Polynomial-Time Classical Simulation of Noisy Quantum Circuits with Naturally Fault-Tolerant Gates

Jon Nelson*[1,2], Joel Rajakumar*[1,2], Dominik Hangleiter[3,1], and Michael J. Gullans[1,2]

[1]*Joint Center for Quantum Information & Computer Science, University of Maryland and NIST*
[2]*Department of Computer Science, University of Maryland*
[3]*Simons Institute for the Theory of Computing, University of California at Berkeley*
[4]*National Institute of Standards and Technology (NIST)*

January 9, 2026

Abstract

We construct a polynomial-time classical algorithm that samples from the output distribution of noisy geometrically local Clifford circuits with any product-state input and single-qubit measurements in any basis. Our results apply to circuits with nearest-neighbor gates on an $O(1)$-D architecture with depolarizing noise after each gate. Importantly, we assume that the circuit does not contain qubit resets or mid-circuit measurements. This class of circuits includes Clifford-magic circuits and Conjugated-Clifford circuits, which are important candidates for demonstrating quantum advantage using non-universal gates. Additionally, our results can be extended to the case of IQP circuits augmented with CNOT gates, which is another class of non-universal circuits that are relevant to current experiments. Importantly, these results do not require randomness assumptions over the circuit families considered (such as anticoncentration properties) and instead hold for *every* circuit in each class as long as the depth is above a constant threshold. This allows us to rule out the possibility of fault-tolerance in these circuit models. As a key technical step, we prove that interspersed noise causes a decay of long-range entanglement at depths beyond a critical threshold. To prove our results, we merge techniques from percolation theory and Pauli path analysis.

1 Introduction

A first step towards understanding the power of quantum computation is to determine the conditions under which it is possible to perform a quantum computation that cannot be classically simulated or, in other words, to demonstrate "quantum advantage". Although there is robust theoretical evidence that this is true for large-scale fault-tolerant quantum computers, this becomes a subtle question when restricted to near-term quantum hardware. These devices are noisy and may lack capabilities that are required for fault tolerance such as the ability to perform intermediate measurements or reset qubits during a computation.

To address these questions, a line of research has considered the task of sampling from the output distribution of (random) quantum circuits as a way to demonstrate quantum advantage [TD04]. This is relevant to near-term hardware because even very restricted models of quantum computation can be used to demonstrate such quantum advantages, and the advantage is not contrived in that it tolerates a constant amount of total error [BJS10, BMS16, AA13, GWD17, HKS+17]. Adding realistic noise with constant rate

*These authors contributed equally to this work

complicates this picture, since the fidelity with which the targeted task can be performed on a noisy quantum device diminishes exponentially with the system size. It is therefore important to study the computational power of quantum computations under realistic noise assumptions [ABOIN96, FT16, BMS17, AGL+23, CCHL22]. Obtaining a precise understanding of the computational power of noisy restricted models of computation can elucidate the ingredients that are necessary for near-term quantum advantage.

In this work, we consider quantum circuits with gate sets that are naturally fault-tolerant. These include Clifford circuits with state preparation and measurement in arbitrary product bases, and IQP circuits augmented with CNOT gates. In particular, this allows for perfect magic state inputs in the case of Clifford circuits and arbitrary diagonal non-Clifford gates in the case of IQP+CNOT circuits. The circuit classes we consider do not allow for intermediate measurement and classical feed-forward and are therefore not universal for quantum computation; however, they constitute circuits which are readily implementable in early-stage fault-tolerant quantum devices, which can execute non-universal gate sets transversally [BEG+24, EK09], making them good candidates for demonstration of quantum advantage. In particular, these circuit classes include Clifford-magic, conjugated Clifford circuits and hypercube block IQP circuits, which have all been proven to be hard to approximately sample from in the nearly noiseless case assuming certain complexity-theoretic conjectures [YJS19, BFK18, HKB+24, PBG20]. Thus, it is important to understand the noise and depth regimes in which these circuits maintain their computational power.

We develop a classical algorithm for sampling from the output of these circuit families when they are subject to mid-circuit depolarizing noise. Applied to geometrically local circuits, our algorithm is efficient at circuit depths above a *constant threshold* that depends on the noise rate and is efficient for *all* circuits in this class. Our results demonstrate that in order to achieve scalable quantum advantage with realistically noisy circuits, an additional ingredient is required such as a universal set of mid-circuit gates, intermediate measurements, or a supply of fresh qubits.

1.1 Overview of Results

We summarize our main result informally in the following theorem.

Theorem 1 (Informal)**.** *There exists an efficient randomized classical algorithm that approximately samples from the output distribution of a noisy quantum circuit C with circuit-level local noise rate γ in the following cases,*

1. *C is a geometrically local Clifford-Magic, Conjugated Clifford, or IQP+CNOT circuit and $d \geq \Omega(\gamma^{-1} \log \gamma^{-1})$*
2. *C is a Clifford-Magic, Conjugated Clifford, or IQP+CNOT circuit and $d \geq \Omega(\gamma^{-1} \log n)$.*

We use the term 'geometrically local' to refer to circuits whose gates are nearest-neighbor on an $O(1)$-D lattice. Our algorithm has many desirable properties: (1) it does not require assumptions on the output distribution or distribution from which the circuit is sampled. In particular, in contrast to the results of [BMS17, GD18, AGL+23, SYGY24] it does not require that the output distributions have the anticoncentration property (see [HE23] for details on this property) and it applies to *worst-case* circuits. (2) It performs *exact* sampling with random runtime which is polynomial in expectation (and this can be modified to give the more common *approximate* sampler with worst-case polynomial runtime). (3) Only the depth threshold but not the runtime depends on the noise strength.

Next, our proof techniques also show that noisy Clifford circuits acting on a random input bit string anticoncentrate in $O(\gamma^{-1} \log n)$ depth, which is proved in Section D. This implies that noisy random quantum circuits on arbitrary architectures anticoncentrate in $O(\log n)$-depth, which was only previously known for $1D$ and all-to-all connectivity [DHB22, DNS+22] [1]. Our proof is conceptually significant because it

[1] recent concurrent work [SHH24] also addresses this gap for the noiseless case

demonstrates that anticoncentration can arise from the noise alone, rather than due to the randomness of the gates (in the noiseless setting, these circuits can be non-anticoncentrated).

1.2 Related Works

The strongest upper bound on the computational complexity of worst-case noisy circuits is due to Aharonov *et al.* [ABOIN96], who showed that the output distribution of any noisy quantum circuit converges to uniform after circuit depth which is logarithmic in system size, and is thus trivially simulatable. The same work [ABOIN96] also demonstrates a lower bound on the computational complexity of general quantum circuits by proving that noisy quantum circuits can solve problems in QNC^1 (decision problems solvable by noiseless logarithmic depth quantum circuits) with a quasipolynomial overhead in system size. At the same time, Shor's factoring algorithm can be implemented using log-depth quantum circuits [CW00]. Together, these results provide some evidence that it is difficult in general to classically simulate noisy quantum circuits at logarithmic depths. This suggests to start by understanding the classical simulatability of restricted families of noisy quantum circuits that correspond to the limitations of near-term hardware, such as geometric locality or limited gate sets.

Various prior works have provided classical simulation algorithms for the setting of noisy Clifford circuits with single-qubit non-Clifford gates [FT16, VHP05, KRUdW08, PV10, SRP+21, GGCT24]. When applied to Clifford-Magic or Conjugated Clifford circuits, these algorithms are only efficient when the depolarizing noise rate is above $\sim 15\%$. Our algorithm expands this classical simulability regime by showing that circuits with *any* constant noise rate are efficiently simulable as long as the depth is above a constant threshold. Our algorithm is also conceptually quite different. In prior approaches, the noise converts a circuit containing non-Clifford ('magic') gates into a probabilistic mixture of Clifford circuits, which can then be simulated using the Gottesman-Knill theorem. In contrast, our algorithm leverages the loss of entanglement, rather than the loss of magic, as the key to classical simulatability. Specifically, the noise transforms a circuit with entangling gates into a mixture of circuits with a tensor-product form, which can be simulated via brute-force methods [2]. Furthermore, we carefully account for the accumulation of noise in the circuit, and find that increasing the depth reduces the entanglement even further, allowing for classical simulability at lower noise strengths. At first glance, this is somewhat unintuitive because one might expect higher depth circuits to possess more entanglement and to be more robust to noise. Our results demonstrate an important idea: as depth increases, the decohering effects of accumulating noise always win out against the entangling/fault-tolerance abilities of quantum gates.

A rich landscape of existing work has also considered noisy circuits where the gates are chosen randomly [AGL+23, BMS17, TTT21, GD18, MAG+24, SYGY24, FRD+23]. The accuracy of these algorithms strongly relies on the randomness property and thus does not exclude the possibility of a *worst-case* circuit which is cleverly designed to maintain robustness to noise (e.g. [ABOIN96, FT16]). Therefore, these results do not have implications for early fault-tolerance. The work of [GGCT24] and [SFG21] consider worst-case circuits; however, these results are incomparable since they only apply in the easier setting of estimating expectation values rather than sampling from the output distribution.

The only prior work achieving both properties (worst-case and sampling before $\log(n)$ depth) is Refs. [RWL24, Oh24], which take advantage of commutation properties of the gate set to prove their bounds. In contrast, our results and our proof techniques suggest that classical simulatability can arise from more general properties such as unitarity and geometric locality.

We also note that phase transitions in entanglement properties of quantum circuits due to noise/percolation have been observed in a variety of other settings (e.g. [SHEF25, Aha00, BEF+08, FT16]); however, a key

[2] Note, each circuit in the mixture can contain up to $\Omega(n)$ non-Clifford gates, which means Gottesman-Knill-type algorithms would likely fail unless the tensor product structure is exploited.

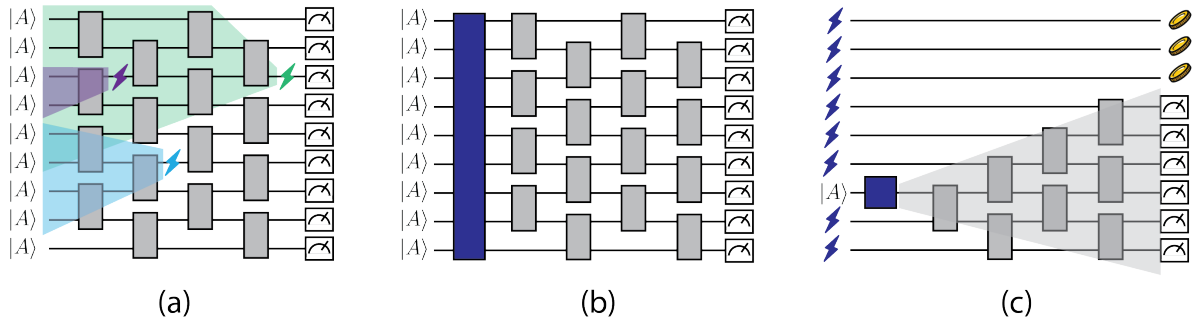

Figure 1: Overview of the simulation algorithm. (a) Depolarizing errors are sampled and propagated to the beginning of the circuit. $|A\rangle$ represents a single-qubit magic-state but can be any single-qubit state. (b) The circuit can now equivalently be represented as one error channel followed by the original noiseless circuit. (c) We prove that this input error channel has the effect of depolarizing many of the input qubits and so it remains to simulate the lightcones of qubits that are not depolarized. When these lightcones intersect they must be simulated together but can otherwise be simulated independently. Any measurements that are not in the lightcone of a depolarized qubit can be simulated by a random coin flip.

distinction in our results is that we demonstrate a phase transition in depth for any noise strength $\gamma = \Omega(1)$ (not just noise which is greater than a critical threshold).

1.3 Proof Techniques

At a high-level, the algorithm works by first propagating errors to the beginning of the circuit. This reformulates the noisy circuit as one layer of noise followed by the ideal circuit, which is pictured in Fig. 1b. The focus of much of the technical work is to show that this propagated error channel (which we denote by the channel Π) effectively depolarizes many of the input qubits, which is pictured in Fig. 1c. This is similar to [RWL24, Oh24]; however, we require new tools to deal with the fact that the noise spreads out non-trivially, rather than commuting through the gates. In particular, we analyze which input Pauli operators can survive the noise. First, we show that

> (1) It is exponentially unlikely (in the depth and the number of qubits) that any input Pauli operator acting on a large island of qubits survives Π.

