

Bell sampling from quantum circuits

Dominik Hangleiter^{*} and Michael J. Gullans[†]

Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park, Maryland 20742, USA

(Dated: February 2, 2024)

A central challenge in the verification of quantum computers is benchmarking their performance as a whole and demonstrating their computational capabilities. In this work, we find a universal model of quantum computation, *Bell sampling*, that can be used for both of those tasks and thus provides an ideal stepping stone towards fault-tolerance. In Bell sampling, we measure two copies of a state prepared by a quantum circuit in the transversal Bell basis. We show that the Bell samples are classically intractable to produce and at the same time constitute what we call a *circuit shadow*: from the Bell samples we can efficiently extract information about the quantum circuit preparing the state, as well as diagnose circuit errors. In addition to known properties that can be efficiently extracted from Bell samples, we give two new and efficient protocols, a test for the depth of the circuit and an algorithm to estimate a lower bound to the number of T gates in the circuit. With some additional measurements, our algorithm learns a full description of states prepared by circuits with low T -count.

Introduction. As technological progress on fault-tolerant quantum processors continues, a central challenge is to demonstrate their computational advantage and to benchmark their performance as a whole. Quantum random sampling experiments serve this double purpose [1–4] and have arguably surpassed the threshold of quantum advantage [5–10]. However, this approach currently suffers several drawbacks. Most importantly, it can only serve its central goals—benchmarking and certification of quantum advantage—in the classically simulable regime. This deficiency arises because evaluating the performance benchmark, the *cross-entropy benchmark*, requires a classical simulation of the ideal quantum computation. What is more, the cross-entropy benchmark suffers from various problems related to the specific nature of the physical noise in the quantum processor [9, 11, 12], and yields limited information about the underlying quantum state. More generally, in near-term quantum computing without error correction, we lack many tools for validating a given quantum computation just using its output samples.

In this work, we consider *Bell sampling*, a universal model of quantum computation in which two identical copies of a state prepared by a quantum circuit are measured in the transversal Bell basis, see Fig. 1. We show that, in this model, the outcomes are (i) simultaneously classically intractable to produce on average over universal random circuits under a standard assumption, (ii) yield diagnostic information about the underlying quantum state, and (iii) allow for detecting and correcting certain errors in the state preparation. Bell sampling from random universal quantum circuits thus overcomes the central practical problems of quantum random sampling as a means to benchmark and demonstrate the computational advantage of near-term quantum processors. It also serves as a stepping stone towards fault-tolerant quantum advantage: not only can we naturally detect certain errors from the Bell samples, but the protocol is also compatible with stabilizer codes in the sense that the Bell measurement between code blocks is transversal for such codes and allows for the fault-tolerant extraction of all error syndromes. Effectively, we may think of

the Bell samples as classical *circuit shadows*, in analogy to the notion of state shadows coined by Aaronson [13] and Huang *et al.* [14], since we can efficiently extract specific information about the generating circuit or a family of generating circuits from them.

Technically, we make the following contributions. We provide complexity-theoretic evidence for the classical intractability of Bell sampling from random universal quantum circuits, following an established hardness argument [4, 15, 16]. We introduce a new test to verify the depth of quantum circuits. Here, we make use of the fact that from the Bell basis samples one can compute correlation properties of the two copies and in particular a swap test on any subsystem. We observe that we can compare the measured average subsystem entropy to the maximal value achievable by a bounded-depth quantum circuit in a given architecture in order to estimate a lower bound on the depth of the circuit. For random circuits, we can refine this test by making use of their average entanglement properties as represented by the Page curve [17, 18]. We further show that the Bell samples can be used to efficiently measure the stabilizer nullity—a magic monotone [19]—and give a protocol to efficiently learn a full description of any quantum state that can be prepared by a circuit with low T -count. Here, we build on a result by Montanaro [20], who has shown that stabilizer states can be learned from Bell samples. Finally, we give a protocol for detecting errors in the state preparation based only on the properties of the Bell samples.

Of course, the idea to sample in the Bell basis to learn about properties of quantum states is as old as the theory of quantum information itself and has found many applications in quantum computing, including learning stabilizer states [20], testing stabilizerness [21], measuring magic [22, 23], and quantum machine learning [24]. The novelty of our approach is to view Bell sampling as a computational model. We then ask the question: What can we learn from the Bell samples about the circuit preparing the underlying quantum state?

Bell sampling. We begin by defining the Bell sampling protocol and noting some simple properties that will be useful in the remainder of this work. Consider a quantum circuit C

^{*} mail@dhangleiter.eu

[†] mgullans@umd.edu




FIG. 1. **The Bell sampling protocol.** In the Bell sampling protocol we prepare the quantum state $C|0^n\rangle \otimes C|0^n\rangle$ using a quantum circuit C , and measure all qubits transversally in the Bell basis across the bipartition of the system.

acting on n qubits, and define the Bell basis of two qubits as

$$|\sigma_r\rangle = (\sigma_r \otimes \mathbb{1})|\Phi^+\rangle, \text{ where } |\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}, \quad (1)$$

and for $r \in \{0, 1\}^2$ we identify

$$\sigma_{00} = \mathbb{1}, \quad \sigma_{01} = X, \quad \sigma_{10} = Z, \quad \sigma_{11} = i\sigma_{01}\sigma_{10} = Y. \quad (2)$$

The Bell sampling protocol proceeds as follows, see Fig. 1.

1. Prepare $|C\rangle := |C\rangle \otimes |C\rangle := C|0^n\rangle \otimes C|0^n\rangle$.
2. Measure all qubit pairs $(i, i+n)$ for $i \in [n] := \{1, 2, \dots, n\}$ in the Bell basis, yielding an outcome $r \in \{0, 1\}^{2n}$.

It is easy to see that the distribution of the outcomes r can be written as

$$P_C(r) = \frac{1}{2^n} | \langle C | \sigma_r | \overline{C} \rangle |^2 \quad (3)$$

where $\sigma_r = \sigma_{r_1 r_{n+1}} \otimes \sigma_{r_2 r_{n+2}} \otimes \dots \otimes \sigma_{r_n r_{2n}}$ is the n -qubit Pauli matrix corresponding to the outcome $r = (r_1, r_2, \dots, r_{2n})$, and \overline{C} denotes complex conjugation of C . In order to perform the measurement in the Bell basis, we need to apply a depth-1 quantum circuit consisting of n transversal cnot -gates followed by Hadamard gates on the control qubits and a measurement of all qubits in the computational basis.

Our first observation is that Bell sampling is a universal model of quantum computation. In particular, we show that Bell sampling from *random circuits* is classically intractable under certain complexity-theoretic assumptions. At the same time, we also show that we can use *the very same samples* to efficiently infer properties of and detect errors in the state preparation. The Bell samples can thus simultaneously be used for quantum computation and act as a classical shadow of the quantum state preparation that may be used to characterize a quantum device.

Computational complexity. We first show that Bell sampling is a universal model of quantum computation. To show this, we observe that we can estimate both the sign and the magnitude of $\langle C | Z | C \rangle$ for any quantum circuit C from Bell samples from a circuit $C'(C)$ in which we use a variant of Ramsey interferometry with a single ancilla qubit in each

copy of the circuit, see the Supplementary Material (SM) [25]. We then show that approximately sampling from the Bell sampling distribution P_C is classically intractable on average for universal random quantum circuits C with $\Omega(n^2)$ gates in a brickwork architecture (as depicted in Fig. 1), assuming certain complexity-theoretic conjectures are satisfied, via a standard proof technique, see the SM [25] for details. The argument puts the complexity-theoretic evidence for the hardness of Bell sampling from random quantum circuits in linear depth on a par with that for standard universal circuit sampling [1, 5, 26–29].

Bell samples as classical circuit shadows. Samples in the computational basis—while difficult to produce for random quantum circuits—yield very little information about the underlying quantum state. In particular, the problem of verification is essentially unsolved since the currently used methods require exponential computing time. In contrast, from the Bell samples, we can *efficiently* infer many properties of the quantum state preparation $|C\rangle \otimes |C\rangle$. Known examples include the overlap $\text{tr}[\rho\sigma]$ of a state preparation $\rho \otimes \sigma$ via a swap test, the magic of the state $|C\rangle$ [22], and the outcome of measuring any Pauli operator $P \otimes P$ [30]. Here, we add two new properties to this family. We give efficient protocols for testing the depth of low-depth quantum circuits, for testing its magic, and for learning quantum states that can be prepared by a circuit with low T -count.

Let us begin by recapping how a swap test can be performed using the Bell samples, and observing some properties that are useful in the context of benchmarking random quantum circuits. To this end, write the two-qubit swap operator $\mathbb{S} = P_{\sqrt{2}} - P_{\wedge^2}$ as the difference between the projectors onto the symmetric subspace $P_{\sqrt{2}} = |\sigma_{00}\rangle\langle\sigma_{00}| + |\sigma_{01}\rangle\langle\sigma_{01}| + |\sigma_{10}\rangle\langle\sigma_{10}|$ and the antisymmetric subspace $P_{\wedge^2} = |\sigma_{11}\rangle\langle\sigma_{11}|$. The overlap $\text{tr}[\rho\sigma] = \text{tr}[(\rho \otimes \sigma)\mathbb{S}]$ can then be directly estimated up to error ϵ from $M \in O(1/\epsilon^2)$ Bell samples as

$$\frac{1}{M} (|\{r : \pi_Y(r) = 0\}| - |\{r : \pi_Y(r) = 1\}|). \quad (4)$$

For quantum state preparations $\rho \otimes \rho$, the overlap quantifies the purity $\text{tr}[\rho^2]$ of ρ . Using randomized compiling implemented independently on two copies of a fixed circuit, we can convert experimentally relevant noise on the two copies into an effective Pauli channel [31, 32]. Errors also decohere into Pauli channels in repeated rounds of syndrome extraction in stabilizer codes [33, 34]. At low noise rates $\eta \ll 1$ we can then approximate $\rho \approx (1 - \eta)|C\rangle\langle C| + \eta\sigma$, and a simple calculation shows that the purity can be used to estimate the fidelity from the relation $1 - \text{tr}[\rho^2] = 2(1 - F) + O(\eta^2)$, where $F = \langle C | \rho | C \rangle$. Moreover, in the case of random circuits, there is an exact mapping between average fidelity and purity for any noise rate; see the SM [25] for details.

We can compare this to other means of estimating the fidelity of a quantum state. One method that has been widely used recently is cross-entropy benchmarking (XEB), which uses only classical samples from computational-basis measurements [5, 35]. This method can yield reliable estimates of the fidelity of the underlying quantum state under a weaker version of the white-noise approximation [5, 36]. But while




FIG. 2. **Depth-dependent Page curves.** (a) The maximal subsystem entanglement entropy depends on the circuit architecture and depth (shades of blue) until the half-cut entanglement reaches its maximal value given by $n/2$. We measure the subsystem entropy at half-cuts to obtain the maximal sensitivity to different circuit depths. (b) We detect errors in the Bell samples by detecting strings that lead to a non-zero estimate of the purity of ρ .

Bell sampling is computationally and sample efficient, XEB requires computation of ideal output probabilities of C , making it infeasible for already moderate numbers of qubits. Another means for fidelity estimation, shadow fidelity estimation [14, 37], requires implementing deep Clifford circuits and is only practical for specific states such as stabilizer states [38]. In the SM [25], we also discuss the relation of Bell sampling to other means of verifying quantum computations.

In the following, we will assume that the purity test has succeeded, and resulted in a value close to unity.

Depth test. We now describe a Bell sampling protocol to measure the depth of a quantum circuit C which is promised to be implemented in a fixed architecture, i.e., with gates applied in layers according to a certain pattern. The basic idea underlying the depth test is to use swap tests on subsystems of different sizes in order to obtain estimates of subsystem purities. For a subsystem A of $[n]$, the subsystem purity is given by $P_A(\rho) = \text{tr}[\rho_A^2]$, where $\rho_A = \text{tr}_{A^c}[\rho]$ is the reduced density matrix on subsystem $A \subset [n]$. It can be estimated from the fraction of outcome strings with even Y -parity $\pi_Y(r_A)$ on the substrings $r_A = (r_i, r_{n+i})_{i \in A}$.

Our test is based on the observation that the amount of entanglement generated by quantum circuits on half-cuts reaches a depth-dependent maximal value until it saturates at a circuit depth that depends on the dimensionality of the circuit architecture, see Fig. 2(a) for an illustration. In order to lower-bound the depth of a circuit family we choose a subsystem size at which the distinguishability between different depths is maximal. This is typically the case at half-cuts, where the Rényi-2 entanglement entropy $S_A(\rho) = -\log P_A(\rho)$ can be at most $n/2$. At the same time, the entanglement entropy is bounded as a function of depth $S_A(d) \leq d|\partial A|$, where ∂A is the number of gates applied across the boundary of A in every layer of the circuit. We now compute an empirical estimate $\hat{S}_{n/2}$ of $S_A(|C\rangle\langle C|)$ for a size- $n/2$ subsystem A using the Bell samples. In order to obtain a lower bound on the depth of the circuit C generating a given set of Bell samples, we com-

pute the maximum d such that $\hat{S}_{n/2} - \epsilon \geq d \cdot |\partial A|$ up to an error tolerance ϵ depending on the number of Bell samples. We can further refine this test for random quantum circuits by exploiting their average subsystem entanglement properties, known as the *Page curve* [17]. Depth-dependent Page curves have been computed analytically [18] and numerically [39] for a few circuit architectures and random ensembles.

We remark that these entanglement-based tests rely on universal features of quantum chaotic dynamics. As a result, they are also expected to be applicable to generic Hamiltonian dynamics, similar to how ideas for standard quantum random sampling have recently been extended to this case [40, 41].

Magic test and Clifford+ T learning algorithm. Another primitive that can be exploited in property tests of quantum states using the Bell samples is the fact that for stabilizer states $|S\rangle$, the Bell distribution is supported on a coset of the stabilizer group of $|S\rangle$ [20]. Leveraging this property allows for efficiently learning stabilizer states [20], testing stabilizerness [21], learning circuits with a single layer of T -gates [42] and estimating measures of magic [22, 23]. Here, we describe a simple, new protocol that, from the Bell samples, allows us to efficiently estimate the stabilizer nullity, a magic monotone [19], and learn states that can be prepared by quantum circuits with $t \in O(\log n)$ T -gates.

Our learning algorithm proceeds in two steps. In the first step, we find a compression of the non-Clifford part of the circuit, similarly to Refs. [43, 44]. To achieve this, using Bell difference sampling [21], we find a Clifford unitary U_C corresponding to a subspace $C \subset \mathbb{F}_2^{2n}$ such that $U_C|\psi\rangle$ has high fidelity with $|x\rangle|\varphi\rangle$ for some computational-basis state $|x\rangle$ on the first $\dim(C)$ qubits, and a state $|\varphi\rangle$ on the remaining qubits containing the non-Clifford information. The dimension of C satisfies $\dim(C) \geq n - t$. The number of T -gates required to prepare $|\psi\rangle$ is therefore lower-bounded by the stabilizer nullity $M(|\psi\rangle) := n - \dim(C)$, which is a magic monotone [19]. We show that only $O(n/\epsilon)$ Bell samples are sufficient to ensure that $|\psi\rangle$ is ϵ -close to a state with exact stabilizer nullity given by the estimate \hat{M} of $M(|\psi\rangle)$. To the best of our knowledge this is the most efficient way of measuring the magic of a quantum state to date.

In the second step of the learning algorithm, we characterize the state $|\varphi\rangle$ on the remaining $n - \dim(C) \leq t$ qubits using pure-state tomography, for example via the scheme of Ref. [45], giving an estimate $|\hat{\varphi}\rangle$. The output of the algorithm is then given by a classical description of $|\hat{\psi}\rangle = U_C|x\rangle|\hat{\varphi}\rangle$. The learning algorithm runs in polynomial time and succeeds with high probability in learning an ϵ -approximation to $|\psi\rangle$ in fidelity using $O(n/\epsilon)$ Bell samples and $O(2^t/\epsilon^2)$ measurements to perform tomography of $|x\rangle|\hat{\varphi}\rangle$.

Using Clifford+ T simulators [e.g. 46–48] we can now produce samples from and compute outcome probabilities of $|\hat{\psi}\rangle$ in time $O(2^t)$. We note that the exponential scaling in t is asymptotically optimal since the description of a state with stabilizer nullity t has $2^t + n - t$ real parameters. Our algorithm generalizes to arbitrary non-Clifford gates.

To summarize, we have given efficient ways to extract properties of the circuit C —its depth and an efficient circuit description for circuits with low T -count—using only

a small number of Bell samples. Further properties of $|C\rangle$ that can be efficiently extracted from the Bell samples include the expectation values of any diagonal two-copy observables $A = \sum_r a_r |\sigma_r\rangle\langle\sigma_r|$ and different measures of magic [22]. The Bell samples thus serve as an efficient classical shadow of C .

Error Detection and Correction In the last part of this paper, we discuss another appealing feature of Bell samples: we can perform error detection and correction. The idea that redundantly encoding quantum information in many copies of a quantum state allows error detection goes back to the early days of quantum computing. Already in 1996, Barenco *et al.* [49] have shown that errors can be reduced by symmetrizing many copies of a noisy quantum state. More recently Refs. [50–52] used measurements on multiple copies to suppress errors in expectation value estimation. In our two-copy setting, some simple *single-sample* error detection properties follow immediately from the tests in the previous section.

First, we observe that an outcome in the antisymmetric subspace, i.e., an outcome r with $\pi_Y(r) = 1$, is certainly due to an error. We can thus reduce the error in the sampled distribution by discarding such outcomes. We show in the SM [25] that such error detection reduces the error rate of a white-noise model by approximately a factor of 2. Quantum computations in the Bell sampling model with error detection can thus achieve equal fidelities to circuit model computations, where no error detection is possible, in spite of the factor of 2 in qubit overhead.

