C.2 Trapped-ion noise model.

The realistic noise model is based on Trout *et al.* [42]. The noisy quantum gates R_X , R_Y , R_Z and XX are defined with noise channels $\mathcal{S}_X$, $\mathcal{S}_Y$, $\mathcal{S}_Z$ and $\mathcal{C}$, respectively:

$$\begin{aligned}\mathcal{S}_X &= \mathcal{W}_{p_d} \circ \mathcal{D}_{p_{\text{dep}}} \circ \mathcal{R}_X^{p_\alpha}, \\ \mathcal{S}_Y &= \mathcal{W}_{p_d} \circ \mathcal{D}_{p_{\text{dep}}} \circ \mathcal{R}_Y^{p_\alpha}, \\ \mathcal{S}_Z &= \mathcal{W}_{p_d} \circ \mathcal{D}_{p_{\text{dep}}} \circ \mathcal{R}_Z^{p_\alpha}, \\ \mathcal{C} &= [\mathcal{W}_1^{p_{d1}} \otimes \mathcal{W}_2^{p_{d2}}] \circ [\mathcal{D}_1^{p_{\text{dep}}} \otimes \mathcal{D}_2^{p_{\text{dep}}}] \circ \mathcal{H}_{p_{\text{xx}}} \circ \mathcal{H}_{p_h}.\end{aligned}\tag{115}$$

Here, $\mathcal{R}_X^{p_\alpha} \rho = (1 - p_\alpha) \rho + p_\alpha X \rho X$, $\mathcal{R}_Y^{p_\alpha} \rho = (1 - p_\alpha) \rho + p_\alpha Y \rho Y$, $\mathcal{R}_Z^{p_\alpha} \rho = (1 - p_\alpha) \rho + p_\alpha Z \rho Z$ are imprecisions of the angle of rotation about the axes X, Y, Z , respectively. $\mathcal{D}_{p_{\text{dep}}}$ is the local depolarizing channel (114) with the error rate p_{dep} , $\mathcal{W}_{p_d} \rho = (1 - p_d) \rho + Z \rho Z$ is the dephasing channel, and $\mathcal{H}_p = (1 - p) \rho + X_1 X_2 \rho X_1 X_2$ is a two qubit channel representing both imprecise rotation when $p = p_{\text{xx}}$ and ion heating when $p = p_h$. The indices 1, 2 refer to two qubits at which acts XX . Furthermore, for an idling dirty qubit, we apply to it $\mathcal{D}(p_{\text{idle}})$. In the computations, the error rates were chosen as

$$\begin{aligned}p_d &= 1.5 \cdot 10^{-4} f, \quad p_{\text{dep}} = 8 \cdot 10^{-4} f, \\ p_{d1} &= p_{d2} = 7.5 \cdot 10^{-4} f, \quad p_\alpha = 1 \cdot 10^{-4} f, \\ p_{\text{xx}} &= 1 \cdot 10^{-3} f, \quad p_h = 1.25 \cdot 10^{-3} f,\end{aligned}\tag{116}$$

with $f = 1$ for the clean and dirty setup and $f = 1/n, 2/n, \dots, (n-1)/n$ for all dirty qubits with the variable error rates.




Figure 13: **Trapped-ion, realistic noise model.** A 4-qubit HVA circuit with realistic, trapped-ion noise and 2 dirty qubits. When a two-qubit gate is applied between two dirty qubits, it is followed by a 2-qubit noise channel $\mathcal{C}$. All the R_Z gates acting at the dirty qubits are followed by their noise channels $\mathcal{S}_Z$. If a two-qubit XX gate acts in between a dirty and a clean qubit we assume that it is noiseless.

In the case of a XX gate acting at a dirty and a clean qubit, we assume that it is a noiseless gate. In practice modeling an XX gate in between an error-corrected qubit and a dirty qubit requires taking into account details of the error correction code. This is outside of the scope of this work, so we assume no error for such a gate. This assumption allows us to upper bound performance of a quantum computer with quantum error-corrected qubits. Figure 13 shows how the noise channels are applied for our HVA simulations.