We prove this statement using a careful counting argument, in place of assuming anticoncentration as done in prior work [BMS17, GD18, AGL+23]. While the Cliffordness of the circuit simplifies the Pauli propagation significantly, the main ingredient of the proof is the more general property of *unitarity* to argue that not all high-weight Pauli operators can be mapped by the circuit to low-weight operators since this is a strictly smaller space and thus violates unitarity. This suggests that high-weight Pauli operators maintain their approximate weight throughout the circuit and thus are evenly suppressed by errors at various timesteps, i.e. they experience an exponential decay in both depth and weight. The next step of the analysis requires the depth to be large enough so that each Pauli operator experiences enough noise to suppress the exponential number of possible high-weight Pauli operators within an island of qubits. It turns out this occurs at a constant depth of around $d^* = O(\gamma^{-1} \log \gamma^{-1})$, which can be understand as a percolation effect that occurs due to *geometric locality*. We then use the probabilistic method to show that

> (2) If a large island of qubits was not depolarized by Π, then there must be at least one Pauli operator with support on these qubits, which survives Π.

Since we proved in (1) that the survival of large input Pauli operators is extremely unlikely under Π, then it must also be true by (2) that no large islands of qubits remain non-depolarized by Π.

Having proved this result on the entanglement structure of the noisy state, we can now simulate the circuit by considering each island of qubits separately. If we were to simulate these qubits by brute-force, then this would only allow for a quasipolynomial time algorithm when $D > 1$ for $\Omega(\log n)$-depth. However, we can improve this dramatically by decomposing the input state in the Pauli basis and only simulating Pauli operators that survive the channel - this is why our algorithm also works when there is all-to-all connectivity and no lightcones, although this requires the larger depth.

1.4 Discussion and Future Work

1.4.1 Naturally Fault-Tolerant Gate Sets

We consider the Clifford and IQP+CNOT gate sets to be 'naturally' fault-tolerant because they are transversal for certain quantum error-correcting codes such as 3D color codes [Bom15], and thus can be implemented fault-tolerantly with relative ease. These facts have been exploited in various experimental implementations [BEG+24, HKB+24]. In addition, Clifford circuits, in particular, are used in many important error correction gadgets. For instance, the encoding and syndrome extraction circuits for any stabilizer code can be implemented as a Clifford circuit. Furthermore, Pauli noise propagates predictably through Clifford gates, which allows for a decoding algorithm to track the error at various times in the circuit in order to diagnose the error. In fact, there have been a variety of proposals for early fault-tolerance [BGKT20, BL25, DP23] involving propagation of interspersed Pauli errors through Clifford circuits where they can be corrected at the end of the computation, provided that the final measurement contains some syndrome information about the space-time location of the error. This is a crucial part of the techniques in [BGKT20], which demonstrates quantum advantage over shallow classical circuits in the presence of noise. These desirable error correction properties have led to the proposal of using only Clifford gates on magic-state input for achieving fault-tolerant universal quantum computation [BK05]. The crucial difference in our setting is that we do not allow for feed-forward adaptive operations, which is responsible for the universality.

1.4.2 Noisy Quantum Advantage without Intermediate Measurement/Reset

Our results rule out certain proposals for quantum advantage in noisy circuits without intermediate measurement or reset operations (which would otherwise enable fault-tolerance through the threshold theorem). In particular, one might hope that it would be possible to construct some clever Clifford/IQP+CNOT circuit to demonstrate quantum advantage, where resource states are prepared perfectly, and interspersed errors are corrected using final measurement data (e.g. like [ALNP25] or [BGKT20]). However, our results imply that such circuits cannot be deeper than constant depth when they are implemented under geometric locality and gate set constraints.

1.4.3 Computational Complexity Phase Transition at Constant Depth

We further conjecture that the depth threshold for classical simulatability in the geometrically local case corresponds directly to a phase transition in computational complexity (as demonstrated in the IQP case in [RWL24]). That is, it may be possible to construct hard-to-sample, geometrically local quantum circuits at depths right below the threshold (at least for exact sampling). Indeed, a primary motivation for developing asymptotically efficient classical simulation algorithms is to narrow down the regime of near-term quantum advantage, and thereby point towards candidates for noise-robust quantum advantage. By process of elimination, our results point to the possibility that Clifford-magic/IQP+CNOT circuits could evade classical simulatability at depths below $O(1)$. Such a construction could have important consequences for demonstrating near-term quantum advantage. [BL25] provides some evidence in favor of this conjecture by demonstrating that constant-depth geometrically local noisy Clifford circuits possess long-range entanglement.

1.4.4 Peaked Quantum Circuits

Our work is among the first classical sampling algorithms for noisy quantum circuits that apply to worst-case circuits, albeit with restrictions on connectivity and gate sets. This is accomplished by removing the need for the anticoncentration property in our analysis. This not only allows for classical simulatability beyond average case, but also shows that our algorithm works even when the circuit output distribution is bounded away from uniform (the distributions we can simulate may be peaked). Peaked circuits have recently been proposed as a potential candidate for verifiable quantum advantage [AZ24] and are by definition not amenable to simulation techniques relying on anticoncentration. Our methods do not have the same fundamental restrictions and so we are able to rule out noisy quantum advantage for a certain class of candidate circuits with peaked distributions[3].

1.4.5 Generality of Percolation

Another ingredient in our simulation result that may apply more generally is the existence of a constant depth threshold for noise percolation in geometrically local circuits. In recent prior work, noise percolation was shown to exist in noisy IQP circuits [RWL24], which also resulted in classical simulation. For these IQP circuits, the analysis relies on the special property that commuting gates cannot spread entanglement very quickly (i.e. each qubit's lightcone grows linearly with depth) and that certain noise channels commute with the gate set. This may give reason to believe that this percolation effect is unique to IQP circuits. However, our results show that in more general settings, such as Clifford circuits, where the noise can spread in much more complicated ways, percolation still occurs. In these cases, the cause of percolation can be attributed instead to geometric locality, which may be a more fundamental limitation to near-term quantum circuits.

1.4.6 Current Experiments

Our results may also provide avenues to classically simulate real quantum experiments. For example, going beyond geometric locality, can we use percolation to prove a sublogarithmic depth threshold for classical simulatability on partially constrained architectures, such as the hypercube connectivity of [HKB$^+$24], which was recently implemented experimentally [BEG$^+$24]? Note that the IQP+CNOT circuits of [HKB$^+$24] require $O(\log n)$ depth. Our algorithm for the all-to-all connectivity case shows that these circuits will be classically simulatable if they are scaled past a $O(\gamma^{-1} \log n)$ depth (for noise strength γ). It would be of practical interest to see if this bound can be improved by considering the locality structure of the hypercube.

1.5 Paper Outline

In Section 2 we give the preliminaries and notation. In Section 3 we describe the algorithm. In Section 4, we describe how these results directly apply to IQP+CNOT circuits and to conjugated Clifford circuits.

2 Preliminaries

2.1 Pauli and Clifford Group

We define the Pauli group as:

$$\mathsf{P}_n := \{1, i, -1, -i\} \times \{\mathbb{I}, X, Y, Z\}^{\otimes n}$$

[3]Note that it is known that constant-depth peaked circuits can be simulated in quasipolynomial time [BGL23] but it is an open question whether they are simulatable at $O(\log n)$-depth and/or with noise.

We define the phaseless Pauli group $\hat{\mathsf{P}}_n \leqslant \mathsf{P}_n$ as:

$$\hat{\mathsf{P}}_n := \{1\} \times \{\mathbb{I}, X, Y, Z\}^{\otimes n}$$

Let P_n^Z denote the subgroup of P_n^Z where each single-qubit operator is in $\{\mathbb{I}, Z\}$. Let P_n^X be defined similarly. Let $X_i \in \mathsf{P}_n$ denote the n-qubit Pauli operator where X acts on qubit i and $\mathbb{I}$ acts elsewhere and define Z_i analogously. The Clifford group is defined as:

$$\mathsf{C}_n := \{U \in \mathsf{U}(2^n) : UPU^\dagger \in \mathsf{P}_n \forall P \in \mathsf{P}_n\}$$

where $\mathsf{U}(2^n)$ is the group of $2^n \times 2^n$ unitary matrices.

2.2 Binary symplectic representation and Tableau matrix

The binary symplectic representation is a mapping $v : \hat{\mathsf{P}}_n \to \mathbb{Z}_2^{2n}$ where the first n bits represent the X-components and the second n bits represent the Z-components of the corresponding Pauli operator. More precisely, for any $s \in \hat{\mathsf{P}}_n$, if the ith qubit of s contains X or Y then the ith bit of $v(s)$ is 1, otherwise it is 0. Similarly, if the ith qubit of s contains Z or Y then the $(i+n)$th bit of $v(s)$ is 1, otherwise it is 0. The symplectic inner product between two vectors u and w is defined as $u^T \Lambda w$, where

$$\Lambda = \begin{pmatrix} 0 & \mathbb{I}_n \\ \mathbb{I}_n & 0 \end{pmatrix}$$

Notice that $u^T \Lambda w = 0$ if and only if the corresponding Pauli operators commute.

Given a set of Pauli operators $M \subseteq \mathsf{P_n}$, we define T_M to be the $|M| \times 2n$ binary matrix where each row represents the binary symplectic vector corresponding to a Pauli operator in M. We call this the 'tableau matrix'. We will later make use of the fact that the nullspace of $\mathrm{T}_M \Lambda$ over $\mathbb{Z}_2$ exactly corresponds to the subgroup of all Pauli operators that commute with all elements of M.

2.3 Channels and Unitaries

We use C to refer to Clifford unitaries and $\mathcal{C}$ to refer to their corresponding channel $\mathcal{C}(\cdot) = C(\cdot)C^\dagger$. Given a Clifford unitary C that can be decomposed into a sequence of layers $C = U_d \dots U_1$, we denote the subsequence of unitaries up to timestep i as:

$$C_i := U_i \dots U_1 \tag{1}$$

Our results apply for single-qubit depolarizing channels:

$$\mathcal{E}(\rho) = (1-\gamma)\rho + \gamma \frac{\mathbb{I}}{2} \operatorname{Tr}(\rho) \tag{2}$$

Throughout, Tr_A denotes the partial trace over qubits in A and Tr_{-A} denotes the partial trace over all qubits not in A. We will often refer to noisy Clifford circuits as $\tilde{\mathcal{C}} = \mathcal{E}^{\otimes n}\mathcal{U}_d\mathcal{E}^{\otimes n} \dots \mathcal{U}_1\mathcal{E}^{\otimes n}$, where each $\mathcal{U}_i$ is a layer of the Clifford circuit.

2.4 Output distributions

If $\{F_x\}_x$ is a POVM representing the measurement basis, ρ is an initial state, and $\tilde{\mathcal{C}}$ is a channel representing a noisy circuit, we will refer to the resulting output distribution as $p_{\tilde{\mathcal{C}}}$,

$$p_{\tilde{\mathcal{C}}}(x) = \mathrm{Tr}\Big(F_x \tilde{\mathcal{C}}(\rho)\Big) \tag{3}$$

Moreover, we use TVD to describe how close two distributions are (i.e. approximation error),

$$\Delta(p,q) := \|p - q\|_1 := \sum_{x \in \{0,1\}^n} |p(x) - q(x)| \tag{4}$$

3 Main Results

Theorem 2. *Suppose $\mathcal{C}$ is a Clifford circuit of depth d on n qubits. Let $\tilde{\mathcal{C}}$ denote the noisy implementation of $\mathcal{C}$, where each layer is interspersed with depolarizing channels $\mathcal{E}$ of strength γ on every qubit. Let $p_{\tilde{\mathcal{C}}}$ be the output distribution obtained from preparing n qubits in an arbitrary product state, applying $\tilde{\mathcal{C}}$, and measuring each qubit in an arbitrary basis. There exists a depth threshold $d^* = O(\gamma^{-1}\log n)$ such that when $d > d^*$, there exists a randomized classical algorithm that exactly samples from $p_{\tilde{\mathcal{C}}}$ with random runtime T of expected value,*

$$E[T] = O(\mathsf{poly}(n)) \tag{5}$$

If $\mathcal{C}$ is further restricted to nearest-neighbor gates on a D-dimensional lattice, there exists a constant depth threshold $d^*_{local} = O(\gamma^{-1}3^{2D} + D\gamma^{-1}\log(D\gamma^{-1}))$ such that when $d > d^*_{local}$,

$$\mathbb{E}[T] = O(d^{3+D}n^4) \tag{6}$$

Using a standard reduction (applying Markov's inequality), we can convert this 'Las Vegas' algorithm into a 'Monte Carlo' algorithm with polynomial worst-case runtime.