Second, we note that Bell samples are compatible with stabilizer codes. For such codes, the Bell measurement between code blocks is a transversal measurement, and allows to extract the syndrome $\sigma \otimes \sigma$ for $\sigma \in P_n$ in the stabilizer of the code [30]. If a detectable/correctable error occurred in one of the code blocks, this syndrome detects/identifies that error up to stabilizer equivalence. The fact that the Bell measurement is transversal implies that an error in the Bell measurement does not spread, so that local error channels or coherent errors in the entangling `cnot` gates in the Bell measurement reduce the overall measurement fidelity by $(1 - \epsilon)^n$, where ϵ is the error rate per Bell pair. Bell sampling is thus feasible in the regime of $\epsilon \ll 1/n$. We also note that antisymmetric errors in the Bell measurement are detectable.

Finally, we observe that quadratic error suppression is possible for estimating the expectation values of diagonal two-copy observable A , through the estimate $\text{tr}[A\rho^{\otimes 2}]/\text{tr}[\mathbb{S}\rho^{\otimes 2}]$, similar to virtual distillation [50–52]. Specifically, this is true for estimating the expectation $\langle E_0 | P | E_0 \rangle^2$ of a Pauli observable P in the ground state $|E_0\rangle$ of a gapped Hamiltonian by choosing $A = (P \otimes P)\mathbb{S}$, see the SM [25] for details.

Discussion and outlook. In this work, we have considered Bell sampling as a model of quantum computation. We have shown that many properties of the quantum circuit preparing the underlying state can be extracted efficiently, and that in particular certain errors in the state preparation can be detected from single shots. Based on this, we have argued that the Bell samples act as classical circuit shadows. Since Bell sampling is universal this allows us to perform universal quantum computations whose outputs also yield information about the quantum circuit. This makes Bell sampling an interesting

computational model, and our main focus in this work is to establish this fact.

We leave it as an open question how much overhead is required when performing computations with Bell sampling. While our BQP-completeness proof requires an additional ancilla qubit, it is conceivable, that one can encode an n -qubit computation in the circuit model using two copies of an m -qubit system with $m \leq n$ via Bell sampling.

Bell sampling is not only interesting conceptually, however. It is also realistic. Since the Bell basis measurement requires only transversal `cnot` and single-qubit gates, it can be naturally implemented in unit depth on various quantum processor architectures with long-range connectivity. These include in particular ion traps [53] and Rydberg atoms in optical tweezers [54]. It is more challenging to implement Bell sampling in geometrically local architectures such as superconducting qubits [5]. In such architectures, one can interleave the two copies in a geometrically local manner such that the Bell measurement is a local circuit; however, this comes at the cost of additional layers of SWAP gates for every unit of circuit depth. Alternatively, one can use looped pipeline architectures to implement the Bell measurement [55].

But is Bell sampling also practical in the near term? To satisfactorily answer this question, various sources of noise need to be analyzed in detail—tasks we defer to future work but mention here. For some of our protocols, including the purity test and the error detection protocols we discuss the effect of noise sources on the state preparation. But how severely does measurement noise affect the outcomes? In other instances, including the depth and magic test, and the low- T count learning algorithms we have restricted ourselves to (nearly) pure state preparations. We can at least certify that these algorithms are applicable because the purity of the state preparation can be independently checked. But in currently realistic scenarios, the state preparation of deep circuits will never be pure. An important question is therefore whether we can formulate noise-robust versions of these protocols.

While we have exploited the purity of the state $|C\rangle$ in our error detection protocol, it is an interesting question whether it is possible to detect additional errors from the Bell samples. For instance, it might be possible to exploit the fact that the subsystem purity of the target state is low for large subsystems, see Fig. 2.

We have shown that classically simulating the Bell sampling protocol with universal random circuits is classically intractable. An exciting question in this context is whether the complexity of Bell sampling might be more noise robust than computational-basis sampling in the asymptotic scenario. For universal circuit sampling in the computational basis Gao and Duan [56] and Aharonov *et al.* [57] developed an algorithm that simulates sufficiently deep random circuits with a constant noise rate in polynomial time. In the Supplementary Material [25] we give some initial evidence that this simulation algorithm fails for Bell measurements. If the hardness of Bell sampling indeed turns out to be robust to large amounts of circuit noise, we face the exciting prospect of a scalable quantum advantage demonstration with classical validation and error mitigation.

Note: While finalizing this work, we became aware of Refs. [58, 59], where the authors independently report algorithms similar to the one we present above for learning quantum states generated by circuits with low T -count. After this work was completed, we collaborated on the physical implementation of Bell sampling in a logical qubit processor, illustrating the feasibility of our results to near-term devices [60].

Acknowledgements D.H. warmly thanks Abhinav Deshpande and Ingo Roth for helpful discussions that aided in the proofs of Lemma 2 and Lemma 4, respectively. We are

also grateful to Dolev Bluvstein, Maddie Cain, Bill Fefferman, Xun Gao, Soumik Ghosh, Alexey Gorshkov, Vojtěch Havlíček, Markus Heinrich, Marcel Hinsche, Marios Ioannou, Marcin Kalinowski, Mikhail Lukin and Brayden Ware for discussions. This research was supported in part by NSF QLCI grant OMA-2120757 and Grant No. NSF PHY-1748958 through the KITP program on “Quantum Many-Body Dynamics and Noisy Intermediate-Scale Quantum Systems.” D.H. acknowledges funding from the US Department of Defense through a QuICS Hartree fellowship.

-
- [1] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, *Characterizing Quantum Supremacy in Near-Term Devices*, *Nature Phys* **14**, 595 (2018), [arxiv:1608.00263](#).
 - [2] Y. Liu, M. Otten, R. Bassirianjahromi, L. Jiang, and B. Fefferman, *Benchmarking Near-Term Quantum Computers via Random Circuit Sampling*, (2022), [arxiv:2105.05232](#).
 - [3] M. Heinrich, M. Kliesch, and I. Roth, *General Guarantees for Randomized Benchmarking with Random Quantum Circuits*, (2022), [arxiv:2212.06181](#).
 - [4] D. Hangleiter and J. Eisert, *Computational Advantage of Quantum Random Sampling*, *Rev. Mod. Phys.* **95**, 035001 (2023), [arxiv:2206.04079](#).
 - [5] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, *Quantum Supremacy Using a Programmable Superconducting Processor*, *Nature* **574**, 505 (2019).
 - [6] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, *Quantum Computational Advantage Using Photons*, *Science* **370**, 1460 (2020).
 - [7] Q. Zhu, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, C. Guo, C. Guo, S. Guo, L. Han, L. Hong, H.-L. Huang, Y.-H. Huo, L. Li, N. Li, S. Li, Y. Li, F. Liang, C. Lin, J. Lin, H. Qian, D. Qiao, H. Rong, H. Su, L. Sun, L. Wang, S. Wang, D. Wu, Y. Wu, Y. Xu, K. Yan, W. Yang, Y. Yang, Y. Ye, J. Yin, C. Ying, J. Yu, C. Zha, C. Zhang, H. Zhang, K. Zhang, Y. Zhang, H. Zhao, Y. Zhao, L. Zhou, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, *Quantum Computational Advantage via 60-Qubit 24-Cycle Random Circuit Sampling*, *Science Bulletin* **67**, 240 (2022), [arxiv:2109.03494](#).
 - [8] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins, A. E. Lita, T. Gerrits, S. W. Nam, V. D. Vaidya, M. Menotti, I. Dhand, Z. Vernon, N. Quesada, and J. Lavoie, *Quantum Computational Advantage with a Programmable Photonic Processor*, *Nature* **606**, 75 (2022).
 - [9] A. Morvan, B. Villalonga, X. Mi, S. Mandrà, A. Bengtsson, P. V. Klimov, Z. Chen, S. Hong, C. Erickson, I. K. Drozdov, J. Chau, G. Laun, R. Movassagh, A. Asfaw, L. T. A. N. Brandão, R. Peralta, D. Abanin, R. Acharya, R. Allen, T. I. Andersen, K. Anderson, M. Ansmann, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, A. Bilmes, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, J. Campero, H. S. Chang, B. Chiaro, D. Chik, C. Chou, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, A. D. T. Barba, S. Demura, A. Di Paolo, A. Dunsworth, L. Faoro, E. Farhi, R. Fatemi, V. S. Ferreira, L. F. Burgos, E. Forati, A. G. Fowler, B. Foxen, G. Garcia, E. Genois, W. Giang, C. Gidney, D. Gilboa, M. Giustina, R. Gosula, A. G. Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. Hansen, M. P. Harrigan, S. D. Harrington, P. Heu, M. R. Hoffmann, T. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, T. Khattar, M. Khezri, M. Kieferová, S. Kim, A. Kitaev, A. R. Klots, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. W. Lee, Y. D. Lensky, B. J. Lester, A. T. Lill, W. Liu, A. Locharla, F. D. Malone, O. Martin, S. Martin, J. R. McClean, M. McEwen, K. C. Miao, A. Mieszala, S. Montazeri, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, T. E. O’Brien, S. Omonije, A. Opremcak, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, D. M. Rhodes, C. Rocque, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, N. Shutty, V. Shvarts, V. Sivak, J. Skrzynny, W. C. Smith, R. D. Somma, G. Sterling, D. Strain, M. Szalay, D. Thor, A. Torres, G. Vidal, C. V. Heidweiller, T. White, B. W. K. Woo, C. Xing, Z. J. Yao, P. Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, N. Zhu, N. Zobrist, E. G. Rieffel, R. Biswas, R. Babbush, D. Bacon, J. Hilton, E. Lucero, H. Neven, A. Megrant, J. Kelly, I. Aleiner, V. Smelyanskiy, K. Kechedzhi, Y. Chen, and S. Boixo, *Phase Transition in Random Circuit Sampling*, (2023), [arxiv:2304.11119](#).
 - [10] Y.-H. Deng, Y.-C. Gu, H.-L. Liu, S.-Q. Gong, H. Su, Z.-J. Zhang, H.-Y. Tang, M.-H. Jia, J.-M. Xu, M.-C. Chen, H.-S. Zhong, J. Qin, H. Wang, L.-C. Peng, J. Yan, Y. Hu, J. Huang, H. Li, Y. Li, Y. Chen, X. Jiang, L. Gan, G. Yang, L. You,

- L. Li, N.-L. Liu, J. J. Renema, C.-Y. Lu, and J.-W. Pan, *Gaussian Boson Sampling with Pseudo-Photon-Number Resolving Detectors and Quantum Computational Advantage*, (2023), [arxiv:2304.12240](#).
- [11] X. Gao, M. Kalinowski, C.-N. Chou, M. D. Lukin, B. Barak, and S. Choi, *Limitations of Linear Cross-Entropy as a Measure for Quantum Advantage*, (2021), [arxiv:2112.01657](#).
- [12] B. Ware, A. Deshpande, D. Hangleiter, P. Niroula, B. Fefferman, A. V. Gorshkov, and M. J. Gullans, *A Sharp Phase Transition in Linear Cross-Entropy Benchmarking*, (2023), [arxiv:2305.04954](#).
- [13] S. Aaronson, *The Learnability of Quantum States*, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **463**, 3089 (2007), [arxiv:quant-ph/0608142](#).
- [14] H.-Y. Huang, R. Kueng, and J. Preskill, *Predicting Many Properties of a Quantum System from Very Few Measurements*, *Nature Physics* **16**, 1050 (2020).
- [15] S. Aaronson and A. Arkhipov, *The Computational Complexity of Linear Optics*, *Th. Comp.* **9**, 143 (2013).
- [16] M. J. Bremner, A. Montanaro, and D. J. Shepherd, *Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations*, *Physical Review Letters* **117**, 080501 (2016).
- [17] D. N. Page, *Average Entropy of a Subsystem*, *Phys. Rev. Lett.* **71**, 1291 (1993).
- [18] A. Nahum, J. Ruhman, S. Vijay, and J. Haah, *Quantum Entanglement Growth under Random Unitary Dynamics*, *Phys. Rev. X* **7**, 031016 (2017).
- [19] M. Beverland, E. Campbell, M. Howard, and V. Kliuchnikov, *Lower Bounds on the Non-Clifford Resources for Quantum Computations*, *Quantum Sci. Technol.* **5**, 035009 (2020).
- [20] A. Montanaro, *Learning Stabilizer States by Bell Sampling*, (2017), [arxiv:1707.04012](#).
- [21] D. Gross, S. Nezami, and M. Walter, *Schur–Weyl Duality for the Clifford Group with Applications: Property Testing, a Robust Hudson Theorem, and de Finetti Representations*, *Commun. Math. Phys.* **385**, 1325 (2021).
- [22] T. Haug and M. Kim, *Scalable Measures of Magic Resource for Quantum Computers*, *PRX Quantum* **4**, 010301 (2023).
- [23] T. Haug, S. Lee, and M. S. Kim, *Efficient Stabilizer Entropies for Quantum Computers*, (2023), [arxiv:2305.19152](#).
- [24] H.-Y. Huang, R. Kueng, and J. Preskill, *Information-Theoretic Bounds on Quantum Advantage in Machine Learning*, *Phys. Rev. Lett.* **126**, 190505 (2021).
- [25] The Supplemental Material is available at [\[INCLUDE URL\]](#), which includes the additional references [61?–67].
- [26] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, *On the Complexity and Verification of Quantum Random Circuit Sampling*, *Nature Phys* **15**, 159 (2019).
- [27] R. Movassagh, *Quantum Supremacy and Random Circuits*, (2020), [arxiv:1909.06210](#).
- [28] Y. Kondo, R. Mori, and R. Movassagh, *Quantum Supremacy and Hardness of Estimating Output Probabilities of Quantum Circuits*, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (2022) pp. 1296–1307, [arxiv:2102.01960](#).
- [29] H. Krovi, *Average-Case Hardness of Estimating Probabilities of Random Quantum Circuits with a Linear Scaling in the Error Exponent*, (2022), [arxiv:2206.05642](#).
- [30] D. Gottesman, *An Introduction to Quantum Error Correction and Fault-Tolerant Quantum Computation*, (2009), [arxiv:0904.2557](#).
- [31] J. J. Wallman and J. Emerson, *Noise Tailoring for Scalable Quantum Computation via Randomized Compiling*, *Phys. Rev. A* **94**, 052325 (2016).
- [32] A. Winick, J. J. Wallman, D. Dahlen, I. Hincks, E. Ospadov, and J. Emerson, *Concepts and Conditions for Error Suppression through Randomized Compiling*, (2022), [arxiv:2212.07500 \[quant-ph\]](#).
- [33] S. J. Beale, J. J. Wallman, M. Gutiérrez, K. R. Brown, and R. Laflamme, *Quantum Error Correction Decohere Noise*, *Phys. Rev. Lett.* **121**, 190501 (2018).
- [34] E. Huang, A. C. Doherty, and S. Flammia, *Performance of Quantum Error Correction with Coherent Errors*, *Phys. Rev. A* **99**, 022313 (2019).
- [35] J. Choi, A. L. Shaw, I. S. Madjarov, X. Xie, R. Finkelstein, J. P. Covey, J. S. Cotler, D. K. Mark, H.-Y. Huang, A. Kale, H. Pichler, F. G. S. L. Brandão, S. Choi, and M. Endres, *Preparing Random States and Benchmarking with Many-Body Quantum Chaos*, *Nature* **613**, 468 (2023).
- [36] A. M. Dalzell, N. Hunter-Jones, and F. G. S. L. Brandão, *Random Quantum Circuits Transform Local Noise into Global White Noise*, (2021), [arxiv:2111.14907](#).
- [37] M. Kliesch and I. Roth, *Theory of Quantum System Certification*, *PRX Quantum* **2**, 010201 (2021), [arxiv:2010.05925](#).
- [38] M. Ringbauer, M. Hinsche, T. Feldker, P. K. Faehrmann, J. Bermejo-Vega, C. Edmunds, L. Postler, R. Stricker, C. D. Marciniak, M. Meth, I. Pogorelov, R. Blatt, P. Schindler, J. Eisert, T. Monz, and D. Hangleiter, *Verifiable Measurement-Based Quantum Random Sampling with Trapped Ions*, (2023), [arxiv:2307.14424](#).
- [39] G. M. Sommers, D. A. Huse, and M. J. Gullans, *Crystalline Quantum Circuits*, (2023), [arxiv:2210.10808](#).
- [40] D. K. Mark, J. Choi, A. L. Shaw, M. Endres, and S. Choi, *Benchmarking Quantum Simulators Using Ergodic Quantum Dynamics*, *Phys. Rev. Lett.* **131**, 110601 (2023).
- [41] A. L. Shaw, Z. Chen, J. Choi, D. K. Mark, P. Scholl, R. Finkelstein, A. Elben, S. Choi, and M. Endres, *Benchmarking highly entangled states on a 60-atom analog quantum simulator*, [arXiv:2308.07914](#) (2023).
- [42] C.-Y. Lai and H.-C. Cheng, *Learning Quantum Circuits of Some ST Gates*, *IEEE Trans. Inform. Theory* **68**, 3951 (2022), [arxiv:2106.12524](#).
- [43] S. Arunachalam, S. Bravyi, C. Nirkhe, and B. O’Gorman, *The Parameterized Complexity of Quantum Verification*, (2022), [arxiv:2202.08119](#).
- [44] L. Leone, S. F. E. Oliviero, S. Lloyd, and A. Hamma, *Learning Efficient Decoders for Quasi-Chaotic Quantum Scramblers*, (2023), [arxiv:2212.11338](#).
- [45] M. Guta, J. Kahn, R. Kueng, and J. A. Tropp, *Fast State Tomography with Optimal Error Bounds*, [arXiv:1809.11162 \[quant-ph\]](#) (2018), [arxiv:1809.11162](#).
- [46] S. Bravyi, D. Browne, P. Calpin, E. Campbell, D. Gosset, and M. Howard, *Simulation of Quantum Circuits by Low-Rank Stabilizer Decompositions*, *Quantum* **3**, 181 (2019), [arxiv:1808.00128](#).
- [47] H. Pashayan, O. Reardon-Smith, K. Korzekwa, and S. D. Bartlett, *Fast Estimation of Outcome Probabilities for Quantum Circuits*, *PRX Quantum* **3**, 020361 (2022).
- [48] E. T. Campbell and M. Howard, *A Unified Framework for Magic State Distillation and Multi-Qubit Gate-Synthesis with Reduced Resource Cost*, *Phys. Rev. A* **95**, 022316 (2017), [arxiv:1606.01904](#).
- [49] A. Barenco, A. Berthiaume, D. Deutsch, A. Ekert, R. Jozsa, and C. Macchiavello, *Stabilisation of Quantum Computations by Symmetrisation*, (1996), [arxiv:quant-ph/9604028](#).
- [50] J. Cotler, S. Choi, A. Lukin, H. Gharibyan, T. Grover, M. E. Tai, M. Rispoli, R. Schittko, P. M. Preiss, A. M. Kaufman,