In the case of our HVA simulations, the choice of perfect XX gates in between dirty and clean qubits causes deviations from a scaling of the gradient magnitude with the n_d/n ratio. For our choice of the dirty qubits forming a contiguous block of qubits, the number of the noisy XX gates is $L(n_d - 1)$ for $1 \leq n_d \leq n - 1$ and Ln_d for $n_d = 0, n$. Therefore, a fraction of the noisy two-qubit gates in the circuit scales as $(n_d - 1)/n$ for $1 \leq n_d \leq n - 1$ causing the gradient magnitude to scale with $(n_d - 1)/n$ rather than n_d/n as can be seen in Figure 6b.

D Additional numerical results

Here we show results for the gradient behavior in the case of 4 and 6-qubit HVA circuits simulated with the realistic and depolarizing noise models. Figure 14 shows the mean absolute value of the derivative of the cost function with respect to a parameter $|\partial_{\theta_k} C|$ averaged over both 28 instances of random HVA parameters and the HVA parameters. plotted versus depth of the ansatz for the depolarizing and realistic noise models. The behavior is very similar to that of the 8-qubit HVA shown in Figure 5.

Figure 15 displays $|\partial_{\theta_k} C|$ versus the total error rate for the 4 and 6-qubit HVA circuits. We find very good quality of the collapse similar to the 8-qubit case shown in Figure 6. We also observe that for the depolarizing noise small deviations from the perfect collapse are growing with increasing n . We leave further investigation of this phenomenon to future work.

Figure 16 shows the behavior of $|\partial_{\theta_k} C|$ for lower layer numbers $L = 1 - 30$, $n = 4, 6, 8, 10$, and the realistic noise model. Again for each triple of n, n_d, L values we use 28 random HVA instances to estimate $|\partial_{\theta_k} C|$. Apart from the very low $L \leq n$, the behavior of the derivative depends primarily on the n_d/n ratio displaying a decay with an increasing n_d similar to the larger L case. A different behavior observed for very small L is not surprising as for such shallow circuits one can expect dependence of the impact of the errors to be strongly localized due to the causal cone effects, unlike for the deeper circuits. Furthermore, we see in that regime strong dependence of the results on L that overshadows subtler n_d effects. For $n_d > 0$ and larger L , the decay of the gradient with L is approximately exponential, though this behavior is less clear than for larger L due to seemingly random fluctuations superimposed on the dominant trend, which might be finite sample effects.

E Repository

We have prepared a repository which can be found at https://github.com/danielbultrini/Clean_Dirty_Qubits. This contains data frames with numerical results and functions required to plot all the figures shown. In addition, to reproduce the depolarizing noise model results, an open-source Qiskit-based code is provided. This code is not optimized but does implement the depolarizing clean and dirty model. Its aim is to make it easier for interested readers to expand or alter it for their own purposes.




Figure 14: **HVA gradient scaling with respect to the ansatz depth.** Here we show the mean absolute value of the partial derivative of the cost function with respect to a parameter $|\partial_{\theta_k} C|$ for 4 (a,b) and 6-qubit (c,d) systems. We plot the results for both the clean and dirty and the variable noise setups in the same way as in Figure 5. As for the 8-qubit case we find that $|\partial_{\theta_k} C|$ decays exponentially with increasing L . Furthermore, we observe that $|\partial_{\theta_k} C|$ decays exponentially with increasing n_d , similarly as for the 8-qubit HVA circuits.







Figure 15: **HVA $|\partial_{\theta_k} C|$ scaling with respect to the total error rate.** We plot here the mean absolute value of the partial derivative of the cost function from Figure 14 versus the total error rate. For the local depolarizing noise the total error rate is proportional to $L_{\text{eff}} = L \frac{n_d}{n}$. As for the 8-qubit case (see Figure 6) we observe very good collapse quality.




(a) Realistic noise model - 4 Qubits



(b) Realistic noise model - 6 Qubits



(c) Realistic noise model - 8 Qubits



(d) Realistic noise model - 10 Qubits

Figure 16: **HVA $|\partial_{\theta_k} C|$ for 4,6 8, and 10 qubits at low layer depths.** Here we have plotted the realistic noise model simulations for qubit counts 4 in (a), 6 in (b), 8 in (c), and 10 in (d) at depths of up to 30 layers, computed at each layer count. We see that the general pattern seen in the main text holds, with higher noisy qubit counts decaying faster than lower ones.