Corollary 3. *Using the same notation as Theorem 2, there exists a depth threshold $d^* = O(\gamma^{-1}\log n)$ such that when $d > d^*$, there exists a randomized classical algorithm that samples from $q_{\tilde{\mathcal{C}}}$, such that $\|q_{\tilde{\mathcal{C}}} - p_{\tilde{\mathcal{C}}}\|_{TVD} \leq \epsilon$, with runtime*

$$T = O(\epsilon^{-1}\mathsf{poly}(n)) \tag{7}$$

If $\mathcal{C}$ is further restricted to nearest-neighbor gates on a D-dimensional lattice, there exists a constant depth threshold $d^*_{local} = O(\gamma^{-1}3^{2D} + D\gamma^{-1}\log(D\gamma^{-1}))$ such that when $d > d^*_{local}$,

$$T = O(\epsilon^{-1}d^{3+D}n^4) \tag{8}$$

3.1 Preprocessing the Noisy Circuit

In this section, we outline a few techniques we use to simplify the noisy circuit, and make it more amenable to classical simulation.

3.1.1 Stochastic Application of Noise Channels

A key idea is to view each depolarizing channel in the circuit as a stochastic process that traces out the target qubit and replaces it with the identity with probability γ (which we refer to as a 'depolarizing error'), and does nothing otherwise. This stochastic method of applying the depolarizing channel is exactly equivalent in expectation to the more common method of treating it as a deterministic CPTP map. To formalize this, we define the following,

Definition 4. Suppose $\tilde{\mathcal{C}} = \mathcal{E}^{\otimes n}\mathcal{U}_d\mathcal{E}^{\otimes n}\ldots\mathcal{U}_1\mathcal{E}^{\otimes n}$ (where each $\mathcal{U}_t$ is a layer of two-qubit Clifford gates). Let $\mathcal{L}$ represent the ensemble of error locations. For any b drawn from $\mathcal{L}$, we use $\tilde{\mathcal{C}}_b$ to denote the channel that corresponds to applying a depolarizing error at each of the locations specified by b in $\mathcal{C}$.

Note by definition,

$$\tilde{\mathcal{C}}(\rho) = \mathop{\mathbb{E}}_b \tilde{\mathcal{C}}_b(\rho) \tag{9}$$

where $\mathbb{E}_b$ is used throughout the paper as shorthand for $\mathbb{E}_{b\sim\mathcal{L}}$. Note, $p_{\tilde{\mathcal{C}}}(x) = \mathbb{E}_b\, p_{\tilde{\mathcal{C}}_b}(x)$ because the expectation is a linear operator that commutes with any POVM element. Thus, we can exactly sample from $p_{\tilde{\mathcal{C}}}$ by first sampling b and then exactly sampling the measurement outcomes of $\tilde{\mathcal{C}}_b$.

3.1.2 Error Propagation through Clifford Circuits

First, we sample a configuration of depolarizing errors $b \sim \mathcal{L}$ as described in Section 3.1.1. We will now show how to sample from $p_{\tilde{\mathcal{C}}_b}$. A crucial step of our algorithm is to convert the depolarizing errors in $\tilde{\mathcal{C}}_b$ to an error channel that acts only on the circuit input. To this end, we introduce the following notation,

Definition 5 (Pauli projection channel)**.**

$$\Pi_P(\rho) = \frac{1}{2}\rho + \frac{1}{2}P\rho P^\dagger \tag{10}$$

where $P \in \mathsf{P}_n$ given an n-qubit state ρ.

Note since $\Pi_X \circ \Pi_Z(\rho) = \frac{1}{4}\rho + \frac{1}{4}X\rho X^\dagger + \frac{1}{4}Y\rho Y^\dagger + \frac{1}{4}Z\rho Z^\dagger = \mathbb{I}/2$, the depolarizing errors in $\tilde{\mathcal{C}}_b$ can be simulated by the application of $\Pi_X \circ \Pi_Z$. The advantage of simulating depolarizing errors by applying these Pauli projection channels is that they can be propagated to the beginning of the Clifford circuit. For instance given a layer of Clifford gates $\mathcal{U}$, then

$$\begin{aligned}
\Pi_P \circ \mathcal{U}(\rho) &= \frac{1}{2}U\rho U^\dagger + \frac{1}{2}PU\rho U^\dagger P^\dagger && (11)\\
&= \frac{1}{2}U\rho U^\dagger + \frac{1}{2}UU^\dagger PU\rho U^\dagger P^\dagger UU^\dagger && (12)\\
&= \mathcal{U} \circ \Pi_{\mathcal{U}^\dagger(P)}(\rho) && (13)
\end{aligned}$$

This motivates us to define the following,

Definition 6 (Error Propagation)**.** Let M_b denote the set of Pauli operators generated by replacing depolarizing errors in $\tilde{\mathcal{C}}_b$ with Pauli projection channels and propagating these to the beginning of the circuit. Define the corresponding channel $\Pi_{M_b}(\rho) := \bigcirc_{P\in M_b}\Pi_P(\rho)$ acting on the circuit input.

Note we have shown that $\tilde{\mathcal{C}}_b(\rho) = \mathcal{C} \circ \Pi_{M_b}(\rho)$. Now, letting $\langle M_b\rangle \leqslant \mathsf{P}_n$ denote the subgroup generated by M_b we can alternatively write $\Pi_{M_b}(\rho)$ as,

$$\Pi_{M_b}(\rho) = \bigcirc_{P\in M_b}\Pi_P(\rho) = \frac{1}{|\langle M_b\rangle|}\sum_{P\in\langle M_b\rangle} P\rho P^\dagger \tag{14}$$

Crucially, the right-hand side of Eq. 14 only depends $\langle M_b\rangle$ rather than the generators M_b themselves. Thus, this channel is equivalent to any composition of Pauli projection channels where the associated Pauli operators also generate $\langle M_b\rangle$. We occasionally refer to $\langle M_b\rangle$ as the "error group".

3.1.3 Action of Propagated Errors on Pauli Operators

Note that the input state can always be written as a sum of Pauli operators using the Pauli basis and so it will be insightful to characterize the action of Π_{M_b} on Pauli operators. This is described by the following fact.

Fact 7. *Let $s \in \hat{\mathsf{P}}_n$. If s commutes with every Pauli operator in M_b, then $\Pi_{M_b}(s) = s$. Otherwise, $\Pi_{M_b}(s) = 0$*

Proof. Recall that by Eq. (14), $\Pi_{M_b}(s) = \bigcirc_{P\in M_b}\Pi_P(s)$. If s commutes with all elements of M_b then we have that for any $P \in M_b$

$$\Pi_P(s) = \frac{1}{2}(s + PsP^\dagger) = \frac{1}{2}(s + s) = s \tag{15}$$

Thus, $\bigcirc_{P\in M_b}\Pi_P(s) = s$. Otherwise, there exists some $P \in M_b$ such that P anti-commutes with s. In this case,

$$\Pi_P(s) = \frac{1}{2}(s + PsP^\dagger) = \frac{1}{2}(s - s) = 0 \tag{16}$$

and so $\bigcirc_{P\in M_b}\Pi_P(s) = 0$. □

The set of Pauli operators that are preserved by the propagated error channel form a Pauli subgroup that is precisely the centralizer of $\langle M_b\rangle$, which we denote as $\mathsf{C}(\langle M_b\rangle)$. Recall that the centralizer of a group is the set of elements that commute with each element of the group. Using the tableau representation, $\mathsf{C}(\langle M_b\rangle)$ exactly corresponds to the nullspace of $\mathrm{T}_{M_b}\Lambda$.

3.1.4 Effectively Depolarized Qubits

We can now determine if an input qubit is depolarized by Π_{M_b} by inspecting the error group $\langle M_b\rangle$. In particular, we have that if $X_i, Z_i \in \langle M_b\rangle$ then

$$\Pi_{M_b}(\rho) = \Pi_{M_b} \circ \Pi_{X_i} \circ \Pi_{Z_i}(\rho) = \Pi_{M_b}\left(\frac{\mathbb{I}_i}{2} \otimes \mathrm{Tr}_i(\rho)\right) \tag{17}$$

where the first equality follows because $X_i, Z_i \in \langle M_b\rangle$ implies $\langle M_b\rangle = \langle M_b \cup X_i \cup Z_i\rangle$ and so by Eq. 14, Π_{M_b} and $\Pi_{M_b} \circ \Pi_{X_i} \circ \Pi_{Z_i}$ implement the same channel. Note that qubit i of the input state has now become maximally mixed. This motivates us to define depolarized qubits as follows.

Definition 8. [Depolarized Qubits] For any qubit i, if $X_i, Z_i \in \langle M_b\rangle$, then we say qubit i is *depolarized*. Otherwise, qubit i is *non-depolarized*.

Our algorithm will rely on the fact that with high probability an overwhelming majority of qubits become depolarized in this way. We will later show that in this case the depolarized qubits partition the circuit into small islands of non-depolarized qubits that can be simulated independently.

3.2 Description of the Algorithm

We can now write out the full algorithm, which is illustrated in Fig. 1. To summarize, we will sample some configuration of error locations, use this to form subsets of non-depolarized qubits that can be simulated independently, and apply a standard sampling-to-computing reduction (e.g.[BMS17]) to sample using the fact that we can compute marginals efficiently. We prove correctness below in Lemma 9, and leave the full proof of runtime to Section 3.5

For convenience, we will call each set of qubits denoted by L_j above as a *non-depolarized island*.

Lemma 9. *(Correctness of Algorithm 1) For $\tilde{\mathcal{C}}$ defined as in Theorem 2, let $q_{\tilde{\mathcal{C}}}(x)$ be the distribution over output bitstrings x produced by Algorithm 1 on input state ρ. Then, $p_{\tilde{\mathcal{C}}} = q_{\tilde{\mathcal{C}}}$.*

Proof. First recall that by Eq. (17), each depolarized qubit can be replaced with the maximally mixed state $\mathbb{I}/2$ at the beginning of the circuit. It then follows that any error or gate acting solely on maximally mixed qubits can be ignored since $\mathcal{E}(\mathbb{I}/2) = \mathbb{I}/2$ and $\mathcal{U}(\frac{\mathbb{I}}{2} \otimes \frac{\mathbb{I}}{2}) = \frac{\mathbb{I}}{2} \otimes \frac{\mathbb{I}}{2}$. Removing these from the circuit only leaves the errors and gates that are within the lightcone of a non-depolarized qubit. In other words, each remaining error and gate is contained within one of the sets in $\{L_j\}_j$ which are each disjoint. Crucially, this implies that the output state is unentangled across these sets of qubits and so each set can be simulated

[4] the explicit expression for each marginal is in Eq. (18)

Algorithm 1: Sampler for noisy Clifford-magic circuits

Input: Depth-d circuit C, n-qubit state ρ, POVM $\{F_x\}_x$, error rate γ

Promise : C is a Clifford circuit, ρ is a product-state, $F_x = \otimes_i \mathcal{U}_i(|x\rangle\langle x|)$ where each $\mathcal{U}_i$ is an abritrary single-qubit gate acting on qubit i

Output: A sample from the probability distribution $p_{\tilde{\mathcal{C}}}$ defined by $p_{\tilde{\mathcal{C}}}(x) = \mathrm{Tr}\Big(F_x \tilde{\mathcal{C}}(\rho)\Big)$

1 Let $M_b = \{\}$

2 For each depolarizing channel acting on qubit i after timestep t, with probability γ add both $\mathcal{C}_t^\dagger(X_i)$ and $\mathcal{C}_t^\dagger(Z_i)$ to M_b.

3 For each qubit i, check if both $X_i, Z_i \in \langle M_b \rangle$. If so, label this qubit as 'depolarized' and as 'non-depolarized' otherwise.