- M. Greiner, H. Pichler, and P. Hayden, *Quantum Virtual Cooling*, *Phys. Rev. X* **9**, 031013 (2019).
- [51] B. Koczor, *Exponential Error Suppression for Near-Term Quantum Devices*, *Phys. Rev. X* **11**, 031057 (2021).
- [52] W. J. Huggins, S. McArdle, T. E. O’Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, *Virtual Distillation for Quantum Error Mitigation*, *Phys. Rev. X* **11**, 041036 (2021).
- [53] A. Bermudez, X. Xu, R. Nigmatullin, J. O’Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller, *Assessing the Progress of Trapped-Ion Processors Towards Fault-Tolerant Quantum Computation*, *Phys. Rev. X* **7**, 041061 (2017).
- [54] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletić, and M. D. Lukin, *A Quantum Processor Based on Coherent Transport of Entangled Atom Arrays*, *Nature* **604**, 451 (2022).
- [55] Z. Cai, A. Siegel, and S. Benjamin, *Looped Pipelines Enabling Effective 3D Qubit Lattices in a Strictly 2D Device*, (2022), [arXiv:2203.13123](https://arxiv.org/abs/2203.13123).
- [56] X. Gao and L. Duan, *Efficient Classical Simulation of Noisy Quantum Computation*, (2018), [arxiv:1810.03176](https://arxiv.org/abs/1810.03176).
- [57] D. Aharonov, X. Gao, Z. Landau, Y. Liu, and U. Vazirani, *A Polynomial-Time Classical Algorithm for Noisy Random Circuit Sampling*, (2022), [arxiv:2211.03999](https://arxiv.org/abs/2211.03999).
- [58] S. Grewal, V. Iyer, W. Kretschmer, and D. Liang, *Efficient Learning of Quantum States Prepared With Few Non-Clifford Gates*, (2023), [arxiv:2305.13409](https://arxiv.org/abs/2305.13409).
- [59] L. Leone, S. F. E. Oliviero, and A. Hamma, *Learning T-Doped Stabilizer States*, (2023), [arxiv:2305.15398](https://arxiv.org/abs/2305.15398).
- [60] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. B. Ataiades, N. Maskara, I. Cong, X. Gao, P. S. Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, *Logical Quantum Processor Based on Reconfigurable Atom Arrays*, *Nature*, **1** (2023), [arxiv:2312.03982](https://arxiv.org/abs/2312.03982).
- [61] P. Erdős and A. Rényi, *Probabilistic Methods in Group Theory*, *J. Anal. Math.* **14**, 127 (1965).
- [62] S. Garnerone, T. R. de Oliveira, and P. Zanardi, *Typicality in Random Matrix Product States*, *Phys. Rev. A* **81**, 032336 (2010).
- [63] B. Collins, C. E. Gonzalez-Guillen, and D. Perez-Garcia, *Matrix Product States, Random Matrix Theory and the Principle of Maximum Entropy*, *Communications in Mathematical Physics* **320**, 663 (2013), [arxiv:1201.6324](https://arxiv.org/abs/1201.6324).
- [64] M. Fukuda and R. Koenig, *Typical Entanglement for Gaussian States*, *Journal of Mathematical Physics* **60**, 112203 (2019), [arxiv:1903.04126](https://arxiv.org/abs/1903.04126).
- [65] H. Zhu, R. Kueng, M. Grassl, and D. Gross, *The Clifford Group Fails Gracefully to Be a Unitary 4-Design*, [arXiv:1609.08172](https://arxiv.org/abs/1609.08172) [quant-ph] (2016), [arxiv:1609.08172](https://arxiv.org/abs/1609.08172).
- [66] B. W. Reichardt, F. Unger, and U. Vazirani, *Classical Command of Quantum Systems*, *Nature* **496**, 456 (2013).
- [67] U. M. Mahadev, *Classical Verification and Blind Delegation of Quantum Computations*, Ph.D. thesis, University of California, Berkeley (2018).

Supplementary Material for “Bell sampling from quantum circuits”

Dominik Hangleiter and Michael J. Gullans

CONTENTS

S1. BQP-completeness of Bell sampling	i
S2. Classical hardness of Bell sampling	i
A. Hiding	i
B. Average-case GapP hardness of approximating outcome probabilities	ii
1. Worst-case GapP hardness of approximating outcome probabilities	ii
2. Near-exact average-case hardness	iii
3. Anticoncentration	iv
S3. Learning circuit properties from Bell sampling	iv
A. Measuring purity and fidelity	v
1. Purity and Fidelity	v
2. Average fidelity	v
3. Relation to verified quantum computation	vi
B. Testing depth	vi
C. Learning a Clifford + T circuit	vii
1. The Bell distribution of low T -count quantum states	viii
2. The learning algorithm	viii
3. Correctness of Algorithm 4	ix
4. Correctness of Algorithm 3	ix
S4. Error detection and correction	x
A. Error reduction by error detection	x
B. Noise in the Bell measurement	x
C. Virtual distillation using the Bell samples	xi
S5. Applying the noisy simulation algorithm to Bell sampling	xi
A. Recap of the algorithm	xi
B. Strategy 1: Upper bounds on the trace distance	xii
C. Strategy 2: The argument in the Bell basis	xiii
References	xiv

S1. BQP-COMPLETENESS OF BELL SAMPLING

In this section, we show that Bell sampling is BQP-complete as a model of quantum computation in spite of the fact that we are restricting to circuits of the form $C \otimes C$ and measurements in the transversal Bell basis.

Lemma 1 (BQP-completeness). *Bell sampling is BQP-complete.*

Proof. Consider an arbitrary quantum circuit C . Then estimating the probability of measuring the first qubit of C in the $|1\rangle$ state up to additive precision is BQP-complete. The idea of the proof is to design a circuit C' such that we can infer $p_1 := \text{tr}[(|1\rangle\langle 1| \otimes \mathbb{1}_{n-1})C|0^n\rangle\langle 0^n|C^\dagger]$ from Bell samples

from C' . Let $\rho = \text{tr}_{[n]\setminus\{1\}}[C|0^n\rangle\langle 0^n|C^\dagger]$ be the reduced density matrix of the first qubit of the state $C|0\rangle$ before the measurement. Then our task is to estimate $p_1 = \text{tr}[\rho|1\rangle\langle 1|]$ from Bell sampling.

To achieve this, we make use of the fact that we can infer the square of any Pauli expectation value $\langle C|P|C\rangle^2$ up to additive precision from Bell samples from $|C\rangle \otimes |C\rangle$ for any circuit C . Starting from an arbitrary quantum circuit C we add an ancillary qubit (labeled by 0) before the first qubit and run the circuit $C' = e^{-i\frac{\pi}{8}(Z_0Z_1+Z_0)}(H \otimes C)$, where H denotes the Hadamard gate. The resulting marginal state on the ancillary qubit is then given by

$$\begin{aligned} & \text{tr}_{[n]}[e^{-i\frac{\pi}{8}(Z_0Z_1+Z_0)}(|+\rangle\langle +| \otimes \rho)e^{i\frac{\pi}{8}(Z_0Z_1+Z_0)} \\ &= \frac{1}{2} \left(\mathbb{1} + \left(\frac{1}{2} - \frac{1}{2} \langle C|Z_1|C\rangle \right) X + \left(\frac{1}{2} + \frac{1}{2} \langle C|Z_1|C\rangle \right) Y \right). \end{aligned} \quad (\text{S1})$$

From Bell samples from $|C'\rangle \otimes |C'\rangle$ we can now estimate $|\langle C'|X_0|C'\rangle| = |1 - \langle C|Z_1|C\rangle|/2 = p_1$. \square

S2. CLASSICAL HARDNESS OF BELL SAMPLING

In this section, we provide complexity-theoretic evidence that sampling from the Bell distribution P_C up to constant total-variation distance (TVD) error is classically intractable.

In order to show this, we follow a standard proof strategy, which has three main ingredients.

1. **Hiding:** The distributions over outcomes and circuit instances are interchangeable.
2. **Average-case GapP hardness** of approximating the output probabilities. While we cannot prove this in any instance, we typically provide evidence for it using three ingredients:
 - (a) Worst-case GapP hardness of approximation,
 - (b) Average-case GapP hardness of near-exact computation, and
 - (c) Anticoncentration.

We defer the reader to Section III of Ref. [1] for a detailed exposition of the proof strategy.

A. Hiding

Consider the Bell sampling distribution

$$P_C(r) = \frac{1}{2^n} |\langle \sigma_r | C \otimes C | 0^{2n} \rangle|^2 \quad (\text{S2})$$

$$= \frac{1}{2^n} |\langle \Phi^+ | (\sigma_r \otimes \mathbb{1})(C \otimes C) | 0^{2n} \rangle|^2. \quad (\text{S3})$$

Then, using that $|\Phi^+\rangle = \text{vec}(\mathbb{1})/\sqrt{2}$, we find that

$$(\sigma_r \otimes \mathbb{1})|\Phi^+\rangle = (\sqrt{\sigma_r} \otimes \sqrt{\sigma_r}^T)|\Phi^+\rangle, \quad (\text{S4})$$

and hence

$$P_C(r) = |\langle \Phi^+ | \sqrt{\sigma_r} C \otimes \sqrt{\sigma_r}^T C | 0^{2n} \rangle|^2, \quad (\text{S5})$$

so that for x such that $\sqrt{\sigma_r} \neq \sqrt{\sigma_r}^T$, we cannot write $P_C(r) = P'_C(0^n)$ for some $C' = C'(C, r)$. And indeed, while $\sqrt{X}^T = \sqrt{X}$ and $\sqrt{Z}^T = \sqrt{Z}$, we have $\sqrt{Y}^T \neq \sqrt{Y}$.

This means that hiding does not hold on the full output distribution. However, we can restrict ourselves to the *a priori* known support of the Bell sampling distribution, which is given by the symmetric subspace, characterized by $\pi_Y(r) = 0$. Indeed, we can view the fact that $\sqrt{Y}^T \neq \sqrt{Y}$ as a signature of the fact that the corresponding Bell state is not symmetric. Conversely, consider an even number of Y -outcomes. Specifically, for a two-qubit state with two Y -outcomes, we can explicitly check that $\sqrt{Y \otimes Y}^T = \sqrt{Y \otimes Y}$.

Hence, hiding holds on the symmetric subspace for any architecture that includes a layer of single-qubit gates that is invariant under $\sqrt{X}, \sqrt{Y}, \sqrt{Z}$ at the end of the circuit.

Furthermore, defining $C_r = \sqrt{\sigma_r} C$ we observe that

$$|\langle \sigma_r | C_s | 0^{2n} \rangle|^2 = |\langle \sigma_{r \oplus s} | C | 0^{2n} \rangle|^2, \quad (\text{S6})$$

since $\sqrt{\sigma_r} \sigma_s \sqrt{\sigma_r} \in P_n$, where P_n is the Pauli group with phases in $\{\pm 1, \pm i\}$.

B. Average-case GapP hardness of approximating outcome probabilities

1. Worst-case GapP hardness of approximating outcome probabilities

Lemma 2. *Given a quantum circuit C , it is GapP-hard to compute $2^{-n} |\langle C | \sigma_r | \bar{C} \rangle|^2$ up to relative error $< 1/2$.*

Proof. In order to perform the Bell basis measurement we apply $\text{cnot}(H \otimes \mathbb{1})$ across all pairs of qubits. Writing $|C\rangle = \sum_x c_x |x\rangle$, the pre-measurement state then transforms as

$$|C\rangle \otimes |C\rangle = \sum_{xy} c_x c_y |x\rangle |y\rangle \quad (\text{S7})$$

$$\xrightarrow{\text{cnot}^{\otimes n}} \sum_{xy} c_x c_y |x\rangle |x \oplus y\rangle \quad (\text{S8})$$

$$\xrightarrow{(H \otimes \mathbb{1})^{\otimes n}} \sum_{xyz} (-1)^{x \cdot z} c_x c_y |z\rangle |x \oplus y\rangle =: |C\rangle. \quad (\text{S9})$$

Consider a state $|C\rangle$ on $n+1$ qubits and the $(0^n 1, 0^{n+1})$ -

Bell amplitude

$$\langle 0^n 1 | \langle 0^{n+1} | C \rangle = \sum_{xyz} (-1)^{x \cdot z} c_x c_y \langle 0^n 1 | z \rangle \langle 0^{n+1} | x \oplus y \rangle \quad (\text{S10})$$

$$= \sum_x (-1)^{x \cdot 0^n 1} c_x^2 \quad (\text{S11})$$

$$= \sum_{x: x_n=0} c_x^2 - \sum_{x: x_n=1} c_x^2. \quad (\text{S12})$$

Let us now specify $|C\rangle = |C_f\rangle \propto \sum_x |x\rangle |f(x)\rangle$ up to normalization, where $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is an efficiently computable Boolean function. Observe that for $b \in \{0, 1\}$

$$(\mathbb{1} \otimes \langle b |) \sum_x |x\rangle |f(x)\rangle = \sum_{x: f(x)=b} |x\rangle. \quad (\text{S13})$$

For $b = 1$ the state (S13) is normalized by the square root of the number of accepting inputs of f , given by $\#f := |\text{Acc}(f)| \equiv |\{x : f(x) = 1\}|$, and for $b = 0$ by $\sqrt{2^n - \#f}$. Hence, the coefficients of $|C_f\rangle = \sum_{x \in \{0, 1\}^{n+1}} c_x |x\rangle$ are given by

$$c_{yb} \equiv (\langle y | \otimes \langle b |) |C_f\rangle = \begin{cases} 1/\sqrt{\#f}, & f(y) = b = 1 \\ 1/\sqrt{2^n - \#f}, & f(y) = b = 0 \\ 0 & \text{otherwise.} \end{cases} \quad (\text{S14})$$

Let us now add another qubit so that we have $n+2$ qubits in total. Given an arbitrary efficiently computable Boolean function $g: \{0, 1\}^n \rightarrow \{0, 1\}$ let us define

$$f_g: \{0, 1\}^{n+1} \rightarrow \{0, 1\} \quad (\text{S15})$$

$$f_g(y) = \begin{cases} 1 & \text{if } y = (x, g(x)) \\ 0 & \text{if } y = (x, \neg g(x)). \end{cases} \quad (\text{S16})$$

We now reversibly compute the function f_g obtaining $n+2$ -bit outcome strings $(y, f_g(y))$ with the property that $f_g(y) = 1$ if $y = (x, g(x))$. So the $(n+1)$ st qubit encodes the outcome of g while the last (i.e., $(n+2)$ nd) qubit encodes the outcome of f_g .

f_g can be efficiently computed: Let D be the circuit that maps $|x\rangle |b\rangle \rightarrow |x\rangle |g(x) \oplus b\rangle$. Then the quantum circuit $C = (\mathbb{1}_n \otimes \text{cnot}(X \otimes \mathbb{1}))(D^\dagger \otimes \mathbb{1})$ computes $|y\rangle |0\rangle \mapsto |y\rangle |f_g(y)\rangle$.

We observe that

$$|\text{Acc}(f_g)| = |\text{Rej}(f_g)| = 2^{n+1}/2 \quad (\text{S17})$$

and, moreover, we have¹

$$|\text{Acc}(g)| = |\{y \in \text{Acc}(f_g), y_n = 1\}| = |\{y \in \text{Rej}(f_g), y_n = 0\}|, \quad (\text{S18})$$

$$|\text{Rej}(g)| = |\{y \in \text{Acc}(f_g), y_n = 0\}| = |\{y \in \text{Rej}(f_g), y_n = 1\}|. \quad (\text{S19})$$

¹ We start labelling indices at 0 and hence x_n is the $n+1$ st bit of x .

To see the latter, just observe that

$$\begin{aligned} |\{y \in \text{Acc}(f_g), y_n = 1\}| &= |\{(x, g(x)) : g(x) = 1\}| \\ &= |\{(x, \neg g(x)) : \neg g(x) = 0\}| = |\text{Acc}(g)|. \end{aligned} \quad (\text{S20})$$

Let us consider the outcome string $(0^n 11, 0^{n+2})$ and compute the corresponding amplitude of the state $|\mathcal{C}_{f_g}\rangle = ((H \otimes \mathbb{1}) \cdot \text{cnot})^{\otimes(n+2)} |C\rangle \otimes |C\rangle$ as

$$\langle 0^n 11, 0^{n+2} | \mathcal{C}_{f_g} \rangle = \sum_{xyz} (-1)^{x \cdot z} c_x c_y \langle 0^n 11 | z \rangle \langle 0^{n+2} | x \oplus y \rangle \quad (\text{S21})$$

$$= \sum_x (-1)^{x \cdot (0^n 11)} c_x^2 \quad (\text{S22})$$

$$= \sum_{x: x_n=0, x_{n+1}=0} c_x^2 - \sum_{x: x_n=1, x_{n+1}=0} c_x^2 - \sum_{x: x_n=0, x_{n+1}=1} c_x^2 + \sum_{x: x_n=1, x_{n+1}=1} c_x^2 \quad (\text{S23})$$

$$\stackrel{(\text{S17})}{=} \frac{1}{2^n} (|\{y \in \text{Rej}(f_g) : y_n = 0\}| - |\{y \in \text{Rej}(f_g) : y_n = 1\}|) \quad (\text{S24})$$

$$+ \frac{1}{2^n} (|\{y \in \text{Acc}(f_g) : y_n = 1\}| - |\{y \in \text{Acc}(f_g) : y_n = 0\}|) \quad (\text{S25})$$

$$\stackrel{(\text{S18})}{=} \frac{2}{2^n} (|\text{Acc}(g)| - |\text{Rej}(g)|) \equiv \frac{1}{2^{n-1}} \text{gap}(g) \quad (\text{S26})$$

This shows that the output amplitudes of Bell sampling from universal quantum circuits can encode the gap of any $\#P$ -function. Finally, we reduce the outcome to the all-zero outcome—and by the hiding property any outcome in the symmetric subspace—by observing that the output string corresponds to the $|\mathbb{1}\rangle^{\otimes n} |Z\rangle^{\otimes 2}$ outcome and hence we can define $\tilde{C}_{f_g} = (\mathbb{1}^{\otimes n} \otimes (Z^{1/2})^{\otimes 2}) C_{f_g}$ to show that $\langle 0^{2(n+2)} | \mathcal{C}_{f_g} \rangle$ is GapP-hard to compute.

By Proposition 8 of Bremner *et al.* [2] (see also Lemma 8 of Ref. [1]), approximating $|\langle 0^{2n} | \tilde{C}_{f_g} \rangle|^2$ up to any relative error $< 1/2$ or additive error $1/2^{2(n+1)}$ is GapP-hard. \square

Notice that this argument also proves that computing certain Pauli coefficients of an n -qubit quantum circuit is GapP hard up to relative error $< 1/2$ since the circuit C_{f_g} we have used in the encoding of the gap of a $\#P$ function is real and $|\langle C | \sigma_r | \bar{C} \rangle|^2 = |\text{tr}[C] \langle C | \sigma_r \rangle|^2$ for a real circuit C . Notice, however, that we cannot reduce to the all-zero string in this case. This is also easy to see since the probability of the all-zero outcome of a real circuit is always one.

Conversely, if we include the $\sqrt{Z} \otimes \sqrt{Z}$ gate at the end of C_{f_g} , this shows that computing the overlap $|\langle C | \bar{C} \rangle|^2$ is GapP-hard in general.

The next step in applying the Stockmeyer argument is to see if approximate average-case hardness is plausible. To this end we can/need to make two arguments.

2. Near-exact average-case hardness

First, we need to show (near-)exact average-case hardness, see [1, Sec. IV.D.5] for the available techniques.

Consider a circuit C with some Haar-random 2-qubit gates. Let us follow the strategy by Krovi [3]. Analogously, we in-

terpolate every 2-qubit gate G_i in the worst-case circuit C to a Haar random gate $H_i G_i$ with H_i drawn from the 2-qubit Haar measure.

$$G_i(\theta) = \exp\left(i\left(1 - \frac{\theta}{2m}\right) \log H_i\right) G_i \quad (\text{S27})$$

$$= V_i^\dagger \sum_{k_i=1}^4 e^{i(1-\frac{\theta}{2m})\phi_{k_i}} |\psi_{k_i}\rangle \langle \psi_{k_i}| V_i G_i \quad (\text{S28})$$

$$=: \sum_{k_j} e^{i(1-\frac{\theta}{2m})\phi_{k_j}} \tilde{G}_{k_j} \quad (\text{S29})$$

where V_i diagonalizes H_i into eigenvectors $|\psi_{k_i}\rangle$ and eigenvalues $\exp(i\phi_{k_i})$. Notice that compared to Krovi [3], we interpolate only by an angle $\theta/2m$ instead of θ/m . Now, we consider the circuit $C(\theta)$ defined by replacing all the gates G_i of C with $G_i(\theta)$. We can write the output probability as

$$\begin{aligned} |\langle 0^{2n} | C(\theta) \rangle|^2 &= |\langle \Phi^+ |^{\otimes n} C(\theta) \otimes C(\theta) | 0^{2n} \rangle| \\ &= \left| \sum_{k_{ij}: i \in [m], j \in \{0,1\}} e^{i(1-\frac{\theta}{2m}) \sum_{ij} \phi_{k_{ij}}} \langle \Phi^+ |^{\otimes n} \tilde{G}_{k_{00}} \dots \tilde{G}_{k_{d1}} | 0^{2n} \rangle \right|^2 \\ &= \sum_{k, k'} e^{i(1-\frac{\theta}{2m}) \Delta \phi_{k, k'}} \langle \Phi^+ | \tilde{G}_k | 0 \rangle \langle 0 | \tilde{G}_{k'}^\dagger | \Phi^+ \rangle, \end{aligned} \quad (\text{S30})$$

where $k = (k_{00}, k_{01}, \dots, k_{m1})$, $\tilde{G}_k = \prod_{ij} G_{k_{ij}}$ and $\Delta \phi_{k, k'} = \sum_{ij} (\phi_{k_{ij}} - \phi_{k'_{ij}})$. Since $|\Delta \phi_{k, k'}|/2m \in O(1)$, we can now follow the argument of Krovi, replacing $m \leftarrow 2m$, to construct a polynomial of degree $d \in O(m/\log(m)^2)$ that approximates the output probabilities of Bell sampling. Using this polynomial, we can run robust polynomial interpolation to obtain near-exact average-case hardness.

3. Anticoncentration

Finally, the question arises whether the output distribution anticoncentrates. Recall that anticoncentration is defined as

$$\Pr_{C \sim \mathcal{C}} \left[P_C(S) \geq \frac{1}{|\Omega|} \right] \geq \gamma, \quad (\text{S31})$$

for constant γ and the sample space Ω . Our standard tool for showing anticoncentration is the Payley-Zygmund inequality which lower bounds the anticoncentrating fraction in terms of second moments as

$$\Pr_{C \sim \mathcal{C}} [Z \geq \alpha \mathbb{E}[Z]] \geq (1 - \alpha)^2 \frac{\mathbb{E}[Z]^2}{\mathbb{E}[Z^2]}, \quad (\text{S32})$$

for a random variable Z with $0 \leq Z \leq 1$ which we take to be $Z = P_C(S)$.

This problem gets interesting for the Bell sampling circuits since already a single copy of the Bell sampling state has two copies of the circuit, so we will need second and fourth moments of the circuit family to show anticoncentration with the standard technique. Recall that the Bell sampling output distribution is given by

$$P_C(r) = |\langle \sigma_r | C \otimes C | 0^{2n} \rangle|^2, \quad (\text{S33})$$

where $\sigma_r = \prod_{i=1}^N \sigma_{x_i}$. Let's assume that we have a four-design and begin with the first moment.

$$\mathbb{E}[P_C(r)] = \langle \sigma_r | \mathbb{E}[C \otimes C | 0^{2n}] \langle 0^{2n} | C^\dagger \otimes C^\dagger | \sigma_r \rangle \quad (\text{S34})$$

$$= \frac{1}{D_{[2]}} \langle \sigma_r | P_{[2]} | \sigma_r \rangle, \quad (\text{S35})$$

where $D_{[t]} = \binom{d+t-1}{t}$ and $P_{[t]}$ is the projector onto the symmetric subspace of t copies (see [4] for a nice intro). We can explicitly compute the expression as

$$\langle \sigma_r | P_{[2]} | \sigma_r \rangle = \frac{1}{2} \langle \sigma_r | (\mathbb{1} + \mathbb{S}) | \sigma_r \rangle \quad (\text{S36})$$

$$= \langle \Phi^+ | (\sigma_r^2 \otimes \mathbb{1}) | \Phi^+ \rangle + \langle \Phi^+ | (\sigma_r \sigma_r^T \otimes \mathbb{1}) | \Phi^+ \rangle \quad (\text{S37})$$

$$= 2^n (1 + (-1)^{\pi_Y(r)}), \quad (\text{S38})$$

where we in abuse of notation we write $|\Phi^+\rangle = |\Phi^+\rangle^{\otimes n}$ and observe that $\text{tr}[XX^T] = \text{tr}[ZZ^T] = \text{tr}[\mathbb{1}] = 2^n$ and $\text{tr}[YY^T] = -2^n$. We find—as expected—that the Bell state with an odd number of singlet states (corresponding to $\sigma_r = Y$) is in the antisymmetric subspace and therefore the projector onto the symmetric subspace evaluates to zero in that case:

$$\mathbb{E}[P_C(r)] = \begin{cases} D_{[2]}^{-1} = 2/(2^{2n}(1+2^{-n})) & \text{if } \pi_Y(r) \text{ even} \\ 0 & \text{if } \pi_Y(r) \text{ odd} \end{cases} \quad (\text{S39})$$

The output distribution is thus supported on the even Y -parity sector on which all outcomes have equal expectation value given by $D_{[2]}^{-1} = 2/(2^n(2^n+1))$. The size of the even- Y -parity sector should be exactly given by this number since these strings correspond to a basis of the symmetric subspace.

The second moment is more complicated. It reads

$$\begin{aligned} \mathbb{E}[\langle \sigma_r |^{\otimes 2} C^{\otimes 4} | 0 \rangle \langle 0 | (C^\dagger)^{\otimes 4} | \sigma_r \rangle^{\otimes 2}] \\ = \frac{1}{D_{[4]}} \langle \sigma_r |^{\otimes 2} P_{[4]} | \sigma_r \rangle^{\otimes 2} \end{aligned} \quad (\text{S40})$$

To compute this overlap on the symmetric subspace on which $\sigma_r = \sigma_r^T$, we write $P_{[f]} = 4!^{-1} \sum_{\sigma \in S_4} P_\sigma$, where P_σ is the permutation matrix corresponding to the element σ of the symmetric group S_4 . We observe that for a $c = 1/3$ -fraction of the permutations in the definition of $P_{[4]}$ the overlap evaluates to $\text{tr}[\sigma_r^2]/2^{2n} = \text{tr}[\mathbb{1}]/2^{2n} = 1$, while for the other $1 - c$ fraction, it evaluates to $\text{tr}[\sigma_r^4]/2^{2n} = \text{tr}[\mathbb{1}]/2^{2n} = 1/2^n$. Viewing the Bell state $|\Phi^+\rangle$ as a vectorization of the identity matrix, these correspond exactly to the cases in which there are one versus two connected components in the resulting graph.

Consequently, we find

$$\begin{aligned} \mathbb{E}[\langle \sigma_r |^{\otimes 2} C^{\otimes 4} | 0 \rangle \langle 0 | (C^\dagger)^{\otimes 4} | \sigma_r \rangle^{\otimes 2}] \\ = \frac{1}{D_{[4]}} \left(c + \frac{1-c}{2^n} \right). \end{aligned} \quad (\text{S41})$$

for all even Y -parity r (satisfying that $\pi_Y(r)$ is even).

From this we find that if \mathcal{C} generates a state 4-design, for all r in the symmetric subspace

$$\Pr \left[P_C(r) \geq \frac{\alpha}{D_{[2]}} \right] \geq (1 - \alpha)^2 \frac{1}{c + \frac{1-c}{2^n}} \frac{D_{[4]}}{D_{[2]}^2} \quad (\text{S42})$$

$$= \frac{(1 - \alpha)^2}{3! \cdot (c + \frac{1-c}{2^n})} \frac{(2^n + 3)(2^n + 2)}{(2^n + 1)2^n} \quad (\text{S43})$$

$$\geq \frac{(1 - \alpha)^2}{6c} \quad (\text{S44})$$

By the result of Brandão *et al.* [5], this means that the Bell sampling distribution anticoncentrates in linear depth. Given the fact that we just copied two random circuits and qualitatively found the same results as for a single copy, we conjecture that the Bell sampling distribution also anticoncentrates in log-depth. To check this, we would have to directly compute the moments of the output distribution, but this becomes more complicated than for the single-copy case. This is because the local degrees of freedom the standard mapping to a statistical-mechanical model [6–9], increase from 2 to 24 permutations.

S3. LEARNING CIRCUIT PROPERTIES FROM BELL SAMPLING

In the following, we will detail the tests that can be performed *using just the samples* from the Bell distribution. Success on those tests—while falling short of loophole-free verification—increases our confidence in the correctness of the experiment.

A. Measuring purity and fidelity

Our first observation is that, given some noisy state preparation $\rho \otimes \sigma$, we can estimate $\text{tr}[\rho\sigma]$ from Bell measurements on $\rho \otimes \sigma$. To see this, observe that we can express the single-qubit swap operator in the Bell basis as

$$\mathbb{S} = |00\rangle\langle 00| + |11\rangle\langle 11| + |01\rangle\langle 10| + |10\rangle\langle 01| \quad (\text{S45})$$

$$= \underbrace{|\sigma_{00}\rangle\langle \sigma_{00}| + |\sigma_{01}\rangle\langle \sigma_{01}| + |\sigma_{10}\rangle\langle \sigma_{10}|}_{P_{\vee 2}} - \underbrace{|\sigma_{11}\rangle\langle \sigma_{11}|}_{P_{\wedge 2}} \quad (\text{S46})$$

Hence, we can estimate the overlap

$$O(\rho, \sigma) = \text{tr}[\rho\sigma] = \text{tr}[(\rho \otimes \sigma)\mathbb{S}] \quad (\text{S47})$$

by taking the difference between the frequency of outcomes with even parity of 11-outcomes and odd parity of 11-outcomes

$$\hat{O} = \frac{1}{M} (|\{r : \pi_Y(r) = 0\}| - |\{r : \pi_Y(r) = 1\}|), \quad (\text{S48})$$

where M is the total number of measurements.

Let us now see what the quantity $O(\rho, \sigma)$ corresponds to for various assumptions on what ρ and σ are.

1. Purity and Fidelity

The weakest assumption is that $\rho = \sigma$. In this case

$$O(\rho, \rho) = \text{tr}[\rho^2] = P(\rho) \quad (\text{S49})$$

is just the purity of ρ .

Using randomized compiling implemented independently on each copy of the Bell sampling circuit, the effective noise channel for the Bell samples can be approximately reduced to a Pauli channel [10]. Strictly speaking, randomized compiling works well in the experimentally relevant setting in which there are some ‘easy’ gates, say, Pauli gates, on which the noise is approximately gate-independent and some ‘hard’ gates which have gate-dependent errors. In that setting, the gate-dependent noise on the hard gates, except for those in the last layer of the circuit, can be reduced to a Pauli channel. In this case, the effective density matrix on each copy can approximately be written as

$$\rho_C(\eta) := (1 - \eta) |C\rangle\langle C| + \eta \sigma, \quad (\text{S50})$$

for another density matrix σ . As we showed in the main text, when $\eta \ll 1$, the purity can then be used to estimate the fidelity

$$F(\rho_C(\eta), |C\rangle\langle C|) = 1 - \eta(1 - \langle C|\sigma|C\rangle). \quad (\text{S51})$$

Generically, we expect that $\langle C|\sigma|C\rangle$ is exponentially small since after randomized compiling, the deviation σ from the ideal state $|C\rangle\langle C|$ is uncorrelated with $|C\rangle\langle C|$.

Indeed, a particularly simple realization of this noise model occurs in the so-called white-noise approximation where $\sigma = \mathbb{1}/2^n$. In that case, the correspondence between fidelity and purity holds for any noise rate. From $P(\rho_C(\eta))$ we can easily obtain an estimate of η since

$$\text{tr}[\rho_C(\eta)^2] = (1 - \eta)^2 + \eta(1 - \eta)/2^n + \eta^2/2^{2n}, \quad (\text{S52})$$

which lets us directly estimate the fidelity

$$F(\rho_C(\eta), |C\rangle\langle C|) = 1 - \eta(1 - 1/2^n). \quad (\text{S53})$$

2. Average fidelity

Finally, following an argument in Ref. [11], consider circuits C constructed from two-qubit gates which form a unitary 2-design, and consider the noise setting in which a single-qubit Pauli noise channel $\mathcal{N} = \sum_{i \in \{0, \dots, 3\}} p_i \sigma_i \cdot \sigma_i$ with acts on every qubit after the application of a two-qubit gate. Let $\mathcal{U} = |U\rangle\langle U| \equiv \text{vec}(U \cdot U^\dagger)$ be the matricization of the adjoint action of the two-qubit gate U , where we denote the vectorization of a matrix A by $|A\rangle$. The full circuit is then composed of two copies of noisy unitary two-qubit channels given by $\mathcal{N}^{\otimes 4} \mathcal{U}^{\otimes 2}$, where the cut is across the bipartition of the Bell measurement, and let the noisy state be $\rho(\epsilon)$.

Then, we can express the fidelity between a state with noise rate ϵ and a state with noise rate 0, as the purity of a state with noise rate ϵ' . To see this, we evaluate the average over gates in a single layer of the circuit. Then

$$\mathbb{E}_U \mathcal{N}^{\otimes 4} \mathcal{U}^{\otimes 2} = \sum_{\pi, \sigma} \mathcal{N}^{\otimes 4} |\pi\rangle\langle \pi|, \quad (\text{S54})$$

and we can evaluate

$$\mathcal{N}^{\otimes 2} |\mathbb{1}\rangle = \mathcal{N}' \otimes \mathbb{1} |\mathbb{1}\rangle = |\mathbb{1}\rangle, \quad (\text{S55})$$

$$\mathcal{N}^{\otimes 2} (\mathbb{S}) = \sum_{ijk \in \{0, \dots, 3\}} p_i p_j (-1)^{s(k)} \sigma_i \otimes \sigma_j |\sigma_k\rangle\langle \sigma_k| \sigma_i \otimes \sigma_j \quad (\text{S56})$$

$$= \sum_{ij \in \{0, \dots, 3\}} p_i p_j (\sigma_i \sigma_j \otimes \mathbb{1}) \mathbb{S} (\sigma_j \sigma_i \otimes \mathbb{1}) \quad (\text{S57})$$

$$= \sum_k q_k (\sigma_k \otimes \mathbb{1}) \mathbb{S} (\sigma_k \otimes \mathbb{1}) \quad (\text{S58})$$

$$\equiv (\mathcal{N}' \otimes \mathbb{1}) (\mathbb{S}). \quad (\text{S59})$$

Here, we have defined $s(k) = 1$ for $k = 4$ (corresponding to $\sigma_k = Y$) and 0 otherwise. We also find the error probabilities of the new noise channel \mathcal{N}' as $q_0 = p_0^2 + \sum_{i \neq 0} p_i^2$ and for $k \neq 0$, $q_k = 2p_k p_0 - \sum_{ij \neq 0} \epsilon_{ijk} p_i p_j$, where ϵ_{ijk} denotes the Levi-Civita symbol. Thus, we have moved noise from one copy of the circuit to the other copy of the circuit, and thereby related the quantities for fidelity (one noisy, one ideal copy) and purity (two noisy copies).