4 Compute a generating set G for $C(\langle M_b \rangle)$ by finding a basis for the nullspace of $\mathrm{T}_{M_b}\Lambda$ over $\mathbb{Z}_2$.

5 For each non-depolarized qubit i, enumerate the set of qubits in its forward lightcone. Merge together all sets that intersect to form new nonintersecting sets $\{L_j\}_j$.

6 For each L_j, let $G_j \subseteq \hat{\mathsf{P}}_{|L_j|}$ be the set of Pauli operators defined only on L_j, which is formed by taking each element of G and truncating away all qubits outside of L_j.

7 For each L_j, store the final state in the Pauli basis as $\mathrm{Tr}_{-L_j}(\mathcal{C} \circ \Pi_{M_b}(\rho)) = \frac{1}{2^{|L_j|}} \sum_{s \in \langle G_j \rangle} \mathrm{Tr}(\rho s)\mathcal{C}(s)$ and use this to compute any marginal of the output distribution on L_j [4]. Sample measurement outcomes by a standard sampling-to-computing reduction [BMS17] .

8 For all qubits $x \notin L_j$ for all j, output a uniformly random bit as their measurement outcome.

independently. Furthermore the qubits outside of any of these sets remain maximally mixed and so their measurement outcomes can be exactly simulated by uniformly random bits.

To simulate $\mathrm{Tr}_{-L_j} \mathcal{C} \circ \Pi_{M_b}(\rho)$, it is helpful to view the initial state in the Pauli basis as $\rho = \frac{1}{2^n} \sum_{s \in \hat{\mathsf{P}}_n} \mathrm{Tr}(\rho s) s$. Next, it can be shown that the only Pauli operators that pass through $\mathrm{Tr}_{-L_j} \mathcal{C} \circ \Pi_{M_b}$ are those in $\langle G_j \rangle$. This can be seen by first observing that only Pauli operators in $\mathsf{C}(\langle M_b \rangle)$ can pass through Π_{M_b} by Fact 7. This notably only includes Pauli operators with support on non-depolarized qubits. Additionally, any Pauli operator with support outside of L_j will not pass through $\mathrm{Tr}_{-L_j} \mathcal{C} \circ \Pi_{M_b}$ since the lightcone of a non-depolarized qubit outside of L_j is disjoint from L_j. Now, note that $\langle G_j \rangle$ clearly contains all elements in $\mathsf{C}(\langle M_b \rangle))$ that only have support in L_j, but it is not obvious that it does not also contain elements that are outside of $\mathsf{C}(\langle M_b \rangle)$. It can be shown that this is also true by noting that each element of M_b cannot have support on two non-depolarized qubits with non-intersecting lightcones. Therefore, each element in $\mathsf{C}(\langle M_b \rangle)$ will still commute with each element in M_b even after replacing qubits outside of L_j with identity. Thus, $\langle G_j \rangle$ exactly represents the Pauli operators that can pass through $\mathrm{Tr}_{-L_j} \mathcal{C} \circ \Pi_{M_b}$ giving the equality: $\mathrm{Tr}_{-L_j}(\mathcal{C} \circ \Pi_M(\rho)) = \frac{1}{2^{|L_j|}} \sum_{s \in \langle G_j \rangle} \mathrm{Tr}(\rho s)\mathcal{C}(s)$.

We can now express the output probability of outcome z as

$$\mathrm{Tr}\big(\otimes_{i \in L_j} \mathcal{U}_i(|z\rangle\langle z|)\, \mathrm{Tr}_{-L_j} \mathcal{C} \circ \Pi_{M_b}(\rho)\big) = \frac{1}{2^{|L_j|}} \sum_{s \in \langle G_j \rangle} \mathrm{Tr}(\rho s)\, \mathrm{Tr}\big(\otimes_{i \in L_j} \mathcal{U}_i(|z\rangle\langle z|)\mathcal{C}(s)\big) \tag{18}$$

□

3.3 Bounding the Number of Surviving Pauli Operators

The primary runtime cost is due to Step 7, which we will show has a runtime that scales as $O(d^{3+D}n^4)$. This asymptotically dominates the runtime for finding a basis for the nullspace of a binary matrix over $\mathbb{Z}_2$

which can be done in $O(n^3)$ timesteps.

The runtime to simulate qubits in the set L_j is $O(d|L_j|^2|\langle G_j\rangle|)$. This is because it takes $O(d|L_j||\langle G_j\rangle|)$ timesteps to first evolve each Pauli operator through the Clifford circuit since there are $d|L_j|$ gates to apply and $|\langle G_j\rangle|$ operators to apply them on. Next, it takes $O(|L_j||\langle G_j\rangle|)$ to compute each output probability since there are $|\langle G_j\rangle|$ terms to sum in Eq. (18) and $|L_j|$ trace products to compute for each term – one for each qubit. Finally we must compute $|L_j|$ output probabilities for the sampling-to-computing reduction giving a total runtime of $O(d|L_j||\langle G_j\rangle| + |L_j|^2|\langle G_j\rangle|) = O(d|L_j|^2|\langle G_j\rangle|)$.

Next we bound $\mathbb{E}\,|\langle G_j\rangle|$. In the noiseless case (i.e. when M_b is empty), the best bound we can get is $|\langle G_j\rangle| \leq 4^{|L_j|}$, since there are four possible single-qubit Pauli operators in each location. However, noise causes the expected number of 'surviving' Pauli operators to decay with depth and noise strength. This is bounded in the following lemma, which is proven in Section A, and is obtained due to the *unitarity* of the circuit.

Lemma 10. *Using the notation of Algorithm 1,*

$$\mathbb{E}\,|\langle G_j\rangle| \leq de^{3(1-\gamma)^d|L_j|} \tag{19}$$

Corollary 11. *Suppose T_j is the runtime required to sample from $\mathrm{Tr}_{-L_j}(\mathcal{C}\circ\Pi_{M_b}(\rho))$. Then, $\mathbb{E}[T_j] \leq cd^2|L_j|^2e^{3(1-\gamma)^d|L_j|}$ where c is some constant representing the cost of applying Clifford unitaries on Pauli operators.*

In the case that there is no geometric restriction on the Clifford circuit, we obtain our result by assuming there is only one connected component L_1 which contains all qubits and applying the above bound for $d < O(\gamma^{-1}\log n)$. In order to improve this depth threshold to constant for the geometrically local case we also need to show that each L_j is of size logarithmic in n by taking advantage of percolation.

3.4 Percolation into Small Components

To simplify the analysis, we coarse-grain our lattice into sublattices of length $2d$ in each dimension where d is the circuit depth and define a connectivity graph $G = (V, E)$ where each vertex represents a sublattice and each edge represents whether two sublattices are adjacent (including diagonally adjacent). Importantly, two qubits can have intersecting forward lightcones if and only if they are in adjacent sublattices because a lightcone can spread at most a Manhattan distance of d and so it can never reach the lightcone of a qubit in a non-adjacent sublattice. Hereafter, we define a non-depolarized sublattice to mean a sublattice that contains at least one non-depolarized qubit. Our main goal is to show that there does not exist any large connected components of non-depolarized sublattices.

To prove this, we show that if there exists a large connected component of non-depolarized sublattices then there exists a high-weight Pauli operator on these sublattices that 'survives' the Pauli projection channels.

Lemma 12. *For any error configuration b, let $A \subseteq V$ be the largest connected component in G such that each sublattice of A has at least one non-depolarized qubit. There exists some $s \in \mathsf{C}(\langle M_b\rangle)$ with support in at least half of the sublattices in A and no support outside of A.*

By the contrapositive of Lemma 12, if there is no Pauli operator with high weight on a large connected component of sublattices that survives the error channel, then there are no large connected components of sublattices that have non-depolarized qubits. We show in Lemma 25 of Section B that when the depth exceeds some constant depth threshold d^*_{local}, it is possible to tightly bound the probability that there exists some Pauli operator with high-weight support on a connected component of size k, using proof techniques from percolation theory. Combining this with Lemma 12 gives us the following bounds on the size of connected components,

Corollary 13. *Let $\{L_j\}_j$ be the partition of qubits into subsets in Alg. 1.*

$$\mathbf{P}(\max_j |L_j| \geq x) \leq nd \exp\left(-x/(2d)^D\right)$$
$$\mathbf{P}(\max_j diam(L_j) \geq x) \leq 2^{3^D} nd \exp\left(-\Omega(\frac{\gamma x}{3^D D})\right)$$

when $d > d^*_{local}$, where $d^*_{local} = O(\gamma^{-1}3^{2D} + \gamma^{-1}D\log(\gamma^{-1}D))$. $diam(L_j)$ denotes the maximum manhattan distance between two qubits in L_j.

3.5 Computing the Runtime

Now, we have the ingredients to prove the main theorem.

Proof. (Proof of Theorem 2 and Corollary 3) Without geometric locality, we assume there is one connected component L_1 with all qubits. By examining Corollary 11, one can see that when $d > O(\gamma^{-1}\log n)$, the expected runtime is polynomial in n. This proves the first part of Theorem 2. For the geometrically local case, we 'split' the noise channels into two parts of strength roughly $\gamma/2$. More specifically, define $\gamma' = 1 - \sqrt{1-\gamma}$, and replace each noise depolarizing noise channel of strength γ with two consecutive independent depolarizing channels of strength γ'. We use one set of these noise channels to induce percolation into small components, and another set of these noise channels to decay the number of surviving Pauli operators within each component. When $d > d^*_{local}$ (where d^*_{local} is defined for γ' rather than γ [5]), we can compute the expected runtime T of our algorithm as a whole, using linearity of expectation, as follows.

$$\mathbb{E}[T] \leq \sum_j \mathbb{E}[T_j] \tag{20}$$
$$\leq \sum_j \sum_{x=0}^{n} cd^2x^2 e^{3(1-\gamma')^d x} \mathbf{P}(|L_j| = x) \tag{21}$$
$$\leq n \sum_{x=0}^{n} cd^2x^2 e^{3(1-\gamma')^d x} nde^{-\frac{x}{(2d)^D}} \tag{22}$$
$$\leq cd^3n^4 \sum_{x=0}^{n} e^{\frac{x}{(2d)^D}(3(2d)^D(1-\gamma')^d - 1)} \tag{23}$$
$$\leq cd^3n^4 \sum_{x=0}^{n} e^{-\frac{2x}{3(2d)^D}} \tag{24}$$
$$\leq cd^3n^4 \frac{1}{1 - e^{-\frac{2}{3(2d)^D}}} \tag{25}$$
$$= O(d^{3+D}n^4) \tag{26}$$

where we have obtained Eq. (24) by noting that $3(2d)^D(1-\gamma')^d < 1/3$ when $d > d^*_{local}$ and $D \geq 1$ using Eq. (53) proved in Section B. The second to last inequality uses the geometric sum equation, and the final inequality uses the fact that $e^{-a} \leq 1 - \frac{a}{2}$ for $a \leq 1/2$. Lemma 9 tells us that the sampling algorithm is exact, and together with our bounds on expected runtime, this completes the proof of the theorem.

We now prove Corollary 3, which is the version of the algorithm more likely to be implemented in practice. By Markov's inequality, the runtime will exceed $\epsilon^{-1}E[T]$ with probability $< \epsilon$ over the possible error configurations. We can check whether such an adversarial error configuration is sampled at Step 7 of

[5] Abusing notation, we refer to this threshold as d^*_{local} as well since it has the same asymptotic scaling.

Algorithm 1 (i.e. if $|\langle G_j \rangle|$ is too large). If this occurs, we can stop the algorithm and output a random bit string. The TVD error incurred in this case is at most 1, which means the expected TVD error (when the expectation is taken over possible error configurations) is bounded by ϵ. By standard arguments, this means the resultant distribution $q_{\tilde{\mathcal{C}}}$ ϵ-approximates $p_{\tilde{\mathcal{C}}}$ in total variation distance [6]. □

4 Relation to Existing Circuit Classes

4.1 Clifford-Magic and Conjugated Clifford Circuits

Definition 14 (Clifford-magic (CM) circuits [YJS19])**.** Clifford-magic circuits are composed of Clifford gates applied to magic-state input. More specifically, the circuit can be denoted as $C = U_d \dots U_1$ where $U_i \in \mathsf{C}_k$. And the input state is $\rho = |A\rangle\langle A|^{\otimes n}$ where $|A\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\pi/4}|1\rangle)$. The output is measured in the computational basis.