For the depolarizing channel with depolarizing parameter ϵ , we have $p_0 = 1 - 3\epsilon/4$, $p_1 = p_2 = p_3 = \epsilon/4$, so $q_0 = (1 - 3\epsilon/4)^2 + 3(\epsilon/4)^2$, $q_1 = q_2 = q_3 = \epsilon/2 - 2(\epsilon/4)^2$. Hence, in this case the type of noise channel even remains the same, and

purity with local depolarizing strength ϵ simply corresponds to fidelity with a different local depolarizing rate given by $2\epsilon - \epsilon^2$:

$$\mathbb{E}_C P(\rho(\epsilon)) = \mathbb{E}_C F(\rho(2\epsilon - \epsilon^2)), |C\rangle\langle C| \quad (\text{S60})$$

For more general types of Pauli noise, the noise channels \mathcal{N} and \mathcal{N}' will be different; however, the mapping from purity to fidelity carries over through appropriate redefinition of the noise rate.

Via Fuchs-van-der-Graaf this gives a bound on the TVD distance of the sampled distribution Q from the Bell distribution P_C as

$$\text{TVD}(P_C, Q) \leq \sqrt{1 - (1 - \eta(1 - 2^{-n}))^2} \quad (\text{S61})$$

$$\stackrel{\eta \ll 1}{\approx} \sqrt{2\eta(1 - 2^{-n})}. \quad (\text{S62})$$

3. Relation to verified quantum computation

The results in this section imply that arbitrary quantum computations in the Bell sampling model (with randomized compiling or in an encoded setting) can be verified efficiently just from the classical Bell samples under specific assumptions about the noise in the physical device. Let us briefly pause and contextualize this point in the context of protocols for the verification of quantum computations in different models of computation. Efficiently verifying a quantum computation in full generality is known to be possible only using comparably complicated schemes that require large overheads compared to a bare quantum circuit.

The specific schemes we have in mind are blind verified quantum computation [12], verification using two spatially separated entangled quantum computers [13], and classically verifiable quantum computation using post-quantum cryptography [14]. All of these settings make strong assumptions about the capabilities of a quantum device. Blind verified computation requires a large space overhead compared to the circuit model since it makes use of measurement-based quantum computation and requires the ability to prepare single-qubits perfectly. The classical leash protocol of Reichardt *et al.* [13] requires two spatially separated quantum computers which share a large number of Bell pairs, since it relies on the rigidity of CHSH inequalities. It also involves a space-time mapping similar to measurement-based quantum computation and therefore incurs a high space-overhead. Finally, the protocol of Mahadev [14] requires a high overhead since the quantum computation needs to be performed in a post-quantum secure homomorphic encryption scheme and is based on the assumption that such schemes exist.

Bell sampling is a model of quantum computation that allows one to efficiently glean information about a quantum state prepared by an arbitrary circuit from measurements in a single basis only. It is a setting that is realistic in the near and intermediate term, namely one where two copies of a circuit are entangled on a single device. The price to be paid for the

small overhead (just a factor of 2) is that the obtained information is conditional on assumptions about the noise occurring in the physical device. For instance, when we use randomized compiling, that assumption boils down to an assumption about single-qubit Pauli gates and the last circuit layer being nearly noise-free, while the remainder of the noise on the other gates is not context-dependent. We think of this setting as comparable to blind-verified computation, which can be viewed as imposing an assumption on the noise in the single-qubit state preparation step. For verification of random circuits on average (through the average state fidelity) the assumption we have used (but which could be weakened) is local gate-independent noise.

While we have made an effort to obtain a good initial understanding of the types and amounts of noise required for the verification protocol to work, an exhaustive characterization would go far beyond the scope of this letter. We hope that our discussion of Bell sampling as a validated model of quantum computation will motivate future research into specific instantiations.

B. Testing depth

Another property that we can efficiently check from the Bell samples is subsystem purity and thereby the entanglement structure of the state. To this end, given a subsystem $A \subset [n]$ we simply consider the substrings $r_A = (r_i : i \in A) \cup (r_{n+i} : i \in A)$ corresponding to this subsystem and then run the purity test.

Now, we can use the maximal entanglement achievable by circuits in a certain architecture, in order to test for their depth. For geometrically local circuits in large local dimension with entangling two-qubit gates, the maximal entanglement which can be generated in any subsystem is simple: it is just given by the number of entangled pairs which can be generated by the circuit across the boundary of the considered region. Thus, for depth- d circuits in a one-dimensional geometry (with closed boundary conditions), the maximal entanglement entropy of a contiguous subsystem A is given by $E_A(d) = \min\{2d, |A|, n - |A|, n/2\}$. In higher dimensions, for a subsystem A this increases to $E_A(d) = \min\{|\partial A|d, |A|, n - |A|, n/2\}$, where ∂A is the number of edges in the interaction graph protruding out of the subsystem A . The maximal entanglement $E_A(d)$ is thus a depth-dependent property.

This picture provides us with a simple way to test the depth of a circuit from the Bell samples. We estimate $P_A(\rho) = \text{tr}[\rho_A^2]$ for random contiguous subsystems of size $|A| = n/2$ and compare the results to the maximal achievable entanglement in each subsystem in a given architecture with circuit depth d . The minimal d for which $P_A(\rho) \geq E_A(d)$ is our lower bound for the circuit depth.

The estimate of the subsystem purity is obtained from the substrings r_A as

$$\hat{P}_A = \frac{1}{M} (|\{r : \pi_Y(r_A) = 0\}| - |\{r : \pi_Y(r_A) = 1\}|), \quad (\text{S63})$$

and $|\hat{P}_A - \text{tr}[\rho_A^2]| \leq \sqrt{2 \log(\delta/2)/M}$ with probability at least

$1 - \delta$. Hence, in order to obtain an estimate of the Rényi-2 entanglement entropy for a circuit of depth $d \leq n/2$ —given by $2^{-O(d)}$ — $2^{O(d)}$ samples are required. This means that log-depth circuits can be efficiently validated, while larger depth requires superpolynomially many Bell samples.

Algorithm 1 Depth test from maximal entanglement

Input: Bell samples $b^0, \dots, b^M \leftarrow P_\psi$, error tolerance $\epsilon > 0$.

- 1: Estimate $\hat{P}_A(|\psi\rangle\langle\psi|)$ as defined in Eq. (S63) for a contiguous subsystem A with $|A| = n/2$ obtaining an estimate $\bar{P}_{n/2}$, for example using a median-of-means estimator.

Output: $d_l = \max \{d : -\log(\bar{P}_{n/2} + \epsilon) \geq E_{n/2}(d)\}$

We can further refine the depth test by using properties of *random* quantum circuits. Consider first a Haar-random pure state. It has been known since the seminal work of Page [15] that the average entanglement of subsystems of increasing size obeys the now so-called *Page curve*: as the subsystem size increases to $n/2$, its average entanglement entropy increases, as it increases beyond $n/2$, it decreases back to 0 at subsystem size n . The precise shape of the Page curve including the so-called Page correction—the deviation of the entanglement at subsystem size $n/2$ from maximal entanglement—is known for various families of random quantum states [16–19]. Notably, Garnerone *et al.* [16] show that random MPSs exhibit typical entanglement whenever the bond dimension scales at least as a square of the system size.

However, depth-dependent Page curves have to the best of our knowledge not been computed analytically yet. Nonetheless, empirically one finds that for various circuit families, the bounds $E_A(d)$ can be nearly exhausted, see for example Ref. [20].

Rather than comparing to the maximal achievable entanglement $E_A(d)$, if we have a depth-dependent Page curve, we can compare the result to the anticipated value of the Page curve $T_A(d) = -\log(\mathbb{E}_{C \sim \mathcal{C}_d} \text{tr}[(\text{tr}_{A^c} |C\rangle\langle C|)^2])$ for a size- $n/2$ subsystem A for the circuit family \mathcal{C}_d of a given depth d .

Algorithm 2 Depth test from average entanglement

Input: Bell samples $b_0^{C_1}, \dots, b_M^{C_L} \leftarrow P_C$ for $C_1, \dots, C_L \leftarrow \mathcal{C}$, error tolerance $\epsilon > 0$.

- 1: Estimate $\frac{1}{K} \sum_{i=1}^L \hat{P}_A(|\psi\rangle\langle\psi|)$ with \hat{P}_A defined as in Eq. (S63), from choosing a contiguous subsystem with $|A| = n/2$ obtaining an estimate $\bar{P}_{n/2}$, for example using a median-of-means estimator.

Output: $d_l = \max \{d : -\log(\bar{P}_{n/2} + \epsilon) \geq t_A(d)\}$

Small amounts of noise The presence of a small amount of noise only slightly distorts our estimate of the purity and subsystem purity. We can always write the noisy state as

$$\rho = (1 - \eta) |\psi\rangle\langle\psi| + \eta \chi, \quad (\text{S64})$$

where $\chi \perp |\psi\rangle\langle\psi|$. The deviation of its purity from unity is then given by

$$|1 - \text{tr} \rho^2| = 2\eta + O(\eta^2) \quad (\text{S65})$$

and likewise for the subsystem purity

$$|\text{tr} |\psi\rangle\langle\psi|_A^2 - \text{tr} \rho_A^2| \quad (\text{S66})$$

$$\leq 2\eta |\text{tr} |\psi\rangle\langle\psi|_A^2 - \text{tr} |\psi\rangle\langle\psi|_A \chi_A| + O(\eta^2) \quad (\text{S67})$$

$$\leq 2\eta (\text{tr} |\psi\rangle\langle\psi|_A^2 + 1) + O(\eta^2) \quad (\text{S68})$$

$$\leq \eta(3 + 2^{-O(d)}) + O(\eta^2). \quad (\text{S69})$$

Hence to verify depth we need states with exponentially small impurity $2^{-O(d)}$ in the circuit depth d .

Larger amounts of noise The discrepancy of the Page curve for a pure state and the Page curve of a mixed state can be intuitively understood: For a maximally mixed state the Rényi-2 entropy of subsystems of size k is exactly given by $-\log 2^{-k}$, and hence, as a function of subsystem size, it is just given by the subsystem size. The entanglement entropy of a white-noise state $\rho = (1 - \eta) |\psi\rangle\langle\psi| + \eta \mathbb{1}/2^n$ is thus given by

$$\begin{aligned} -\log \text{tr} \rho_A^2 &= \\ &= -\log [(1 - \eta)^2 \text{tr}(|\psi\rangle\langle\psi|_A^2) + (2\eta(1 - \eta) + \eta^2) 2^{-k}] \\ &= -\log [(1 - \eta)^2 \text{tr}(|\psi\rangle\langle\psi|_A^2)] \\ &\quad + \frac{(2\eta(1 - \eta) + \eta^2) 2^{-k}}{(1 - \eta)^2 \text{tr}(|\psi\rangle\langle\psi|_A^2)} + O(2^{-2k}), \end{aligned} \quad (\text{S70})$$

which yields a good approximation for $\text{tr}(|\psi\rangle\langle\psi|_A^2) \gg 2^{-k}$.

C. Learning a Clifford + T circuit

In this section, we show that a quantum state prepared by a circuit with few non-Clifford gates can be learned efficiently. We separate the proof into several steps. First, we derive the expansion of the operator $|\psi\rangle\langle\bar{\psi}|$ corresponding to a quantum state $|\psi\rangle$ generated by a Clifford circuit with few T gates. This operator determines the Bell sampling distribution, and in deriving it, we will already see the central concepts we will use in the learning algorithm. Motivated by this decomposition, we elaborate the algorithm. Finally, we analyze statistical errors when running the algorithm and show the output of the algorithm is close to $|\psi\rangle$ in fidelity.

Before we describe the protocol, let us recap some simple properties of the Bell samples from the stabilizer state $|S\rangle\langle S| = 2^{-n} \sum_{\sigma \in \mathcal{S}} \sigma$ with n -dimensional stabilizer group $\mathcal{S} \subset \text{P}_n$, i.e., a commuting subgroup of the n -qubit Pauli group P_n . For stabilizer states $|S\rangle$, the complex conjugation $|\bar{S}\rangle = \sigma_k |S\rangle$ is described by a Pauli operator σ_k that depends on $|S\rangle$ [21]. Let us denote by roman letters the binary symplectic subspace $S \subset \mathbb{F}_2^{2n}$ corresponding to a subgroup \mathcal{S} of P_n , which includes all but the phase information about S . For $a, b \in \mathbb{F}_2^{2n}$, the symplectic inner product is given by $\omega(a, b) := (\sum_{i=1}^n a_i b_{n+i} - b_i a_{n+i}) \bmod 2$.

From Eq. (3) it immediately follows that the output distribution of Bell sampling from $|S\rangle \otimes |S\rangle$ is supported on the affine space $S \oplus k := \{s \oplus k : s \in S\}$. We can therefore learn

S from differences of the Bell samples $b^i \oplus b^j \in S$, and the missing phases of the stabilizers from a measurement of the corresponding stabilizer operators [21].

1. The Bell distribution of low T -count quantum states

Consider a state $|\psi\rangle = C_t T_{x_t} C_{t-1} T_{x_{t-1}} \cdots T_{x_1} C_0 |0\rangle$ generated by a circuit comprising $t+1$ Clifford layers C_i and t T gates at positions $x_i \in [n]$ for $i \in [t]$. Then, we can shift the T -gates to the end of the circuit as

$$|\psi\rangle = \tilde{T}_t \cdots \tilde{T}_1 C_t \cdots C_0 |0\rangle =: \tilde{T}_t \cdots \tilde{T}_1 |S\rangle, \quad (\text{S71})$$

where we have defined $\tilde{T}_i = C_t \cdots C_i T_{x_i} C_i^\dagger \cdots C_t^\dagger$. Hence, we can write

$$|\psi\rangle \langle \bar{\psi}| = \prod_{i=1}^t (\alpha \mathbb{1} + i\beta P_i) |S\rangle \langle S| \prod_{j=1}^t (\alpha \mathbb{1} - i\beta P_j) K^\dagger, \quad (\text{S72})$$

where $P_i = C_t \cdots C_i Z_{x_i} C_i^\dagger \cdots C_t^\dagger$, $\alpha = \cos(\pi/8)$ and $\beta = \sin(\pi/8)$, and $K : |\psi\rangle \mapsto |\bar{\psi}\rangle$ is the complex-conjugation operator. In the next step, we rewrite $|\psi\rangle \langle \bar{\psi}|$ in terms of the stabilizer groups \mathcal{S} which stabilizes $|S\rangle$ and $\mathcal{G} = \langle P_1, \dots, P_d \rangle$ as

$$|\psi\rangle \langle \bar{\psi}| = \frac{1}{2^n} \left(\sum_{\sigma_g \in \mathcal{G}} \alpha_g \sigma_g \right) \left(\sum_{\sigma_s \in \mathcal{S}} \sigma_s \right) \left(\sum_{\sigma_{g'} \in \mathcal{G}} \bar{\alpha}_{g'} \sigma_{g'} \right) K^\dagger, \quad (\text{S73})$$

where the prefactors are given by

$$\alpha_g = \sum_{x \in \{0,1\}^d : \prod_i P_i^{x_i} = \text{sign}(x)} \alpha^{d-|x|} (i\beta)^{|x|} \text{sign}(x). \quad (\text{S74})$$

Now, we make use of the fact that the complex conjugation operator K acting on the stabilizer state $|S\rangle$ can be written as a Pauli- Z matrix $\sigma_k := \sigma_{10}^{\otimes s}$ for some $s \in \{0,1\}^n$ depending on $|S\rangle$ as $K|S\rangle = \sigma_k|S\rangle$. This gives us

$$\begin{aligned} |\psi\rangle \langle \bar{\psi}| &= \frac{1}{2^n} \left(\sum_{\sigma_g \in \mathcal{G}} \alpha_g \sigma_g \right) \left(\sum_{\sigma_s \in \mathcal{S}} \sigma_s \right) \sigma_k \left(\sum_{\sigma_{g'} \in \mathcal{G}} \beta_{g'} \sigma_{g'} \right) \\ &= \sum_{\sigma_g, \sigma_{g'} \in \mathcal{G}} \alpha_g \beta_{g'} \sigma_g \Pi_{\mathcal{S}} \sigma_k \sigma_{g'} \\ &= \sum_{\sigma_g, \sigma_{g'} \in \mathcal{G}} \alpha_g \beta_{g'} (-1)^{\langle g', k \rangle} \sigma_g \Pi_{\mathcal{S}} \sigma_{g'} \sigma_k, \end{aligned} \quad (\text{S75})$$

where $\Pi_{\mathcal{S}} = \sum_{\sigma_s \in \mathcal{S}} \sigma_s / 2^n \equiv |S\rangle \langle S|$ denotes the projector onto the ground space of \mathcal{S} . Moreover, we let $\beta_g = \bar{\alpha}_g \cdot (-1)^{\pi_Y(g)}$.