Definition 15 (Conjugated Clifford circuits (CCC) [BFK18])**.** A conjugated Clifford gate is a Clifford conjugated with an arbitrary single-qubit rotation gate U on every qubit. Conjugated Clifford circuits (CCCs) are composed of such gates, where the input state is $|0\rangle\langle 0|^{\otimes n}$ and the output is measured in the computational basis.

Noisy Clifford-magic circuits trivially fall into the model of circuits we simulate. To see how noisy Conjugated Clifford circuits do as well, conjugate each depolarizing channel with U (which leaves it unchanged). Then, each U and $U^\dagger$ between gates and noise channels in the circuits can be cancelled out leaving a circuit of the form $\mathcal{U}^{\otimes n} \circ \tilde{\mathcal{C}} \circ \mathcal{U}^{\dagger \otimes n}$, where $\tilde{\mathcal{C}}$ is some noisy Clifford circuit which can be simulated under our model.

4.2 IQP Circuits augmented with CNOTs

With a few modifications, our techniques also work for simulating noisy IQP circuits with intermediate CNOT gates. Like the circuit classes in the previous section, these have also been proposed for near-term quantum advantage experiments [HKB$^+$24, BEG$^+$24]. We define this circuit family more formally below:

Definition 16 (Instantaneous Quantum Polynomial Circuits augmented with CNOTs (IQP+CNOT) [HKB$^+$24, BEG$^+$24])**.** An IQP+CNOT circuit is a circuit composed of any gates that are diagonal in the computational basis along with CNOT gates, where state preparation and measurement is done in the Hadamard basis.

In the case of IQP+CNOT circuits, our results apply to a broader class of noise channels, which notably include depolarizing and dephasing channels [7]. We state this in the following theorem.

Theorem 17. *The results of Theorem 2 apply in the case that $\rho = |+\rangle\langle +|^{\otimes n}$ and $\tilde{C}$ is composed of only diagonal gates and CNOT gates. Moreover, $\mathcal{E}$ can be any noise channel in the following class of single-qubit Pauli noise channels,*

$$\mathcal{E}(\rho) = (1 - p_X - p_Y - p_Z) + p_X X\rho X + p_Y Y\rho Y + p_Z Z\rho Z \tag{27}$$

where $\gamma = p_Z + \min(p_Y, p_Z)$.

[6] See [RWL24], corollary 23, for a more explicit proof

[7] This is the same class of noise channels for which noisy IQP circuits without CNOT gates become classical simulatable due to [RWL24]

Once again our strategy is to simulate mid-circuit errors by propagating them to the beginning of the circuit, but in this case we only propagate the Z component of the error. These remain as Z errors on the input since they commute with diagonal gates and are mapped to other Z errors by the CNOT gates. Since X errors do not commute with diagonal gates, we will not propagate them to the beginning, and instead stochastically simulate them mid circuit. Notice that the projection channel resulting from these propagated errors is exactly the same as if we had replaced each IQP gate with identity and propagated both X and Z errors to the beginning as done previously. In this modified circuit, Z errors would propagate in the same way as the original IQP+CNOT circuit and X errors would propagate to other X errors which would then act trivially on the input since $|+\rangle\langle+|^n$ is a $+1$ eigenstate of any $s \in \mathsf{P}_n^X$. Since replacing diagonal gates with identity makes the circuit entirely Clifford, all of our previous analysis can be applied to show that many of the qubits become depolarized by this propagated error channel acting on the input. Intuitively, percolation of depolarizing errors still occurs even when only considering Z errors because only a single Z error is needed to depolarize an input qubit (i.e. $\Pi_Z(|+\rangle\langle+|) = \frac{\mathbb{I}}{2}$). We leave the full proof to Section C.

Appendix

A Counting Arguments for Pauli Paths in Clifford Circuits

Note that since $\{X, Y, Z\}$ have trace zero, for any $s \in \mathsf{P}_n$, we have the property that $\mathcal{E}^{\otimes n}(s) = (1-\gamma)^{|s|}s$ where $|s|$ is the number of non-identities in s. This convenient property motivates studying the evolution under the noisy circuit in the Pauli basis. Similar techniques were developed and used in a large body of existing literature ([BMS17, GD18, AGL+23, SYGY24, FRD+23, MAG+24] are examples and maybe an incomplete list). We refer to this broad class of techniques as the Pauli path framework. One of our technical contributions is to reframe these techniques under the application of stochastic processes rather than under the application of a deterministic CPTP map (in earlier work, the probability bound of Eq. (28) is treated as a decay factor for the associated Pauli path coefficient). In the setting of Clifford circuits, the analysis can be greatly simplified, and for any $s \in \mathsf{P}_n$, we have the following fact,

$$\mathbf{P}_b(\tilde{\mathcal{C}}_b(s) \neq 0) = (1-\gamma)^{\sum_i |\mathcal{C}_i(s)|} \tag{28}$$

This can be seen by applying $\mathcal{E}^{\otimes n}$ on $\mathcal{C}_i(s)$ at each layer.

One fact that is unique to Clifford circuits is that any Pauli operator $s \in \mathsf{P}_n$ cannot 'split' into different Pauli paths during the circuit (as is the case with general quantum circuits, e.g. random quantum circuits). It will be useful to examine the point in the circuit at which an input Pauli operator is at its minimum weight during its path through the Clifford circuit. We define this below,

Definition 18. Given any set of qubits A, let $S_w^A(\mathcal{C}, G) \subseteq \hat{\mathsf{P}}_n$ be the set $S_w^A(\mathcal{C}, G) := \{s \in G : \min_i \mathcal{C}_i(s) = w$, and $\forall_i\ \mathcal{C}_i(s)$ is supported only in $A\}$. When A is not defined, we take it to be the set of all qubits. When w is not defined, we take it to include all $w : 0 \leq w \leq |A|$. When $\mathcal{C}$ is not defined we take it to be the entire circuit, which can be inferred from context. When G is not defined we take it to be $\hat{\mathsf{P}}_n$.

By applying Eq. (28), we have that,

Fact 19. *For $s \in S_w^A$ and noisy Clifford circuit $\tilde{\mathcal{C}}$ of depth d*

$$\mathbf{P}_b(\tilde{\mathcal{C}}_b(s) \neq 0) = (1-\gamma)^{\sum_i |\mathcal{C}_i(s)|} \leq (1-\gamma)^{dw} \tag{29}$$

Crucially, this probability decays exponentially with depth and weight. Now, we can also bound the size of S_w as follows

Lemma 20.

$$|S_w^A| \leq d\binom{|A|}{w}3^w \tag{30}$$

Proof. Note, for each $s \in S_w^A$, there exists some i at which $|\mathcal{C}_i(s)| = w$ and its support is contained within A. Thus, for all $s \in S_w^A$, the set of all possible weight-w Pauli operators that are contained within A contains $\mathcal{C}_i(s)$ for some i. By *unitarity*, each such $\mathcal{C}_i(s)$ corresponds to exactly one input Pauli s. Therefore we can upper bound $|S_w^A|$ by counting all weight-w Pauli operators contained in A for each depth. We can ignore phases when performing this counting since $S_w^A \subseteq \hat{\mathsf{P}}_n$ and for each $P \in \hat{\mathsf{P}}_n$ there is exactly one phase $\ell \in \{0,1,2,3\}$ such that $\mathcal{C}_i^\dagger(i^\ell P) \in \hat{\mathsf{P}}_n$. The counting argument proceeds as follows. First, there are d ways to choose the timestep i at which $|\mathcal{C}(i)| = w$. Next, there are $\binom{|A|}{w}$ ways to choose the locations of w non-identities, and 3^w ways to assign those locations to single-qubit Pauli operators (X,Y, or Z). □

This allows us to obtain the following bound:

Lemma 21. *(Restatement of Lemma 10) Using the notation of Algorithm 1, we have*

$$\mathbb{E}\,|\langle G_j\rangle| \leq de^{3(1-\gamma)^d|L_j|} \tag{31}$$

Proof. Note that once we have sampled some b, $\langle G_j\rangle = S^{L_j} \cap \mathsf{C}(\langle M_b\rangle)$. Moreover, this is equivalent to the set $\{s \in S^{L_j} : \tilde{\mathcal{C}}_b(s) \neq 0\}$. Now, applying linearity of expectation, we have the following,

$$\begin{aligned}
\mathbb{E}\,|\langle G_j\rangle| &= \mathop{\mathbb{E}}_b |\{s \in S^{L_j} : \tilde{\mathcal{C}}_b(s) \neq 0\}| && (32)\\
&\leq \sum_{s\in S^{L_j}} \mathbf{P}_b(\tilde{\mathcal{C}}_b(s) \neq 0) && (33)\\
&\leq \sum_{w=0}^{|L_j|} \sum_{s\in S_w^{L_j}} \mathbf{P}_b(\tilde{\mathcal{C}}_b(s) \neq 0) && (34)\\
&\leq \sum_{w=0}^{|L_j|} d\binom{|L_j|}{w}3^w(1-\gamma)^{wd} && \text{by Fact 19 and Lemma 20} \quad (35)\\
&\leq d(1+3(1-\gamma)^d)^{|L_j|} && \text{by binomial theorem} \quad (36)\\
&\leq d\exp\Big(3(1-\gamma)^d|L_j|\Big) && (37)
\end{aligned}$$

□

B Depth Threshold for Percolation

A useful fact that we will soon make use of is that for a Pauli subgroup $G \leqslant \hat{\mathsf{P}}_n$

$$\mathsf{C}(\mathsf{C}(G)) = G \tag{38}$$

Although, we do not include a proof here, this can be seen by mapping from phaseless Pauli subgroups to symplectic subspaces where the equivalent statement is shown in Theorem 11.8 of [Rom08].

Lemma 22. *(Restatement of Lemma 12) For any error configuration b, let $A \subseteq V$ be the largest connected component in G such that each sublattice of A has at least one non-depolarized qubit. There exists some $s \in \mathsf{C}(\langle M_b\rangle)$ with support in at least half of the sublattices in A and no support outside of A.*

Proof. We start by showing that for each non-depolarized qubit i in each sublattice of A there is a Pauli operator s_i such that

1. $\Pi_{M_b}(s_i) \neq 0$
2. s_i has support on qubit i
3. Its support is contained within A

Note that by Fact 7, (1) is equivalent to $s_i \in \mathsf{C}(\langle M_b \rangle)$.

We first show there exists an s_i with properties (1) and (2). For sake of contradiction assume there does not a exist an s_i such that (1) and (2) are true. Next, add X_i and Z_i to M_b to produce M_b'. Since every element of $\mathsf{C}(\langle M_b \rangle)$ has identity on qubit i by assumption, they each will commute with X_i and Z_i, and so $\mathsf{C}(\langle M_b \rangle) = \mathsf{C}(\langle M_b' \rangle)$. This implies $\langle M_b \rangle = \langle M_b' \rangle$ (by Eq. (38)), which implies that $X_i, Z_i \in M_b$. By definition this means that qubit i is depolarized which is a contradiction. Thus, (1) and (2) must be true.

Next, we show that there exists an s_i that additionally has property (3). To do this we start with an s_i' with properties (1) and (2), and construct s_i by setting all qubits of s_i' that are outside of A to identity. s_i now satisfies (2) and (3). We next show that it also satisfies (1). Let $N(A)$ denote the set of sublattices that are adjacent to A and including A. First, notice that s_i' and s_i are both not supported in adjacent sublattices of A as these qubits are depolarized (otherwise they would be included in A) and so no Pauli with non-identity on $N(A) - A$ can pass through Π_{M_b}. Now, choose any $P \in M_b$. If P has support on A, then it cannot have support outside of $N(A)$ because each P is contained within a lightcone of the circuit and lightcones can only spread to adjacent sublattices. Therefore, its commutation relation with s_i will be the same as its commutation relation with s_i' (as they are identical on $N(A)$). Alternatively, if P does not have support in A, then they must commute as well since s_i's support is contained entirely in A.

Equipped with this set of Pauli operators $\{s_i\}_{i \in A}$ with the above three properties, we finish the proof using the probabilistic method. Define a randomly chosen Pauli operator s as follows,

$$s = \prod_{i \in A} s_i^{r_i} \tag{39}$$

where each r_i is drawn from an independent random variable R_i taking values 0 or 1 with equal probability. Notice that s maintains properties (1) and (3). For each qubit $i \in A$ define the following event:

$$E_i = \begin{cases} 1 & \text{if } s \text{ has a non-identity on qubit } i \\ 0 & \text{otherwise} \end{cases} \tag{40}$$

Observe that $\mathbb{E}\, E_i \geq 1/2$. This can be seen by the following argument. Assume without loss of generality that s_i contains an X on qubit i. The argument works the same if it contains a Y or Z instead and we are guaranteed one of the three is true by property (2). Then let $S_X = \{j \in A : s_j \text{ contains an } X \text{ on qubit } i\}$. Notice that if $\oplus_{j \in S_X} r_j = 1$ then s has non-identity on qubit i. $\Pr(\oplus_{j \in S_X} R_j = 1) = 1/2$ as long as S_X is non-empty, which is true by assumption, and so $\mathbb{E}\, E_i \geq 1/2$. Using linearity of expectation, we have,

$$\mathbb{E} \sum_{i \in A} E_i = \sum_{i \in A} \mathbb{E}\, E_i \geq |A|/2 \tag{41}$$

There must be at least one assignment to the random variables $\{R_i\}$ that achieves at least this expectation. Therefore, there exists some Pauli operator $s \in \mathsf{C}(\langle M_b \rangle)$ with support on at least $|A|/2$ of the sublattices and no support outside of A. □

Next, we must show that with high probability there does not exist any such Pauli operator that has support on half of the sublattices of a large connected component. This set of Pauli operators is formally defined below.

Definition 23. Let $T_k \subseteq \hat{\mathsf{P}}_n$ be the subset of Pauli operators such that $s \in T_k$ if s is contained within a connected component of k sublattices and has support on at least $k/2$ sublattices within this connected component.

Next, using a similar argument to Section A, we bound the number of Pauli operators in T_k that have a minimum weight of w in the circuit.

Lemma 24.

$$|T_k \cap S_w| \leq nd2^{3^D k}\binom{k(6d)^D}{w}3^w \tag{42}$$

Proof. Recall that by definition each $s \in T_k$ is contained within a connected component of k sublattices. There are at most $n2^{\Delta k}$ ways to choose a k-sized connected component (in which s has support) in a Δ-regular graph [Kri14]. Here $\Delta = 3^D - 1$ although we drop the minus one for convenience. When the circuit is applied, the support of $\mathcal{C}_i(s)$ is contained within the neighborhood of the original connected component due to lightcones. This consists of $3^D k$ sublattices since each sublattice has 3^D neighbors including itself. We label the set of these qubits as A, where $|A| = 3^D k(2d)^D = (6d)^D k$. Since $T_k \cap S_w \subseteq S_w^A$, it remains to upper bound $|S_w^A|$. By Lemma 20, $|S_w^A| \leq d\binom{k(6d)^D}{w}3^w$. □

Applying this upper bound on $|T_k \cap S_w|$, we can now bound the probability that any $s \in T_k$ survives $\tilde{\mathcal{C}}$.

Lemma 25.

$$\mathbf{P}_b(\bigcup_{s \in T_k} \Pi_{M_b}(s) \neq 0) \leq nd\exp(-k) \tag{43}$$

when $d > d^*_{local}$, *where* $d^*_{local} = O(\gamma^{-1}3^{2D} + \gamma^{-1}D\log(\gamma^{-1}D))$

Proof. First note that for any timestep i and any $s \in T_k$, $|\mathcal{C}_i(s)| \geq \frac{k}{2\cdot 3^D}$. To show this, we apply a similar lightcone argument to the proof of Lemma 24. If $\mathcal{C}_i(s)$ has support on x sublattices, then $\mathcal{C}_i^\dagger(\mathcal{C}_i(s)) = s$ has support on at most $3^D x$ sublattices due to lightcones. Therefore, since s has support on $k/2$ sublattices, this guarantees that $\mathcal{C}_i(s)$ has support on at least $k/(2\cdot 3^D)$ sublattices.

$$\mathbf{P}_b(\bigcup_{s\in T_k}\Pi_{M_b}(s)\neq 0)\leq\sum_{s\in T_k}\mathbf{P}_b(\Pi_{M_b}(s)\neq 0) \tag{44}$$

$$=\sum_{w=k/(2\cdot 3^D)}^{k(6d)^D}\sum_{s\in T_k\cap S_w}(1-\gamma)^{\sum_i|\mathcal{C}_i(s)|} \tag{45}$$

$$\leq\sum_{w=k/(2\cdot 3^D)}^{k(6d)^D}|T_k\cap S_w|(1-\gamma)^{wd} \tag{46}$$

$$\leq(1-\gamma)^{kd/(4\cdot 3^D)}\sum_{w=1}^{k(6d)^D}|T_k\cap S_w|(1-\gamma)^{wd/2} \tag{47}$$

$$\leq(1-\gamma)^{kd/(4\cdot 3^D)}\sum_{w=1}^{k(6d)^D}nd2^{3^Dk}\binom{k(6d)^D}{w}3^w(1-\gamma)^{wd/2} \tag{48}$$

$$=(1-\gamma)^{kd/(4\cdot 3^D)}nd2^{3^Dk}(1+3(1-\gamma)^{d/2})^{k(6d)^D} \tag{49}$$

$$\leq nd\exp\left(\frac{k(-\gamma d)}{4\cdot 3^D}+k3^D\ln 2+3(1-\gamma)^{d/2}k(6d)^D\right) \tag{50}$$

$$\leq nd\exp\left(-k[\frac{\gamma d}{4\cdot 3^D}-3^D\ln 2-3(1-\gamma)^{d/2}(6d)^D]\right) \tag{51}$$

Eq. (45) and Eq. (46) follow from Fact 19, Eq. (47) splits the error term into two factors, Eq. (48) applies Lemma 24, and Eq. (49) applies the binomial theorem. We now want to show that the following term is ≥ 1 in some depth regimes,

$$\frac{\gamma d}{4\cdot 3^D}-3^D\ln 2-3(1-\gamma)^{d/2}(6d)^D \tag{52}$$

We can lower bound the first two terms $\frac{\gamma d}{4\cdot 3^D}-3^D\ln 2\geq 2$ when $d\geq\gamma^{-1}(4\cdot 3^D)(3^D\ln 2+2)$. We can write the third term as,

$$3(1-\gamma)^{d/2}(6d)^D\leq 3(6de^{-\frac{\gamma d}{2D}})^D \tag{53}$$

$$=3(6d\frac{4D}{\gamma}\frac{\gamma}{4D}e^{-\frac{\gamma d}{2D}})^D \tag{54}$$

$$<3(6\frac{4D}{\gamma}e^{-\frac{\gamma d}{4D}})^D \qquad \text{using the fact that } xe^{-x}<1 \tag{55}$$

$$=3(6\frac{4D}{\gamma})^De^{-\frac{\gamma d}{4}} \tag{56}$$

$$\leq 1 \tag{57}$$

where the last inequality holds true when $d\geq O(\gamma^{-1}D\log(\gamma^{-1}D))$. Thus, we have shown that $\frac{\gamma d}{4\cdot 3^D}-3^D\ln 2-3(1-\gamma)^{d/2}(6d)^D\geq 1$ when $d>d^*_{local}$ where $d^*_{local}=O(\gamma^{-1}3^{2D}+\gamma^{-1}D\log(\gamma^{-1}D))$.

□

Corollary 26. *(Restatement of Corollary 13) Let $\{L_j\}_j$ be the partition of qubits into subsets in Alg. 1.*

$$\mathbf{P}(\max_j|L_j|\geq x)\leq nd\exp\left(-x/(2d)^D\right)$$

$$\mathbf{P}(\max_j diam(L_j)\geq x)\leq 2^{3^D}nd\exp\left(-\Omega(\frac{\gamma x}{3^DD})\right)$$

when $d > d^*_{local}$, where $d^*_{local} = O(\gamma^{-1}3^{2D} + \gamma^{-1}D\log(\gamma^{-1}D))$. $diam(L_j)$ denotes the maximum manhattan distance between two qubits in L_j.

Proof. Note that each set L_j in Algorithm 1 spans at least $|L_j|/(2d)^D$ sublattices. Therefore by Lemma 12 there exists $s \in T_k$ for $k \geq \max_j |L_j|/(2d)^D$ such that $\Pi_{M_b}(s) \neq 0$. Using Lemma 25, this suffices as proof of the first inequality. When considering the diameter of L_j, we can make a few improvements. Firstly, when $d > d^*_{local}$, Lemma 25 can be tightened to

$$\mathbf{P}_b(\bigcup_{s\in T_k} \Pi_{M_b}(s) \neq 0) \leq 2^{3^D} nd \exp\left(-\Omega(\frac{\gamma dk}{3^D})\right) \tag{58}$$

by a simple examination of Eq. (51). Next, if L_j contains two qubits of manhattan distance x, then it spans at least $k \geq x/2dD$ sublattices (regardless of D) by considering the case where they are diagonal. Combining these two facts gives the second inequality. □

C Subroutine for IQP + CNOT Circuits

Theorem 27. *(Restatement of Theorem 17) The results of Theorem 2 apply in the case that $\rho = |+\rangle^{\otimes n}$ and $\tilde{C}$ is composed of diagonal gates and CNOT gates. Moreover, $\mathcal{E}$ can be any noise channel in the following class of single-qubit Pauli noise channels,*

$$\mathcal{E}(\rho) = (1 - p_X - p_Y - p_Z) + p_X X\rho X + p_Y Y\rho Y + p_Z Z\rho Z \tag{59}$$

where $\gamma = p_Z + \min(p_Y, p_Z)$. Notably, this includes depolarizing and dephasing channels.

Proof. We make the following edits to the steps of Algorithm 1.

- Step 2: Add only $\mathcal{C}_t^\dagger(Z_i)$ to M_b instead of both $\mathcal{C}_t^\dagger(Z_i)$ and $\mathcal{C}_t^\dagger(X_i)$. Simulate the Π_{X_i} in place by adding either an X_i or an identity to the circuit with equal probability. As noted in [RWL24], this technique allows us to generalize to dephasing noise channels and channels which can be written as a composition of dephasing channels with other Pauli noise channels.
- Step 3: We will say qubit i is 'depolarized' if $Z_i \in \langle M_b\rangle$ (since $\Pi_Z(|+\rangle\langle+|) = \frac{\mathbb{I}}{2}$).
- Step 4: Compute a generating set G for $\mathsf{C}(\langle M_b\rangle) \cap \mathsf{P}_n^X$ rather than the full group $\mathsf{C}(\langle M_b\rangle)$, since the Pauli basis decomposition of the input state $|+\rangle\langle+|^n$ contains only operators in P_n^X

After the above modifications, $M_b \subseteq \mathsf{P}_n^Z$, while $G \subseteq \mathsf{P}_n^X$. Due to this fact, most of the analysis for the Clifford case imports directly over. The only step of the analysis that is not as trivial is computing the output probabilities for the distribution on L_j after writing the output state in the Pauli basis. In the Clifford case, this is simple because one can efficiently evolve each Pauli basis element as $\mathcal{C}(s)$ for each $s \in \langle G_j\rangle$. In the IQP+CNOT setting, the non-Clifford diagonal gates may act non-trivially on the Pauli operators, so simulation in the Pauli basis may not be ideal. In other words, going from Lemma 10 to Corollary 11 is not simple.

First we note that the input state can be written in the Pauli basis as:

$$\rho = \frac{1}{2^n}\sum_{s\in\langle G_j\rangle} \mathrm{Tr}(|+\rangle\langle+|^n s)s = \frac{1}{2^n}\sum_{s\in\langle G_j\rangle} s \tag{60}$$

Next, Lemma 28 (proved below) tells us that states of this form also have a sparse representation in the computational basis. That is,

$$\rho = \mathop{\mathbb{E}}_{r\in\{0,1\}^{n-|G|}} |\psi_r\rangle\langle\psi_r| \tag{61}$$

where each $|\psi_b\rangle$ is a superposition of only $2^{|G_j|}$ computational basis states. Once we have this representation, we can sample $r \in \{0,1\}^{n-|G_j|}$ randomly, and simulate $C|\psi_r\rangle$ in the computational basis by updating its state vector description after each CNOT and diagonal gate. Note that neither CNOTs nor diagonal gates can create a superposition of computational basis states, so we will only need to store $2^{|G_j|} = |\langle G_j\rangle|$ computational basis states throughout the entire circuit evolution. We can then sample single-qubit measurement outcomes on the final state using standard state vector methods.