Let us denote by $\langle \mathcal{G}, \mathcal{S} \rangle = \{\sigma_g \sigma_s : \sigma_g \in \mathcal{G}, \sigma_s \in \mathcal{S}\}$ the group generated by \mathcal{G} and \mathcal{S} , and by $\mathcal{G} \oplus \sigma = \{\sigma_g \sigma : \sigma_g \in \mathcal{G}\}$ the shift of \mathcal{G} by σ . We notice that all Pauli operators appearing in the sum are elements of $\langle \mathcal{G}, \mathcal{S} \rangle \oplus \sigma_k$ with dimension $\dim(\langle \mathcal{G}, \mathcal{S} \rangle) \leq n+t$. Let us decompose $\langle \mathcal{G}, \mathcal{S} \rangle = \langle \mathcal{C}, \mathcal{L} \rangle$ into a maximally commuting subgroup \mathcal{C} , i.e., the maximal subgroup with the property that $[\sigma_s, \sigma_l] = 0$, $\forall \sigma_s \in \mathcal{C}, \sigma_l \in \langle \mathcal{G}, \mathcal{S} \rangle$ (the stabilizer group), and a ‘logical’ subgroup \mathcal{L} (which is

ambiguous because it can be shifted by any element of \mathcal{C} . On the level of the corresponding symplectic vector spaces, $C = H \cap H^\perp$, where $H^\perp = \{s \in \mathbb{F}_2^{2n} : \omega(s, h) = 0 \forall h \in H\}$, and $H := \text{span}(G, S)$. To find C one thus simply needs to solve a linear system of equations. Thus, $\dim(C) \geq n-t$ while $\dim(\mathcal{L}) \leq 2t$. Then we can decompose every element σ of $\langle \mathcal{G}, \mathcal{S} \rangle$ as $\sigma_g = \sigma_l(g) \sigma_c(g) = \sigma_c(g) \sigma_l(g)$ for $\sigma_l(g) \in \mathcal{L}$ and $\sigma_c(g) \in \mathcal{C}$.

We can then rewrite Eq. (S75) as

$$|\psi\rangle \langle \bar{\psi}| = \sum_{\sigma_g, \sigma_{g'} \in \mathcal{G}} \alpha_g \beta_{g'} \sigma_l(g) \sigma_c(g) \Pi_{\mathcal{S} \setminus \mathcal{C}} \Pi_{\mathcal{C}} \sigma_c(g') \sigma_l(g') \sigma_k \quad (\text{S76})$$

$$= \sum_{\sigma_g, \sigma_{g'} \in \mathcal{G}} \alpha_g \beta_{g'} \sigma_l(g) \Pi_{\mathcal{S} \setminus \mathcal{C}} \underbrace{\sigma_c(g) \Pi_{\mathcal{C}} \sigma_c(g')}_{=\Pi_{\mathcal{C}}} \sigma_l(g') \sigma_k \quad (\text{S77})$$

$$= \sum_{\sigma_g, \sigma_{g'} \in \mathcal{G}} \alpha_g \beta_{g'} \sigma_l(g) \Pi_{\mathcal{S} \setminus \mathcal{C}} \underbrace{\sigma_c(g) \Pi_{\mathcal{C}} \sigma_c(g')}_{=\Pi_{\mathcal{C}}} \sigma_l(g') \sigma_k \quad (\text{S78})$$

$$= \sum_{\sigma_g, \sigma_{g'} \in \mathcal{G}} \sum_{\sigma \in \mathcal{S} \setminus \mathcal{S}'} \frac{\alpha_g \beta_{g'}}{2^{\dim(\mathcal{L})}} \underbrace{\sigma_l(g) \sigma_l(g')}_{\in \mathcal{L}} \Pi_{\mathcal{C}} \sigma_k \quad (\text{S79})$$

$$=: \sum_{l \in \mathcal{L}} \lambda_l \sigma_l \Pi_{\mathcal{C}} \sigma_k =: \sum_{h \in \text{span}(L, C) \oplus k} q(h) \sigma_h. \quad (\text{S80})$$

In the first equality we have used that $\Pi_{\mathcal{S}} = \Pi_{\mathcal{C}} \Pi_{\mathcal{S} \setminus \mathcal{C}}$, and in the last equality we have grouped the operators in \mathcal{L} and defined its coefficients as $\lambda_l \in \mathbb{C}$, $l = 1, \dots, 4^t$. Fixing an L such that $\text{span}(L, C) = H$, and given $l \in L$ we fix the phase of the corresponding $\sigma_l \in \mathcal{L}$ in the logical (non-commuting) subgroup \mathcal{L} to be $+1$, making the coefficients λ_l unique. Since L decomposes C into $2^{\dim(L)}$ disjoint cosets, for $h = l \oplus s \oplus k$, we also define $q(h) = \lambda_l / 2^{\dim(C)}$.

When we perform Bell sampling from $|\psi\rangle \otimes |\psi\rangle$, a Bell sample b will then be distributed as

$$P_\psi(b) = |\text{tr}[|\psi\rangle \langle \bar{\psi}| \sigma_b]|^2 / 2^n = 2^n |q(b)|^2. \quad (\text{S81})$$

In particular the distribution is supported on $\text{span}(C, L) \oplus k$.

A subspace $S \subset \{0,1\}^n$ is *isotropic* if for all $s, t \in S$, $\omega(s, t) = 0$. An isotropic subspace S thus corresponds to a commuting subgroup \mathcal{S} , since $\omega(s, t) = 0$ iff $[\sigma_s, \sigma_t] = 0$. The maximal isotropic subspace C of H is called the *radical* of H , which is given by $\text{rad } H = C \cap C^\perp$.

2. The learning algorithm

Our learning algorithm is based on the decomposition (S80) and the resulting Bell distribution (S81). In the algorithm, we assume access to state preparations of $|\psi\rangle$

Algorithm 3 Magic estimation

Input: $M \in \mathbb{N}$

- 1: Perform Bell sampling from $|\psi\rangle$, obtaining samples $b^0, \dots, b^M \leftarrow P_\psi$.
 - 2: Compute all Bell differences $b^{(i,j)} = b^j \oplus b^i$.
 - 3: Define $G' = \text{span}(\{b^{(i,j)}\}_{i,j})$, and $\hat{t} = \dim(G') - n$.
-

Algorithm 4 Clifford+ T learning algorithm

Input: Error threshold ϵ , failure probability δ .

- 1: Run Algorithm 3 with $M = 2n \log(1/\delta)/\epsilon$, yielding \hat{t}, G' .
- 2: Find a set of generators $s^1, \dots, s^{\hat{t}}$ of the radical $C' = \text{rad } G'$ of G' by solving a linear system of equations.
- 3: Find a Clifford unitary $U_{C'}$ such that $U_{C'} \sigma_{s^i} U_{C'}^\dagger = \pm Z_i$ for all $i = 1, \dots, \hat{t}$.
- 4: Let $M' = O(2^k \log(1/\delta)/\epsilon^2)$. Prepare M' copies of $U_{C'} |\psi\rangle$. Measure qubits $1, \dots, n - \hat{t}$ in the computational basis, and qubits $n - \hat{t} + 1, \dots, n$ according to the pure-state tomography scheme of Ref. [22], yielding data sets $X = \{x^1, \dots, x^{M'}\}$, $D = \{d^1, \dots, d^{M'}\}$. Let $x \in \{0, 1\}^{n-\hat{t}}$ be the majority outcome of the computational-basis measurements.
- 5: Remove those elements d^k for which $x^k \neq x$ from D .
- 6: Run the recovery algorithm of Ref. [22] on D .

Output: $U_{C'}, x, |\hat{\varphi}\rangle$.

3. Correctness of Algorithm 4

To show the correctness of the algorithm, we proceed in several steps. In the first step, we show that the state $U_{C'} |\psi\rangle$ is close to a product state $|x\rangle \otimes |\varphi\rangle$, since all elements of C' are close to stabilizers of $|\psi\rangle$. In the second step, we show that the tomographic estimates we obtain from measuring the first $n - \hat{t}$ qubits in the computational basis, and performing state tomography on others, yield an overall reconstruction $|\hat{\psi}\rangle$ with fidelity at least $1 - \epsilon$ with the target state $|\psi\rangle$.

Lemma 3. *Let C' be the radical of the subspace spanned by $M \in O(n \log(1/\delta)/\epsilon)$ Bell samples from a pure state $|\psi\rangle$, and let $U_{C'}$ be the associated Clifford unitary. Then with probability at least $1 - \delta$ over the Bell samples, there is a bit string $x \in \{0, 1\}^{\dim(C')}$ and a pure state $|\varphi\rangle \in (\mathbb{C}^2)^{\otimes n - \dim(C')}$ such that $|\langle x | \langle \varphi | \psi \rangle|^2 \geq 1 - \epsilon$.*

Proof. We begin by observing that all elements of C' are approximate stabilizers of $|\psi\rangle$. To see this, we observe that for any Pauli operator P we can measure $\langle \psi | P | \psi \rangle^2$ using the Bell samples $B = \{b^1, \dots, b^M\}$ as [23]

$$\langle \psi | P | \psi \rangle^2 \approx \hat{e}(P) := \frac{1}{M} \left(|\{b \in B : P \otimes P | \sigma_b\rangle = +1 | \sigma_b\rangle\}| - |\{b \in B : P \otimes P | \sigma_b\rangle = -1 | \sigma_b\rangle\}| \right). \quad (\text{S82})$$

Let us write the Pauli operators corresponding to the Bell samples $\sigma_{b^i} = \sigma_{g^i} \sigma_k$ for $g^i \in G'$ and σ_k the Pauli operator corresponding to complex conjugation. Then, for all $c \in C'$ and $i \in [M]$, $[\sigma_c, \sigma_{g^i}] = 0$, since C' is the radical of G' . Hence,

$$\sigma_c \otimes \sigma_c | \sigma_{b^i} \rangle = \sigma_c \sigma_{b^i} \otimes \sigma_c | \Phi^+ \rangle \quad (\text{S83})$$

$$= \pm (\sigma_c \sigma_{g^i} \sigma_c \sigma_k \otimes \mathbb{1}) | \Phi^+ \rangle \quad (\text{S84})$$

$$= \pm (\sigma_{g^i} \sigma_k \otimes \mathbb{1}) | \Phi^+ \rangle = \pm | \sigma_{b^i} \rangle, \quad (\text{S85})$$

where the sign depends on whether $\sigma_c^T = \pm \sigma_c$ and $\sigma_k \sigma_c = \pm \sigma_c \sigma_k$, but not on the Bell sample b^i . Hence, all estimated expectation values $e(\sigma_c) = 1$. This implies that our estimate for $\langle \psi | U_{C'}^\dagger Z_1^{z_1} \dots Z_{n-\hat{t}}^{z_{n-\hat{t}}} U_{C'} | \psi \rangle^2$ equals 1 for all $z \in \{0, 1\}^{n-\hat{t}}$,

and therefore the first $n - \hat{t}$ qubits of $U_{C'} |\psi\rangle$ are close to a computational-basis state.

We now bound the distance of $\rho_L = \text{tr}_{[n-\hat{t}]^c} [|\psi\rangle \langle \psi|]$ from a computational-basis state. By the union bound with failure probability δ , $M \geq 2 \log(1/\delta)/\epsilon$ Bell samples are sufficient to ensure that $|e(\sigma_c) - \langle \psi | \sigma_c | \psi \rangle|^2 < \epsilon$ and hence $|\langle \psi | \sigma_c | \psi \rangle|^2 > 1 - \epsilon$ [24]. By another union bound, this implies that $M \geq 2n \log(1/\delta)/\epsilon$ samples are sufficient to ensure that this is the case for all $c \in C'$. Let $s(\sigma_c) = \text{sign}(\langle \psi | \sigma_c | \psi \rangle)$. Then can use direct fidelity estimation [25] to compute the fidelity with the computational-basis state $|x\rangle$ that satisfies $\langle x | Z_c | x \rangle = s(\sigma_c)$, where $Z_c = U_{C'} \sigma_c U_{C'}^\dagger$. We find

$$\text{tr}[\rho_L |x\rangle \langle x|] = \frac{1}{2^{n-\hat{t}}} \sum_{c \in C'} [s(\sigma_c) \text{tr}[\rho_L Z_c]] \quad (\text{S86})$$

$$\geq \frac{1}{2^{n-\hat{t}}} \sum_{c \in C'} [s(\sigma_c) s(\sigma_c) (1 - \epsilon)] \quad (\text{S87})$$

$$= 1 - \epsilon, \quad (\text{S88})$$

since $\sqrt{1 - \epsilon} \geq 1 - \epsilon$.

What remains to be shown is that the state $U_{C'} |\psi\rangle$ is ϵ -close to the product state $|x\rangle |\varphi\rangle$ with some $|\varphi\rangle$ in fidelity. To see this, we observe that we can write $U_{C'} |\psi\rangle = \sum_y a_y |y\rangle |\varphi_y\rangle$ for some post-measurement states $|\varphi_y\rangle = \langle y | \otimes \mathbb{1} | \psi \rangle / \|(\langle y | \otimes \mathbb{1}) | \psi \rangle\|$. We find $|\alpha_x|^2 \geq 1 - \epsilon$, and setting $|\varphi\rangle := |\varphi_x\rangle$ proves the claim. \square

The tomography steps 4 returns x with exponentially low failure probability. By discarding those elements from D for which $x^i \neq x$ in step 5, we ensure that the remaining elements will yield an estimate $|\hat{\varphi}\rangle$ that has fidelity at least $1 - \epsilon$ with $|\varphi\rangle_x$, which yields the claim.

The overall runtime of the algorithm is polynomial in n and $2^{\hat{t}}$, since finding $U_{C'}$ is achieved by solving a linear system of equations, and we perform quantum state tomography on at most \hat{t} qubits in the last step.

4. Correctness of Algorithm 3

The estimate \hat{t} of the stabilizer nullity t of $|\psi\rangle$ clearly satisfies $\hat{t} \leq t$. It immediately follows from Lemma 3 that there is a state $|\phi\rangle = U_{C'}^\dagger |x\rangle |\varphi\rangle$ with fidelity $|\langle \phi | \psi \rangle|^2 \geq 1 - \epsilon$ and stabilizer nullity exactly \hat{t} . To see this, observe that $M(|\phi\rangle) = M(|x\rangle) + M(|\varphi\rangle) = 0 + \hat{t}$, since the stabilizer nullity is additive for product states.

We also note that we the subspace G' carries probability weight $1 - \epsilon$ of both the Bell distribution P_ψ and the characteristic distribution C_ψ of $|\psi\rangle$. The characteristic distribution of $|\psi\rangle$ is defined as $C_\psi(b) = 2^n |p(b)|^2$, where we write

$$|\psi\rangle \langle \psi| = \sum_b p(b) \sigma_b. \quad (\text{S89})$$

To show this directly, we formulate the following Lemma, which generalizes a well-known result of Erdős and Rényi [26] to nonuniform distributions over subspaces of \mathbb{F}_2^n .

Lemma 4 (Weighted subspace generation). *Let $S \subset \{0, 1\}^n$ be a binary subspace of dimension at most n and $P : S \rightarrow [0, 1]$ be a probability distribution over that subspace. Then the subspace $S' := \text{span}(s^1, \dots, s^M) \subset S$ spanned by M samples $s^i \leftarrow P$ with label $i = 1, \dots, M$ has the property that $\sum_{s \in S \setminus S'} P(s) < \epsilon$ with probability at least $1 - \delta$ whenever*

$$M \geq 2n \log(n) \log(2/\delta) / \epsilon. \quad (\text{S90})$$

Proof. We construct the subspace iteratively, observing that the first i samples which have been drawn generate a subspace $S_i = \text{span}(s^1, \dots, s^i)$. The $(i+1)$ th sample s^{i+1} now either lies in S_i or in its complement. If it lies in the complement the dimension of $S_{i+1} = \text{span}(S_i, s^{i+1})$ is increased by 1 compared to S_i , if it lies in S_i its dimension remains unchanged. Suppose the complement of S_i has probability weight at least ϵ . Then the probability that after drawing k additional samples, none of them lies in the complement of S_i is given by $1 - (1 - \epsilon)^k$ and hence $k = \log \delta / \log(1 - \epsilon) \geq 2 \log(1/\delta) / \epsilon$ samples are sufficient that $\dim(S_{i+k}) \geq \dim(S_i) + 1$ with probability at least $1 - \delta$. Repeating this argument n times, the total success probability is given by $(1 - \delta)^n$ unless $P(S \setminus S_i) < \epsilon$ for some $i \in [nk]$. Choosing the failure probability in every step as δ/n and observing that $(1 - \delta/n)^n \geq 1 - 2\delta$ proves the claim. \square

Now, we have that $\text{span}(\{b^i\}_i) \oplus k \subset G'$, since for $M > n+t$ the Bell samples are linearly dependent and hence at least one sample b^{j_0} is in the span of all others. Let b^M be such a sample. Then $\text{span}(\{b^{(i,M)}\}_i \in \text{span}(\{b^i\}_i)$ and therefore $\text{span}(\{b^{(i,j)}\}_i) \supset \text{span}(\{b^i\}_i)$. But this implies that $P_\psi(G' \oplus k) \geq 1 - \epsilon$.

Let $Q_\psi(a) = \sum_b P_\psi(b) P_\psi(b+a)$ be the distribution of the Bell difference samples. Then the differences $b^{(2i, 2i+1)}$ for $i \in \{0, \dots, (M-1)/2\}$ are distributed according to Q_ψ . It follows from Lemma 4 that if $M \geq 4n \log(n) \log(2/\delta) / \epsilon$, $Q_\psi(G') \geq 1 - \epsilon$ with probability at least $1 - \delta$. It follows from Proposition 10 of Ref. [27] that $C_\psi(G') \geq 1 - \epsilon$.

S4. ERROR DETECTION AND CORRECTION

In this section, we will outline some details of the error detection procedure, and discuss potential issues that arise due to noise in the Bell measurement itself.

A. Error reduction by error detection

In this section, we will calculate the amount of error reduction that it is possible by error detection using the global symmetry. Recall that our error detection procedure runs as follows.

Algorithm 5 Error detection

Input: Bell sample $r \leftarrow P_\rho$.

- 1: **if** $\pi_Y(r) = 1$ **then**
- 2: Declare an error and abort.
- 3: **end if**

Output: r

Let us now quantify the amount of error reduction that is possible using Algorithm 5. First, let us recall why the error detection is correct: Since we know that the purity of the ideal state is unity, all samples which lead to a purity away from unity must have been due to an error. These are exactly the samples from the antisymmetric subspace, i.e., samples r satisfying $\pi_Y(r) = 1$.

In the next step, we can compute the error reduction capabilities of the algorithm. We do so in the simplest possible model of noise: global white noise. In this model we write each copy of the state as

$$\rho = \rho_C(\eta) = (1 - \eta) |C\rangle \langle C| + \eta \mathbb{1}/2^n, \quad (\text{S91})$$

and hence the pre-measurement state as

$$\begin{aligned} \rho \otimes \rho = (1 - \eta)^2 |C\rangle \langle C|^{\otimes 2} + \frac{\eta(1 - \eta)}{2^n} (|C\rangle \langle C| \otimes \mathbb{1} \\ + \mathbb{1} \otimes |C\rangle \langle C|) + \frac{\eta^2}{2^{2n}} \mathbb{1}. \end{aligned} \quad (\text{S92})$$

Again, we start with the case of $s = 0$. We observe that the Bell distribution for the noisy state can be written as

$$P_\rho(r) = (1 - \eta)^2 P_C(r) + \frac{1}{4^n} P_e(\eta), \quad (\text{S93})$$

with the error probability $P_e(\eta) = 2\eta(1 - \eta) + \eta^2$, i.e., the distribution of the errors is uniform. Now, we observe that an error falls into the antisymmetric subspace with probability $D_{[2]}/2^{2n} = \frac{1}{2}(1 + 1/2^n)$. Hence, the probability that an error is detected is given by $P_e(\eta) \cdot \frac{1}{2}(1 + 1/2^n)$.

Now, we can consider the unnormalized postselected distribution $\tilde{Q}_\rho(r) = (1 - \eta)^2 P_C(r) + \delta_{[2]}(r)/D_{[2]} P_e(\eta)$, and normalize it as $\tilde{P}_\rho(r) = \tilde{Q}_\rho(r) / (\sum_r \tilde{Q}_\rho(r))$, where $\delta_{[2]}(r) = 1$ if $\pi_Y(r) = 0$ and 0 otherwise. In order to compute the effective error reduction, we find

$$\arg \min_\epsilon \|P_{\rho(\epsilon)} - \tilde{P}_{\rho(\eta)}\|_{\ell_1} = \eta/2 + O(\eta^2). \quad (\text{S94})$$

For $\eta \ll 1$ we thus find an error reduction by a factor of 2 from the first stage of the error detection algorithm. Error detection based just on the entire subsystem therefore recovers the probability of no error compared to running computations on a single copy.

In spite of doubling the number of qubits compared to standard-basis sampling, we have thus achieved the same overall post-selected fidelity at a given error rate.

B. Noise in the Bell measurement

Above, we have considered (local) noise in the state preparation while keeping the Bell measurements themselves error-

free. Of course, in an actual implementation of Bell sampling, the Bell measurements themselves will be noisy as well. The Bell measurement is constituted of transversal `cnot`-entangling gates and single-qubit measurements. The potential sources of error are therefore errors in the entangling gates and errors in the measurement apparatus.

Since we can move all the errors before the Bell measurement to the state preparation, which we have already discussed, we only have to consider errors *after* the entangling gates. First, consider single-qubit noise channels \mathcal{N} with noise strength ϵ after each two-qubit entangling gates before the measurements in the Hadamard and computational basis, i.e.,

$$|\Phi\rangle\langle\Phi|^+ \rightarrow \text{cnot } \mathcal{N}^{\otimes 2}(|+\rangle\langle+| \otimes |0\rangle\langle 0|) \text{cnot}, \quad (\text{S95})$$

and hence, the fidelity of the noisy measurement with the ideal measurement is just given by $\text{tr}[\mathcal{N}(|0\rangle\langle 0|)|0\rangle\langle 0|] \text{tr}[\mathcal{N}(|+\rangle\langle+|)|+\rangle\langle+|] \approx (1-\epsilon)^2$. The global measurement fidelity is then just $(1-\epsilon)^{2n}$, and hence for noise rate $\epsilon \ll 1/n$, the fidelity is sufficiently high. Notice that this is also the regime in which we can meaningfully use the cross-entropy benchmark [11]. Coherent errors such as `cnot`-over- or underrotations do not change the overall picture, since all entangling gates are carried out in parallel and hence the measurement fidelity just factorizes into the local fidelities.

Of course, antisymmetric errors in the measurement will also be detected by our detection procedure (assuming that they do not combine with antisymmetric errors from the state preparation). At a high level, these favourable properties of the Bell measurement with regards to their susceptibility to errors and their capabilities to detect them is what makes them fault-tolerant gadgets in stabilizer codes as well.

Let us note, however, that this analysis will be drastically different in architectures in which the `cnot`-entangling gates cannot be carried out in a single circuit layer at the end of the circuit such as geometrically local architectures. In such architectures the error contribution from the Bell measurement itself will be significant.

C. Virtual distillation using the Bell samples

We observe that even further error suppression is possible in the white-noise model for estimation of expectation values of observables that are diagonal in the Bell basis. Such observables can be written as $A = \sum_{i,j \in \{0,1\}} a_{ij} |\sigma_{ij}\rangle\langle\sigma_{ij}|$.

Writing an arbitrary noisy state preparation of $|\psi\rangle$ as

$$\rho_\psi(\epsilon) = (1-\epsilon)|\psi\rangle\langle\psi| + \epsilon\rho_\perp, \quad (\text{S96})$$

where $\text{tr}[\rho_\perp|\psi\rangle\langle\psi|] = 0$, we can write the expectation value with respect to $\rho = \rho_\psi(\epsilon)$ as

$$\begin{aligned} \text{tr}[A(\rho \otimes \rho)] &= (1-\epsilon)^2 \text{tr}[A|\psi\rangle\langle\psi|^{\otimes 2}] \\ &+ \epsilon(1-\epsilon) \text{tr}[A(\rho_\perp \otimes |\psi\rangle\langle\psi| + |\psi\rangle\langle\psi| \otimes \rho_\perp)] + \epsilon^2 \text{tr}[A\rho_\perp^{\otimes 2}]. \end{aligned} \quad (\text{S97})$$

Combined with the purity estimate, we can then estimate the ideal expectation value $\langle A \rangle_\psi = \text{tr}[A|\psi\rangle\langle\psi|^{\otimes 2}]$ from the noisy Bell sampling data from $\rho_\psi(\epsilon)$

$$\begin{aligned} \langle \hat{A} \rangle_\psi &= \frac{\text{tr}[A(\rho \otimes \rho)]}{\text{tr}[\mathbb{S}(\rho \otimes \rho)]} = (1 - \text{tr}[\rho_\perp^2]\epsilon^2) \langle A \rangle_\psi + \epsilon^2 \text{tr}[A\rho_\perp^{\otimes 2}] \\ &+ (\epsilon + \epsilon^2) \text{tr}[A(\rho_\perp \otimes |\psi\rangle\langle\psi| + |\psi\rangle\langle\psi| \otimes \rho_\perp)] + O(\epsilon^3). \end{aligned} \quad (\text{S98})$$

Here, we have used that $\text{tr}[\rho_\psi(\epsilon)^2] = (1-\epsilon)^2 + \epsilon^2 \text{tr}[\rho_\perp^2]$.

In particular, choosing $A = (P \otimes P)\mathbb{S}$, we can estimate $\langle\psi|P|\psi\rangle^2$ with error suppression ϵ^2 . To see this, observe that $\text{tr}[A(\rho \otimes \rho)] = \text{tr}[P\rho P\rho]$ and

$$\begin{aligned} \langle \hat{P} \rangle_\psi^2 &= \frac{\text{tr}[P\rho P\rho]}{\text{tr}[\rho^2]} = (1 - \text{tr}[\rho_\perp]^2\epsilon^2) \langle P \rangle_\psi^2 \\ &+ (\epsilon + \epsilon^2) \langle\psi|P\rho_\perp P|\psi\rangle + \epsilon^2 \text{tr}[P\rho_\perp P\rho_\perp] + O(\epsilon^3). \end{aligned} \quad (\text{S99})$$

Generically, the term $\langle\psi|P\rho_\perp P|\psi\rangle$ will be exponentially suppressed and hence we obtain an error suppression from ϵ to $\min\{\epsilon/2^n, \epsilon^2\}$.

As a concrete example, consider the case where $\rho \propto e^{-\beta H}$ for a local gapped Hamiltonian with gap Δ satisfying $\beta\Delta \gg 1$, and we are interested in estimating local Pauli expectation values in the ground state $|E_0\rangle$. Using the two-copy observable $A = (P \otimes P)\mathbb{S}$ we obtain the estimator $\text{tr}[P\rho P\rho]/\text{tr}[\rho^2] = |\langle E_0|P|E_0\rangle|^2 + O(e^{-2\beta\Delta})$. This identity follows because $\langle E_n|P|E_0\rangle$ for $n \neq 0$ is generically suppressed as an inverse polynomial in n for local gapped Hamiltonians and local operators P [28]. As a result, we can use post-processing of the Bell samples to virtually “cool” the system to half the temperature of the initial state.

S5. APPLYING THE NOISY SIMULATION ALGORITHM TO BELL SAMPLING

In this section, we consider whether the noisy simulation algorithm of Gao and Duan [29] and Aharonov *et al.* [30] applies to Bell sampling. Before we start, let us briefly recap the algorithm. We will use the notation of Ref. [30].

A. Recap of the algorithm

The key idea is to write an output probability $P(C, x)$ of a random circuit $C = U_d U_{d-1} \dots U_1$ with Haar-random two-qubit gates U_i as a path integral

$$\begin{aligned} P(C, x) &= \sum_{s_0, \dots, s_d \in \mathbb{P}_n} \text{tr}[|x\rangle\langle x| s_d] \text{tr}[s_d U_d s_{d-1} U_d^\dagger] \dots \\ &\dots \text{tr}[s_1 U_1 s_0 U_1^\dagger] \text{tr}[s_0 |0^n\rangle\langle 0^n|] \end{aligned} \quad (\text{S100})$$

$$\equiv \sum_{s \in \mathbb{P}_n^{d+1}} \langle x | s_d \rangle \langle s_d | \mathcal{U}_d | s_{d-1} \rangle \dots \langle s_1 | \mathcal{U}_1 | s_0 \rangle \langle s_0 | 0^n \rangle \quad (\text{S101})$$

$$= \sum_{s \in \mathbb{P}_n^{d+1}} f(C, s, x), \quad (\text{S102})$$

where P_n is the n -qubit Pauli group. This expression can be easily seen from the fact that the Pauli matrices form a complete operator basis and therefore $\text{tr}[U\rho U^\dagger s] = \sum_{t \in P_n} \text{tr}[UtU^\dagger s] \text{tr}[\rho t]$. We also write $\mathcal{U} := U \cdot U^\dagger$ and $\text{tr}[ab] = \langle\langle a|b \rangle\rangle$. Notice that in writing the path integral (S129), we have normalized the Pauli matrices to $\text{tr}[pp'] = \delta_{p,p'}$ for $p, p' \in P_n$.

We can also think of the Pauli path integral as a Fourier decomposition of the output probabilities. In this Fourier representation, the effect of local depolarizing noise can be easily analyzed since it just acts as $\mathcal{E}(\rho) = (1 - \epsilon)\rho + \epsilon \text{tr}[\rho] \mathbb{1}/2^n$. The contribution of a Pauli path of a noisy quantum circuit to the total output probability thus decays with the number of non-identity Pauli operators in it (the Hamming weight of s) as

$$\tilde{P}(C, x) = \sum_{s \in P_n^{d+1}} (1 - \epsilon)^{|s|} f(C, s, x). \quad (\text{S103})$$

The algorithm of Aharonov *et al.* [30] is based on approximating this sum by truncating it to paths with weight $|s| \leq \ell$ for some $\ell \in \mathbb{N}$. The approximation can then be computed in time $2^{O(\ell)}$. Furthermore, since the output string just appears as $\langle\langle x|s_d \rangle\rangle$, any marginal can be computed at the same complexity. Replacing the outcome string $x \in \{0, 1\}^n$ by a string y with characters $\{0, 1, \bullet\}^n$, whenever there is a \bullet at position i of y we write $\mathbb{1}_2$ at i , and eventually trace over $\prod_{i: y_i \in \{0, 1\}} |y_i\rangle \langle y_i|_i \otimes \mathbb{1}$, which is efficient. The algorithm then samples from the truncated distribution $\tilde{p}_C^{(\ell)}$ using marginal sampling. What remains to be shown is that the total-variation distance $\Delta = \text{TVD}(\tilde{p}_C^{(\ell)}, \tilde{p}_C^{(\ell)})$ between the truncated and the noisy distribution is sufficiently small in ℓ to give an efficient algorithm. To this end, they provide an upper bound on $\mathbb{E}_C[\Delta^2]$ using the Cauchy-Schwarz inequality, see Secs. 2 & 3 of Ref. [30].

Here, we consider two possible strategies to prove the algorithm remains efficient for Bell sampling. We show that the first strategy fails, and give evidence that the second strategy fails. Together, this provides some evidence that the algorithm does not work, making Bell sampling a compelling candidate for noise-resilient sampling.

B. Strategy 1: Upper bounds on the trace distance

The first strategy we consider is to adapt the upper bound of Aharonov *et al.* [30] on TVD to an upper bound on the trace distance. An upper bound ϵ on the trace distance of the sampled quantum state shows that the optimal single-copy measurement distinguishing probability is given by ϵ . In Bell sampling we perform a two-copy measurement, and we can bound the two-copy trace distance in terms of the single-copy trace distance as

$$\|\rho \otimes \rho - \sigma \otimes \sigma\|_1 = \|\rho \otimes \rho - \rho \otimes \sigma + \rho \otimes \sigma - \sigma \otimes \sigma\|_1 \quad (\text{S104})$$

$$\leq \|\rho \otimes (\rho - \sigma)\|_1 + \|(\rho - \sigma) \otimes \sigma\|_1 \quad (\text{S105})$$

$$= 2\|\rho - \sigma\|_1 \quad (\text{S106})$$

Consider the ideal pre-measurement state in the path integral formulation

$$\rho(C) = \sum_{s \in P_n^{d+1}} |s_d\rangle \langle\langle s_d|\mathcal{U}_d|s_{d-1}\rangle\rangle \cdots \langle\langle s_1|\mathcal{U}_1|s_0\rangle\rangle \langle\langle s_0|0^n\rangle\rangle \quad (\text{S107})$$

$$=: \sum_{s \in P_n^{d+1}} g(C, s) |s_d\rangle. \quad (\text{S108})$$

Then we can write the noisy state as well as the state which is effectively generated when truncating the noisy path integral to paths of weight $\leq \ell$ as

$$\tilde{\rho}(C) := \sum_{s \in P_n^{d+1}} (1 - \gamma)^{|s|} g(C, s) |s_d\rangle, \quad (\text{S109})$$

$$\tilde{\rho}^\ell(C) := \sum_{s: |s| \leq \ell} (1 - \gamma)^{|s|} g(C, s) |s_d\rangle \quad (\text{S110})$$

$$\Delta \tilde{\rho}(C) := \sum_{s: |s| > \ell} (1 - \gamma)^{|s|} g(C, s) |s_d\rangle. \quad (\text{S111})$$

Analogously to Eq. (25) of Aharonov *et al.* [30], we can bound the trace distance

$$\|\Delta \tilde{\rho}\|_1^2 \equiv \mathbb{E}_C \|\tilde{\rho}^\ell(C) - \tilde{\rho}(C)\|_1^2 \quad (\text{S112})$$

$$\leq 2^n \mathbb{E}_C \|\Delta \tilde{\rho}(C)\|_2^2 \quad (\text{S113})$$

$$= 2^n \mathbb{E}_C \text{tr}[(\Delta \tilde{\rho}(C))^\dagger \Delta \tilde{\rho}(C)] \quad (\text{S114})$$

$$= 2^n \mathbb{E}_C \sum_{s, s': |s|, |s'| > \ell} (1 - \gamma)^{2|s|} \langle\langle s'_d|s_d \rangle\rangle g(C, s) g(C, s') \quad (\text{S115})$$

$$= 2^n \sum_{s: |s| > \ell} (1 - \gamma)^{2|s|} \mathbb{E}_C g(C, s)^2, \quad (\text{S116})$$

using the Cauchy-Schwarz inequality and orthogonality of the coefficients $g(C, s)$ (note that this holds since it is just a local property of $\mathbb{E}_U[\langle\langle q|\mathcal{U}|p \rangle\rangle \langle\langle r|\mathcal{U}|s \rangle\rangle]$).

Hence, in order to upper bound $\|\Delta \tilde{\rho}\|_1$ we need to upper bound $\sum_s \mathbb{E}_C [g(C, s)^2]$ for certain values of s . Aharonov *et al.* [30] achieve this by upper bounding the total sum over all s .

Following their strategy, we define a Fourier weight

$$V_k = 2^n \mathbb{E}_C \sum_{s: |s|=k} g(C, s)^2, \quad (\text{S117})$$

and compute its properties.

We certainly have

$$V_0 = 1 \quad (\text{S118})$$

$$V_k = 0 \quad \forall 0 < k \leq d. \quad (\text{S119})$$

To see this, we just follow the argument of Aharonov *et al.* [30]. In particular $V_0 = 2^n \langle\langle \mathbb{1}|\mathbb{1} \rangle\rangle^{2d} \langle\langle \mathbb{1}|0 \rangle\rangle^2 = 1$ since $\langle\langle \mathbb{1}|0 \rangle\rangle = 1/\sqrt{2^n}$ and $\langle\langle \mathbb{1}|\mathbb{1} \rangle\rangle = 1$.