□

Lemma 28. *For any set of Pauli operators $G \leq \mathsf{P}_n^X$ such that $\rho = \frac{1}{2^n}\sum_{s\in\langle G\rangle} s$ is a valid quantum state, one can find a decomposition of ρ of the following form*

$$\rho = \mathop{\mathbb{E}}_{r\in\{0,1\}^{n-|G|}} |\psi_r\rangle\langle\psi_r| \tag{62}$$

where $|\psi_b\rangle = \sum_{i\in\{0,1\}^{|G|}} \frac{1}{\sqrt{2^{|G|}}}|f(i,r)\rangle$ for some efficiently computable function $f : \{0,1\}^n \to \{0,1\}^n$.

Proof. Let T_G^X be the first n columns of the matrix T_G. In other words it is the X component of the tableau matrix for G, which only contains Pauli operators in P_n^X anyways. It can be shown that there exists an $n\times n$ binary matrix A such that

$$T_G^X A = \begin{pmatrix}\mathbb{I}_{|G|} & 0\end{pmatrix} \tag{63}$$

where A is composed of the elementary column operations associated with performing Gauss-Jordan elimination on the transpose of T_G^X. The above form follows from the fact that T_G^X has rank $|G|$ since G is a basis of independent operators.

Note, elementary column operations on binary matrices are either (1) swapping column i and column j or (2) setting column i to be column $i \oplus$ column j. Since each column represents a qubit, operation (1) corresponds to a SWAP operation between qubit i and j while operation (2) corresponds to a CNOT operation targetted on qubit i and controlled on qubit j. Therefore, A can be associated with a unitary matrix U_A which performs these operations.

Notice now that,

$$\begin{aligned}
\mathcal{U}_A(\rho) &= \mathcal{U}_A\left(\frac{1}{2^n}\sum_{b_1,\ldots,b_{|G|}\in\{0,1\}}\prod_{s\in G} s^{b_i}\right) && (64)\\
&= \frac{1}{2^n}\sum_{b_1,\ldots,b_{|G|}\in\{0,1\}}\prod_{s\in G}\mathcal{U}_A(s^{b_i}) && (65)\\
&= \frac{1}{2^n}\sum_{b_1,\ldots,b_{|G|}\in\{0,1\}}\prod_{i\in|G|} X_i^{b_i} && (66)\\
&= \left(\frac{\mathbb{I}+X}{2}\right)^{\otimes|G|}\otimes\left(\frac{\mathbb{I}}{2}\right)^{\otimes n-|G|} && (67)\\
&= |+\rangle\langle+|^{\otimes|G|}\otimes\left(\frac{\mathbb{I}}{2}\right)^{\otimes n-|G|} && (68)\\
&= \mathop{\mathbb{E}}_{r\in\{0,1\}^{n-|G|}} |+\rangle\langle+|^{|G|}\otimes|r\rangle\langle r| && (69)\\
& && (70)
\end{aligned}$$

Therefore, $\rho = \mathcal{U}_A^\dagger(\mathbb{E}_{r\in\{0,1\}^{n-|G|}} |+\rangle\langle+|^{|G|} \otimes |r\rangle\langle r|)$. Since $U_A^\dagger$ is composed of SWAP and CNOT gates, it maps computational basis states to computational basis states. Defining $f : \{0,1\}^n \to \{0,1\}^n$ to be the function such that $f(x) = \mathcal{U}_A^\dagger(|x\rangle\langle x|)$ concludes the proof. □

D Noise-induced Anticoncentration in Clifford Circuits

While the circuits we consider are known to be hard to exactly sample from in the worst-case, an additional property known as anticoncentration is usually required in order to show hardness of approximate sampling. In particular, Clifford-Magic, Conjugated Clifford, and IQP+CNOT circuits were all initially proposed as an ensemble of random circuits which anticoncentrate [BFK18, YJS19, HKB+24, BEG+24]. This anticoncentration property enables one to prove hardness of approximate sampling for these circuit families (up to complexity-theoretic assumptions), but at the same time, it allows for polynomial-time classical simulation when the circuit is noisy [BMS17, AGL+23, SYGY24]. Thus, an important line of work is to understand the regimes in which anticoncentration sets in. Here, we show that our counting arguments can be adapted to show that *any* noisy Clifford circuit on random input bit-strings anticoncentrates in $O(\log n)$ depth. When considering an ensemble of circuits where Clifford gates are drawn randomly and states are prepared and measured in an arbitrary product basis, the randomness in the input bit-strings can be absorbed into the circuit. Finally, since random Clifford gates form a 2-design for Haar random gates [DCEL09, DLT02, Web16, ZKGG16], this result also implies that noisy Haar random circuits anticoncentrate in $\Theta(\log n)$-depth for all architectures, due to noise. This was only previously known for $1D$ and all-to-all architectures.

Theorem 29. *Using the same notation as Theorem 2, suppose we use $p_{\tilde{\mathcal{C}},y}$ to denote the probability distribution generated by $\tilde{\mathcal{C}}$ applied to some computational basis state $|y\rangle\langle y|$ and measured in an arbitrary basis. Then, when $d \geq O(\gamma^{-1}\log n)$:*

$$\mathop{\mathbb{E}}_{y\sim\{0,1\}^n} \sum_{x\in\{0,1\}^n} p_{\tilde{\mathcal{C}},y}(x)^2 = \frac{O(1)}{2^n} \tag{71}$$

where the expectation is taken over the uniform distribution.

Proof. We will make use of the following orthogonality property. For any $s, s' \in \hat{\mathsf{P}}_n^Z$, if $s \neq s'$, then

$$\mathop{\mathbb{E}}_{y\sim\{0,1\}^n} \mathrm{Tr}(s\,|y\rangle\langle y|)\,\mathrm{Tr}\big(s'\,|y\rangle\langle y|\big) = 0 \tag{72}$$

Letting $v_z(\cdot)$ be the map to the binary symplectic representation where we have dropped the X-component, the above property can be seen as follows:

$$\begin{aligned}
\mathop{\mathbb{E}}_{y\sim\{0,1\}^n} \mathrm{Tr}(s\,|y\rangle\langle y|)\,\mathrm{Tr}\big(s'\,|y\rangle\langle y|\big) &= \mathop{\mathbb{E}}_{y\sim\{0,1\}^n} -1^{v_z(s)\cdot y + v_z(s')\cdot y} && (73)\\
&= \mathop{\mathbb{E}}_{y\sim\{0,1\}^n} -1^{(v_z(s)+v_z(s'))\cdot y} && (74)\\
&= 0 \qquad\qquad \text{if } v_z(s)+v_z(s') \neq 0 && (75)
\end{aligned}$$

We will relate anticoncentration to the expected number of Pauli operators which 'survive' the errors. Recall that given a configuration of errors b in circuit $\tilde{C}_b$, we denote the group of propagated errors as M_b and the corresponding centralizer ('surviving' Pauli operators) as $\mathsf{C}(\langle M_b\rangle)$. Using U to denote the unitary which

transforms the measurement basis into the computational basis, we have,

$$\mathop{\mathbb{E}}_{y\sim\{0,1\}^n}\sum_{x\in\{0,1\}^n} p_{\tilde{\mathcal{C}},y}(x)^2 = \mathop{\mathbb{E}}_{y\sim\{0,1\}^n}\sum_{x\in\{0,1\}^n}\mathrm{Tr}\Big(|x\rangle\langle x|\,\mathcal{U}\circ\tilde{\mathcal{C}}(|y\rangle\langle y|)\Big)^2 \tag{76}$$

$$= \mathop{\mathbb{E}}_{y\sim\{0,1\}^n}\sum_{x\in\{0,1\}^n}\mathrm{Tr}\left(|x\rangle\langle x|\,\mathop{\mathbb{E}}_{b}\left[\mathcal{U}\circ\mathcal{C}\circ\Pi_{M_b}(|y\rangle\langle y|)\right]\right)^2 \tag{77}$$

$$= \mathop{\mathbb{E}}_{y\sim\{0,1\}^n}\sum_{x\in\{0,1\}^n}\left(\frac{1}{2^n}\sum_{s\in\hat{\mathsf{P}}_n^Z}\mathop{\mathbb{E}}_{b}\left[\mathrm{Tr}(|x\rangle\langle x|\,\mathcal{U}\circ\mathcal{C}\circ\Pi_{M_b}(s))\,\mathrm{Tr}(s\,|y\rangle\langle y|)\right]\right)^2 \tag{78}$$

$$= \frac{1}{2^{2n}}\mathop{\mathbb{E}}_{y\sim\{0,1\}^n}\sum_{x\in\{0,1\}^n}\left(\sum_{s\in\hat{\mathsf{P}}_n^Z}\mathbf{P}_b(\Pi_{M_b}(s)\neq 0)\,\mathrm{Tr}\Big(\mathcal{U}^\dagger(|x\rangle\langle x|)\mathcal{C}(s)\Big)\,\mathrm{Tr}(s\,|y\rangle\langle y|)\right)^2 \tag{79}$$

$$= \frac{1}{2^{2n}}\mathop{\mathbb{E}}_{y\sim\{0,1\}^n}\sum_{x\in\{0,1\}^n}\sum_{s\in\hat{\mathsf{P}}_n^Z}\mathbf{P}_b(\Pi_{M_b}(s)\neq 0)^2\,\mathrm{Tr}\Big(\mathcal{U}^\dagger(|x\rangle\langle x|)\mathcal{C}(s)\Big)^2\,\mathrm{Tr}(s\,|y\rangle\langle y|))^2 \tag{80}$$

$$\leq \frac{1}{2^{2n}}\sum_{x\in\{0,1\}^n}\sum_{s\in\hat{\mathsf{P}}_n^Z}\mathbf{P}_b(\Pi_{M_b}(s)\neq 0)^2 \tag{81}$$

$$= \frac{1}{2^n}\sum_{s\in\hat{\mathsf{P}}_n^Z}\mathbf{P}_b(\Pi_{M_b}(s)\neq 0)^2 \tag{82}$$

$$= \frac{1}{2^n}\left[1+\sum_{s\in\{\hat{\mathsf{P}}_n^Z - I^{\otimes n}\}}\mathbf{P}_b(\Pi_{M_b}(s)\neq 0)^2\right] \tag{83}$$

$$\leq \frac{1}{2^n}\left[1+(1-\gamma)^d\sum_{s\in\{\hat{\mathsf{P}}_n^Z - I^{\otimes n}\}}\mathbf{P}_b(\Pi_{M_b}(s)\neq 0)\right] \tag{84}$$

$$= \frac{1}{2^n}\left[1+(1-\gamma)^d\mathop{\mathbb{E}}_{b}|\mathsf{C}(\langle M_b\rangle)|\right] \tag{85}$$

$$\leq \frac{1}{2^n}\left[1+(1-\gamma)^d de^{3(1-\gamma)^d n}\right] \tag{86}$$

where we have used orthogonality in Eq. (80), the fact that Pauli observables on states are bounded by 1 Eq. (81), the fact that Pauli operators with weight at least 1 survive with probability at most $(1-\gamma)^d$ in Eq. (84), and Lemma 10 in Eq. (86). Therefore, we achieve anticoncentration when $d > O(\gamma^{-1}\log n)$. □

References

[AA13] Scott Aaronson and Alex Arkhipov. The Computational Complexity of Linear Optics. *Th. Comp.*, 9(4):143–252, 2013. doi:10.4086/toc.2013.v009a004. 1

[ABOIN96] D. Aharonov, M. Ben-Or, R. Impagliazzo, and N. Nisan. Limitations of noisy reversible computation, 1996. arXiv:quant-ph/9611028. 1, 1.2

[AGL+23] Dorit Aharonov, Xun Gao, Zeph Landau, Yunchao Liu, and Umesh Vazirani. A polynomial-time classical algorithm for noisy random circuit sampling. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 945–957, June 2023. arXiv:2211.03999, doi:10.1145/3564246.3585234. 1, 1.1, 1.2, 1.3, A, D

[Aha00] Dorit Aharonov. A Quantum to Classical Phase Transition in Noisy Quantum Computers. *Phys. Rev. A*, 62(6):062311, November 2000. arXiv:quant-ph/9910081, doi:10.1103/PhysRevA.62.062311. 1.2

[ALNP25] A. Anshu, Y. Liu, Q. Nguyen, and C. Pattison, 2025. Simons talk available at: https://simons.berkeley.edu/talks/quynh-nguyen-harvard-university-2025-05-28. 1.4.2

[AZ24] Scott Aaronson and Yuxuan Zhang. On verifiable quantum advantage with peaked circuit sampling. April 2024. arXiv:2404.14493. 1.4.4

[BEF+08] Daniel E. Browne, Matthew B. Elliott, Steven T. Flammia, Seth T. Merkel, Akimasa Miyake, and Anthony J. Short. Phase transition of computational power in the resource states for one-way quantum computation. *New J. Phys.*, 10(2):023010, February 2008. doi:10.1088/1367-2630/10/2/023010. 1.2

[BEG+24] Dolev Bluvstein, Simon J. Evered, Alexandra A. Geim, Sophie H. Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J. Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J. Gullans, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. Logical quantum processor based on reconfigurable atom arrays. *Nature*, 626(7997):58–65, February 2024. doi:10.1038/s41586-023-06927-3. 1, 1.4.1, 1.4.6, 4.2, 16, D

[BFK18] Adam Bouland, Joseph F. Fitzsimons, and Dax Enshan Koh. Complexity classification of conjugated clifford circuits. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.CCC.2018.21, doi:10.4230/LIPICS.CCC.2018.21. 1, 15, D

[BGKT20] Sergey Bravyi, David Gosset, Robert König, and Marco Tomamichel. Quantum advantage with noisy shallow circuits. *Nat. Phys.*, 16(10):1040–1045, October 2020. doi:10.1038/s41567-020-0948-z. 1.4.1, 1.4.2

[BGL23] Sergey Bravyi, David Gosset, and Yinchen Liu. Classical simulation of peaked shallow quantum circuits. (arXiv:2309.08405), September 2023. arXiv:2309.08405. 3

[BJS10] M. J. Bremner, R. Jozsa, and D. J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2010. doi:10.1098/rspa.2010.0301. 1

[BK05] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A*, 71(2):022316, February 2005. doi:10.1103/PhysRevA.71.022316. 1.4.1

[BL25] Thiago Bergamaschi and Yunchao Liu. On fault tolerant single-shot logical state preparation and robust long-range entanglement. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2025.16, doi:10.4230/LIPICS.ITCS.2025.16. 1.4.1, 1.4.3

[BMS16] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations. *Physical Review Letters*, 117(8):080501, August 2016. doi:10.1103/PhysRevLett.117.080501. 1

[BMS17] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Achieving quantum supremacy with sparse and noisy commuting quantum computations. *Quantum*, 1:8, April 2017. doi:10.22331/q-2017-04-25-8. 1, 1.1, 1.2, 1.3, 3.2, 7, A, D

[Bom15] H. Bombin. Gauge color codes: Optimal transversal gates and gauge fixing in topological stabilizer codes, 2015. URL: https://arxiv.org/abs/1311.0879, arXiv:1311.0879. 1.4.1

[CCHL22] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. The complexity of nisq, 2022. URL: https://arxiv.org/abs/2210.07234, arXiv:2210.07234. 1

[CW00] Richard Cleve and John Watrous. Fast parallel circuits for the quantum Fourier transform. (arXiv:quant-ph/0006004), June 2000. arXiv:quant-ph/0006004. 1.2

[DCEL09] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1), July 2009. URL: http://dx.doi.org/10.1103/PhysRevA.80.012304, doi:10.1103/physreva.80.012304. D

[DHB22] Alexander M. Dalzell, Nicholas Hunter-Jones, and Fernando G. S. L. Brandão. Random Quantum Circuits Anticoncentrate in Log Depth. *PRX Quantum*, 3(1):010333, March 2022. arXiv:2011.12277, doi:10.1103/PRXQuantum.3.010333. 1.1

[DLT02] D.P. DiVincenzo, D.W. Leung, and B.M. Terhal. Quantum data hiding. *IEEE Transactions on Information Theory*, 48(3):580–598, March 2002. URL: http://dx.doi.org/10.1109/18.985948, doi:10.1109/18.985948. D

[DNS+22] Abhinav Deshpande, Pradeep Niroula, Oles Shtanko, Alexey V. Gorshkov, Bill Fefferman, and Michael J. Gullans. Tight Bounds on the Convergence of Noisy Random Circuits to the Uniform Distribution. *PRX Quantum*, 3(4):040329, December 2022. arXiv:2112.00716, doi:10.1103/PRXQuantum.3.040329. 1.1

[DP23] Nicolas Delfosse and Adam Paetznick. Spacetime codes of clifford circuits, 2023. URL: https://arxiv.org/abs/2304.05943, arXiv:2304.05943. 1.4.1

[EK09] Bryan Eastin and Emanuel Knill. Restrictions on Transversal Encoded Quantum Gate Sets. *Phys. Rev. Lett.*, 102(11):110502, March 2009. doi:10.1103/PhysRevLett.102.110502. 1

[FRD+23] Enrico Fontana, Manuel S. Rudolph, Ross Duncan, Ivan Rungger, and Cristina Cîrstoiu. Classical simulations of noisy variational quantum circuits, 2023. URL: https://arxiv.org/abs/2306.05400, arXiv:2306.05400. 1.2, A

[FT16] Keisuke Fujii and Shuhei Tamate. Computational quantum-classical boundary of noisy commuting quantum circuits. *Sci Rep*, 6, May 2016. arXiv:1406.6932, doi:10.1038/srep25598. 1, 1.2

[GD18] Xun Gao and Luming Duan. Efficient classical simulation of noisy quantum computation. October 2018. arXiv:1810.03176. 1.1, 1.2, 1.3, A

[GGCT24] Guillermo González-García, J. Ignacio Cirac, and Rahul Trivedi. Pauli path simulations of noisy quantum circuits beyond average case, 2024. URL: https://arxiv.org/abs/2407.16068, arXiv:2407.16068. 1.2

[GWD17] Xun Gao, Sheng-Tao Wang, and L.-M. Duan. Quantum Supremacy for Simulating a Translation-Invariant Ising Spin Model. *Physical Review Letters*, 118(4), January 2017. doi:10.1103/PhysRevLett.118.040502. 1

[HE23] Dominik Hangleiter and Jens Eisert. Computational advantage of quantum random sampling. *Rev. Mod. Phys.*, 95(3):035001, July 2023. arXiv:2206.04079, doi:10.1103/RevModPhys.95.035001. 1.1

[HKB+24] Dominik Hangleiter, Marcin Kalinowski, Dolev Bluvstein, Madelyn Cain, Nishad Maskara, Xun Gao, Aleksander Kubica, Mikhail D. Lukin, and Michael J. Gullans. Fault-tolerant compiling of classically hard IQP circuits on hypercubes. (arXiv:2404.19005), April 2024. arXiv:2404.19005. 1, 1.4.1, 1.4.6, 4.2, 16, D

[HKS+17] Craig S. Hamilton, Regina Kruse, Linda Sansoni, Sonja Barkhofen, Christine Silberhorn, and Igor Jex. Gaussian Boson Sampling. *Phys. Rev. Lett.*, 119(17):170501, October 2017. arXiv:1612.01199, doi:10.1103/PhysRevLett.119.170501. 1

[Kri14] Michael Krivelevich. The phase transition in site percolation on pseudo-random graphs. *arXiv preprint arXiv:1404.5731*, 2014. B

[KRUdW08] Julia Kempe, Oded Regev, Falk Unger, and Ronald de Wolf. Upper bounds on the noise threshold for fault-tolerant quantum computing, 2008. URL: https://arxiv.org/abs/0802.1464, arXiv:0802.1464. 1.2

[MAG+24] Antonio Anna Mele, Armando Angrisani, Soumik Ghosh, Sumeet Khatri, Jens Eisert, Daniel Stilck França, and Yihui Quek. Noise-induced shallow circuits and absence of barren plateaus, 2024. URL: https://arxiv.org/abs/2403.13927, arXiv:2403.13927. 1.2, A

[Oh24] Changhun Oh. Classical simulability of constant-depth linear-optical circuits with noise. (arXiv:2406.08086), June 2024. arXiv:2406.08086. 1.2, 1.3

[PBG20] Hakop Pashayan, Stephen D. Bartlett, and David Gross. From estimation of quantum probabilities to simulation of quantum circuits. *Quantum*, 4:223, January 2020. arXiv:1712.02806, doi:10.22331/q-2020-01-13-223. 1

[PV10] M B Plenio and S Virmani. Upper bounds on fault tolerance thresholds of noisy clifford-based quantum computers. *New Journal of Physics*, 12(3):033012, March 2010. URL: http://dx.doi.org/10.1088/1367-2630/12/3/033012, doi:10.1088/1367-2630/12/3/033012. 1.2

[Rom08] Steven Roman. *Metric Spaces*, pages 301–324. Springer New York, New York, NY, 2008. doi:10.1007/978-0-387-72831-5_12. B

[RWL24] Joel Rajakumar, James D. Watson, and Yi-Kai Liu. Polynomial-Time Classical Simulation of Noisy IQP Circuits with Constant Depth. (arXiv:2403.14607), March 2024. arXiv:2403.14607. 1.2, 1.3, 1.4.3, 1.4.5, 6, 7, C

[SFG21] Daniel Stilck França and Raul García-Patrón. Limitations of optimization algorithms on noisy quantum devices. *Nat. Phys.*, 17(11):1221–1227, November 2021. doi:10.1038/s41567-021-01356-3. 1.2

[SHEF25] Ryotaro Suzuki, Jonas Haferkamp, Jens Eisert, and Philippe Faist. Quantum complexity phase transitions in monitored random circuits. *Quantum*, 9:1627, February 2025. URL: http://dx.doi.org/10.22331/q-2025-02-10-1627, doi:10.22331/q-2025-02-10-1627. 1.2

[SHH24] Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. Random unitaries in extremely low depth. (arXiv:2407.07754), July 2024. arXiv:2407.07754. 1

[SRP+21] James R. Seddon, Bartosz Regula, Hakop Pashayan, Yingkai Ouyang, and Earl T. Campbell. Quantifying Quantum Speedups: Improved Classical Simulation From Tighter Magic Monotones. *PRX Quantum*, 2(1):010345, March 2021. doi:10.1103/PRXQuantum.2.010345. 1.2

[SYGY24] Thomas Schuster, Chao Yin, Xun Gao, and Norman Y. Yao. A polynomial-time classical algorithm for noisy quantum circuits. (arXiv:2407.12768), July 2024. arXiv:2407.12768. 1.1, 1.2, A, D

[TD04] Barbara M. Terhal and David P. DiVincenzo. Adaptive Quantum Computation, Constant Depth Quantum Circuits and Arthur-Merlin Games. *Quant. Inf. Comp.*, 4(2):134–145, 2004. arXiv:quant-ph/0205133. 1

[TTT21] Yasuhiro Takahashi, Yuki Takeuchi, and Seiichiro Tani. Classically simulating quantum circuits with local depolarizing noise. *Theoretical Computer Science*, 893:117–132, 2021. URL: https://www.sciencedirect.com/science/article/pii/S0304397521004291, doi:10.1016/j.tcs.2021.07.025. 1.2

[VHP05] S. Virmani, Susana F. Huelga, and Martin B. Plenio. Classical simulability, entanglement breaking, and quantum computation thresholds. *Physical Review A*, 71(4), April 2005. URL: http://dx.doi.org/10.1103/PhysRevA.71.042328, doi:10.1103/physreva.71.042328. 1.2

[Web16] Zak Webb. The clifford group forms a unitary 3-design, 2016. URL: https://arxiv.org/abs/1510.02769, arXiv:1510.02769. D

[YJS19] Mithuna Yoganathan, Richard Jozsa, and Sergii Strelchuk. Quantum advantage of unitary Clifford circuits with magic state inputs. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 475(2225):20180427, May 2019. doi:10.1098/rspa.2018.0427. 1, 14, D

[ZKGG16] Huangjun Zhu, Richard Kueng, Markus Grassl, and David Gross. The Clifford group fails gracefully to be a unitary 4-design. *arXiv:1609.08172 [quant-ph]*, September 2016. arXiv:1609.08172. D