What remains is to compute $\sum_{k \geq d+1} V_k$. To this end, we can compute—using a 2-design assumption on C —the total

Fourier weight

$$\sum_{k \geq 0} V_k = 2^n \mathbb{E}_C \sum_{k \geq 0} \sum_{s: |s|=k} g(C, s)^2 \quad (\text{S120})$$

$$= 2^n \mathbb{E}_C \sum_{s', s} g(C, s) g(C, s') \quad (\text{S121})$$

$$= 2^n \mathbb{E}_C \left(\sum_{s \in \mathbb{P}_n^{d+1}} g(C, s) \right)^2 \quad (\text{S122})$$

$$= 2^n \mathbb{E}_C \left(\sum_{p \in \mathbb{P}_n} \langle C | p | C \rangle \right)^2 \quad (\text{S123})$$

$$= 2^n \sum_{p, p'} \text{tr} \left[\mathbb{E}_C |C\rangle \langle C|^{\otimes 2} (p \otimes p') \right] \quad (\text{S124})$$

$$\stackrel{C \text{ 2-des}}{=} \frac{2^n}{2D_{[2]}} \sum_{p, p'} \text{tr}[(\mathbb{1} + \mathbb{S})(p \otimes p')] \quad (\text{S125})$$

$$= \frac{2^n}{2D_{[2]}} \sum_{p, p'} (2^n \delta_{p, \mathbb{1}} \delta_{p', \mathbb{1}} + \delta_{p, p'}) \quad (\text{S126})$$

$$= \frac{2^n}{D_{[2]}} \frac{2^n + 4^n}{2} = 2^n \quad (\text{S127})$$

Here, we have used orthogonality in reverse, and in lines (S122) and (S123), we have used that $\rho = \sum_p p \text{tr}[\rho p] = \sum_s s_d g(C, s)$ and hence $\text{tr}[\rho p] = \sum_{s: s_d=p} g(C, s)$.

Putting everything together, analogously to Aharonov *et al.* [30, Eq. (29)] we find that

$$\|\Delta \tilde{\rho}\|_1^2 \leq 2^n (1 - \gamma)^{2\ell}, \quad (\text{S128})$$

which remains trivial for $\ell \in o(n)$.

The same argument as in Ref. [30] thus cannot be used to show that the algorithm works for any measurement strategy. One might wonder if we can tighten the bound. We argue that we cannot.

First, observe that the only strict inequality we have used is to bound the trace distance by the Frobenius norm in Eq. (S114). We can also hardly hope to remove the factor of 2^n incurred in this bound, however, because we expect the state to be spread out in Hilbert space and the bound is tight for the uniform distribution.

Second, observe that the trace distance upper bound is dominated by the sum over all 4^n Pauli matrices. Compare this to the original algorithm, where the only non-zero contributions to the final measurement outcomes were Pauli paths which ended in a Z -type string since the overlap with a computational basis state was computed. This is a reduction by precisely the factor of 2^n which we find in our upper bound on the trace distance. Since the Bell distribution overlaps with almost all Pauli strings, we expect the upper bound to be similar when done directly in the Bell basis.

C. Strategy 2: The argument in the Bell basis

Indeed, an alternative proof strategy is to directly upper bound the total-variation distance of the truncated Bell distribution. To this end we write the Bell-basis Pauli path integral

as

$$P(C, r) = \sum_{s \in \mathbb{P}_{2n}^{d+1}} \langle \Phi^+ | \mathcal{P}_r^\dagger \otimes \mathbb{1}_n | s_d \rangle \langle s_d | \mathcal{U}_d^{\otimes 2} | s_{d-1} \rangle \cdots \quad (\text{S129})$$

$$\cdots \langle s_1 | \mathcal{U}_1^{\otimes 2} | s_0 \rangle \langle s_0 | 0^{2n} \rangle \quad (\text{S130})$$

$$=: \sum_{s \in \mathbb{P}_{2n}^{d+1}} f(C, s, r) \quad (\text{S131})$$

Writing $s_i = (s_i^0, s_i^1)$ with $s_i^j \in \mathbb{P}_n$, we observe that $\langle \Phi^+ | \mathcal{P}_r^\dagger \otimes \mathbb{1}_n | s_d \rangle = (-1)^{\langle P_r, s_i^0 \rangle + \delta(s_i^0, Y)} \delta(s_i^0, s_i^1)$, where $\langle p, q \rangle = 1$ if p and q anticommute and zero otherwise. The boundary condition for the end of the Pauli path is therefore that both branches—copies of the system—must end at the same n -qubit Pauli.

Valid paths are therefore of the form

$$(s_0^0, s_0^1) \rightarrow (s_1^0, s_1^1) \rightarrow \cdots \rightarrow (s_{d-1}^0, s_{d-1}^1) \rightarrow (s_d^0, s_d^0), \quad (\text{S132})$$

and there are $(4^{2n})^d \cdot 4^n = 4^{n(2d+1)}$ of them.

Next when we add noise to the circuit, the sum again transforms to

$$\tilde{P}(C, r) = \sum_{s \in \mathbb{P}_{2n}^{d+1}} (1 - \gamma)^{|s|} f(C, s, r). \quad (\text{S133})$$

Let us now bound the total-variation distance between an ℓ -truncated noisy path integral \tilde{P}_ℓ and the non-truncated path integral (S133) as

$$\mathbb{E}_C[\Delta^2] = \mathbb{E}_C \left(\sum_r |\tilde{P}(C, r) - \tilde{P}_\ell(C, r)| \right)^2 \quad (\text{S134})$$

$$\leq 2^{2n} \mathbb{E}_C \sum_r (\tilde{P}(C, r) - \tilde{P}_\ell(C, r))^2 \quad (\text{S135})$$

$$\leq 2^{2n} \mathbb{E}_C \sum_r \left(\sum_{s: |s| > \ell} (1 - \gamma)^{|s|} f(C, s, r) \right)^2 \quad (\text{S136})$$

Aharonov *et al.* [30] now go forward to bound their expression analogous to (S136) (up to scaling and replacing $r \leftarrow x$. Letting

$$W_k = 2^{2n} \mathbb{E}_C \sum_{s \in \mathbb{P}_n^{d+1}: |s|=k} f(C, s, 0^n)^2 \quad (\text{S137})$$

be the total Fourier weight of a circuit at degree k , they use the following properties:

- orthogonality of the Fourier coefficients, i.e.,

$$\mathbb{E}_C[f(C, s, x) f(C, s', x)] = 0, \quad \forall s \neq s' \quad (\text{S138})$$

- Bounds on the total Fourier weight

$$W_0 = 1 \quad (\text{S139})$$

$$W_k = 0, \quad \forall 0 < k < d, \quad (\text{S140})$$

$$\sum_{k \geq d+1} W_k \in O(1), \quad (\text{S141})$$

which follow from anticoncentration,

$$2^n \mathbb{E}_C[p(C, x)^2] \in O(2^{-n}). \quad (\text{S142})$$

Notice that all of these properties are second-moment properties. In the Bell sampling, these become fourth moment properties.

a. Orthogonality Let us begin by considering the orthogonality property.

To this end, consider a single gate in the circuit. The exact expression to compute is

$$\mathbb{E}_{U \sim \mu} \langle p^0 | \mathcal{U} | q^0 \rangle \langle r^0 | \mathcal{U} | s^0 \rangle \langle p^1 | \mathcal{U} | q^1 \rangle \langle r^1 | \mathcal{U} | s^1 \rangle, \quad (\text{S143})$$

$p, q, r, s \in \mathbb{P}_{2^4}$. If μ is the Haar measure, using Weingarten calculus, we can rewrite the expression as

$$\sum_{\pi, \sigma \in S_4} \text{Wg}(\pi \sigma^{-1}) \text{tr}(W_\pi(p^0 \otimes p^1 \otimes r^0 \otimes r^1)) \times \text{tr}(W_\sigma(q^0 \otimes q^1 \otimes s^0 \otimes s^1)). \quad (\text{S144})$$

which does not obviously simplify.

Fortunately, to show orthogonality in the single-copy case (which involves only second moments), we can make use of the right-invariance of our gate set under multiplication with Pauli matrices, i.e., $\mathcal{G} \cdot p = \mathcal{G}$ for any $p \in \mathbb{P}_2$. Hence, we can insert an expectation value over Pauli matrices,

$$\mathbb{E}_{U \sim \mathcal{G}} f(U) = \mathbb{E}_{U \sim \mathcal{G}} \mathbb{E}_{v \sim \mathbb{P}_2} f(Uv). \quad (\text{S145})$$

In our (two-copy) case, we therefore get

$$\mathbb{E}_{U \sim \mathcal{H}} \langle p^0 | \mathcal{U} | q^0 \rangle \langle r^0 | \mathcal{U} | s^0 \rangle \langle p^1 | \mathcal{U} | q^1 \rangle \langle r^1 | \mathcal{U} | s^1 \rangle \quad (\text{S146})$$

$$= \mathbb{E}_{U \sim \mathcal{H}} \mathbb{E}_{v \sim \mathbb{P}_2} \langle p^0 | \mathcal{U} v | q^0 \rangle \langle r^0 | \mathcal{U} v | s^0 \rangle \langle p^1 | \mathcal{U} v | q^1 \rangle \langle r^1 | \mathcal{U} v | s^1 \rangle \quad (\text{S147})$$

$$= \mathbb{E}_{U \sim \mathcal{H}} \mathbb{E}_{v \sim \mathbb{P}_2} \text{tr}[p^0 U v q^0 v^\dagger U^\dagger] \text{tr}[r^0 U v s^0 v^\dagger U^\dagger] \quad (\text{S148})$$

$$\text{tr}[p^1 U v q^1 v^\dagger U^\dagger] \text{tr}[r^1 U v s^1 v^\dagger U^\dagger]. \quad (\text{S149})$$

Hence, for orthogonality to hold it suffices that

$$\mathbb{E}_{v \in \mathbb{P}_2} [v^{\otimes 4} (p \otimes q \otimes r \otimes s) (v^\dagger)^{\otimes 4}] = 0. \quad (\text{S150})$$

Indeed,

$$\mathbb{E}_{v \in \mathbb{P}_2} [v^{\otimes 4} (p \otimes q \otimes r \otimes s) (v^\dagger)^{\otimes 4}] \quad (\text{S151})$$

$$= \frac{1}{16} \sum_{v \in \mathbb{P}_2} (-1)^{\langle v, p \rangle + \langle v, q \rangle + \langle v, r \rangle + \langle v, s \rangle} p \otimes q \otimes r \otimes s \quad (\text{S152})$$

$$= \frac{1}{16} \sum_{v \in \mathbb{P}_2} (-1)^{\langle v, pqrs \rangle} p \otimes q \otimes r \otimes s \quad (\text{S153})$$

$$= 0, \quad \forall pq \neq i^k rs, \quad (\text{S154})$$

Eq. (S154) gives us the orthogonality property

$$\mathbb{E}_{C \sim \mathcal{D}} f(C, q, r) f(C, s, r) = 0, \quad \forall q^0 s^0 \neq i^k q^1 s^1. \quad (\text{S155})$$

The proof follows straightforwardly from (S154), noting that the last layer of Paulis is always the same across q and s .

While the condition $q \neq s$ that arises in the single-copy case considered by Aharonov *et al.* [30] reduces summation over \mathbb{P}_2^2 to summation over \mathbb{P}_2 , the condition $pq \neq rs$ we find for the 2-copy case reduces summation over \mathbb{P}_4^2 to summation over $\mathcal{S} = \{p, q, r, s \in \mathbb{P}_2 : pq = i^k rs\}$. We have $|\mathcal{S}| = 4^{2n} \cdot 4^n = 4^{3n}$ since we can choose p, q freely and then the pair r, s is constrained to $rs = i^k pq$ for some k of which there are 4^n choices. As a result, similar to the trace-distance calculation, we find an additional exponential summation over Pauli strings, which blows up the sum by a factor of 4^n .

Notice, however, that the orthogonality condition (S154) we have found is only a sufficient condition for the expectation (S143) to vanish, and it could be that the full Haar average has more zero terms. We expect, however, that condition (S154) is the only condition. To prove that this is indeed the case, one needs to compute the fourth moments (S143)—a task that we leave to future work.

To summarize, in both strategies we run into an exponential blow-up in the summation that, as of now, results in an exponential upper bound on the TVD between the sampled distribution and the target distribution.

-
- [1] D. Hangleiter and J. Eisert, *Computational Advantage of Quantum Random Sampling*, *Rev. Mod. Phys.* **95**, 035001 (2023), [arxiv:2206.04079](#).
 - [2] M. J. Bremner, A. Montanaro, and D. J. Shepherd, *Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations*, *Physical Review Letters* **117**, 080501 (2016).
 - [3] H. Krovi, *Average-Case Hardness of Estimating Probabilities of Random Quantum Circuits with a Linear Scaling in the Error*

- Exponent*, (2022), [arxiv:2206.05642](#).
- [4] H. Zhu, R. Kueng, M. Grassl, and D. Gross, *The Clifford Group Fails Gracefully to Be a Unitary 4-Design*, arXiv:1609.08172 [quant-ph] (2016), [arxiv:1609.08172](#).
- [5] F. G. S. L. Brandão, A. W. Harrow, and M. Horodecki, *Local Random Quantum Circuits Are Approximate Polynomial-Designs*, *Commun. Math. Phys.* **346**, 397 (2016), [arxiv:1208.0692](#).
- [6] T. Zhou and A. Nahum, *Emergent Statistical Mechanics of*

- Entanglement in Random Unitary Circuits*, *Phys. Rev. B* **99**, 174205 (2019).
- [7] N. Hunter-Jones, *Unitary Designs from Statistical Mechanics in Random Quantum Circuits*, (2019), [arxiv:1905.12053](#).
- [8] B. Barak, C.-N. Chou, and X. Gao, Spoofing Linear Cross-Entropy Benchmarking in Shallow Quantum Circuits, in *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 185, edited by J. R. Lee (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2021) pp. 30:1–30:20, [arxiv:2005.02421](#).
- [9] A. M. Dalzell, N. Hunter-Jones, and F. G. S. L. Brandão, *Random Quantum Circuits Anticoncentrate in Log Depth*, *PRX Quantum* **3**, 010333 (2022), [arxiv:2011.12277](#).
- [10] J. J. Wallman and J. Emerson, *Noise Tailoring for Scalable Quantum Computation via Randomized Compiling*, *Phys. Rev. A* **94**, 052325 (2016).
- [11] B. Ware, A. Deshpande, D. Hangleiter, P. Niroula, B. Fefferman, A. V. Gorshkov, and M. J. Gullans, *A Sharp Phase Transition in Linear Cross-Entropy Benchmarking*, (2023), [arxiv:2305.04954](#).
- [12] J. F. Fitzsimons and E. Kashefi, *Unconditionally Verifiable Blind Quantum Computation*, *Physical Review A* **96** (2017), 10.1103/PhysRevA.96.012303, [arxiv:1203.5217](#).
- [13] B. W. Reichardt, F. Unger, and U. Vazirani, *Classical Command of Quantum Systems*, *Nature* **496**, 456 (2013).
- [14] U. Mahadev, *Classical Verification of Quantum Computations*, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (2018) pp. 259–267, [arxiv:1804.01082](#).
- [15] D. N. Page, *Average Entropy of a Subsystem*, *Phys. Rev. Lett.* **71**, 1291 (1993).
- [16] S. Garnerone, T. R. de Oliveira, and P. Zanardi, *Typicality in Random Matrix Product States*, *Phys. Rev. A* **81**, 032336 (2010).
- [17] B. Collins, C. E. Gonzalez-Guillen, and D. Perez-Garcia, *Matrix Product States, Random Matrix Theory and the Principle of Maximum Entropy*, *Communications in Mathematical Physics* **320**, 663 (2013), [arxiv:1201.6324](#).
- [18] M. Fukuda and R. Koenig, *Typical Entanglement for Gaussian States*, *Journal of Mathematical Physics* **60**, 112203 (2019), [arxiv:1903.04126](#).
- [19] J. T. Iosue, A. Ehrenberg, D. Hangleiter, A. Deshpande, and A. V. Gorshkov, *Page Curves and Typical Entanglement in Linear Optics*, *Quantum* **7**, 1017 (2023), [arxiv:2209.06838](#).
- [20] G. M. Sommers, D. A. Huse, and M. J. Gullans, *Crystalline Quantum Circuits*, (2023), [arxiv:2210.10808](#).
- [21] A. Montanaro, *Learning Stabilizer States by Bell Sampling*, (2017), [arxiv:1707.04012](#).
- [22] M. Guță, J. Kahn, R. Kueng, and J. A. Tropp, *Fast State Tomography with Optimal Error Bounds*, *J. Phys. A: Math. Theor.* **53**, 204001 (2020).
- [23] H.-Y. Huang, R. Kueng, and J. Preskill, *Information-Theoretic Bounds on Quantum Advantage in Machine Learning*, *Phys. Rev. Lett.* **126**, 190505 (2021).
- [24] V. Mnih, C. Szepesvári, and J.-Y. Audibert, *Empirical Bernstein Stopping*, in *Proceedings of the 25th International Conference on Machine Learning*, ICML '08 (Association for Computing Machinery, Helsinki, Finland, 2008) pp. 672–679.
- [25] S. T. Flammia and Y.-K. Liu, *Direct Fidelity Estimation from Few Pauli Measurements*, *Physical Review Letters* **106** (2011), 10.1103/PhysRevLett.106.230501.
- [26] P. Erdős and A. Rényi, *Probabilistic Methods in Group Theory*, *J. Anal. Math.* **14**, 127 (1965).
- [27] S. Grewal, V. Iyer, W. Kretschmer, and D. Liang, *Efficient Learning of Quantum States Prepared With Few Non-Clifford Gates II: Single-Copy Measurements*, (2023), [arxiv:2308.07175](#).
- [28] M. B. Hastings, *Locality in Quantum Systems*, [arXiv:1008.5137](#) (2010).
- [29] X. Gao and L. Duan, *Efficient Classical Simulation of Noisy Quantum Computation*, (2018), [arxiv:1810.03176](#).
- [30] D. Aharonov, X. Gao, Z. Landau, Y. Liu, and U. Vazirani, *A Polynomial-Time Classical Algorithm for Noisy Random Circuit Sampling*, (2022), [arxiv:2211.03999](#).