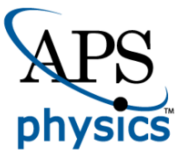

Sample Complexity of Device-Independently Certified "Quantum Supremacy"

Dominik Hangleiter, Martin Kliesch, Jens Eisert, and Christian Gogolin

Phys. Rev. Lett. **122**, 210502 — Published 29 May 2019

DOI: 10.1103/PhysRevLett.122.210502

Sample complexity of device-independently certified "quantum supremacy"

Dominik Hangleiter,[1] Martin Kliesch,[2] Jens Eisert,[1,3] and Christian Gogolin[4,5,6]

[1]*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*
[2]*Institute for Theoretical Physics, Heinrich Heine University Düsseldorf, 40225 Düsseldorf, Germany*
[3]*Department of Mathematics and Computer Science, Freie Universität Berlin, 14195 Berlin, Germany*
[4]*ICFO-Institut de Ciencies Fotoniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels (Barcelona), Spain*
[5]*Institute for Theoretical Physics, University of Cologne, 50937 Köln, Germany*
[6]*Xanadu, 372 Richmond St W, Toronto, M5V 1X6, Canada*

Results on the hardness of approximate sampling are seen as important stepping stones towards a convincing demonstration of the superior computational power of quantum devices. The most prominent suggestions for such experiments include boson sampling, IQP circuit sampling, and universal random circuit sampling. A key challenge for any such demonstration is to certify the correct implementation. For all these examples, and in fact for all sufficiently flat distributions, we show that any non-interactive certification from classical samples and a description of the target distribution requires exponentially many uses of the device. Our proofs rely on the same property that is a central ingredient for the approximate hardness results: namely, that the sampling distributions, as random variables depending on the random unitaries defining the problem instances, have small second moments.

Quantum sampling devices have been hailed as promising candidates for the demonstration of "quantum (computational) supremacy"[1] [1]. The goal of any such experiment is to unambiguously demonstrate that quantum devices can solve some tasks both faster and with a more favourable scaling of the computational effort than any classical machine. At the same time, in the near term it is bound to use those small and computationally restricted quantum devices that are available before the arrival of universal, scalable, and fault-tolerant quantum computers. This challenge has sparked a flurry of experimental activity [2–7] and prompted the development of better classical sampling schemes for exact [8, 9] and imperfect realizations [10–13]. Due to the reality of experimental imperfections, the key theoretical challenge — achieved in Refs. [14–22] using Stockmeyer's approximate counting algorithm [23] — is to prove that even *approximately* sampling from the output distribution of the quantum device is classically hard.

In any such demonstration, the issue of certification is of outstanding importance [10, 17, 24–27]: To demonstrate something non-trivial, one not only needs to build a device that is designed to sample approximately from a classically hard distribution but at the same time, one needs to ensure from a feasible number of uses of the device (or its parts) that it actually achieves the targeted task. How can one convince a skeptical certifier that a quantum device, which supposedly does something no classical machine can do, actually samples from a distribution that is close enough to the ideal target distribution?

The arguably most elegant and most convincing certification would be one based on purely classical data, ideally only the samples produced by the device and a description of the target distribution. Such certification would be free of additional complexity-theoretic assumptions and device-independent, in that it would be agnostic to all implementation details of the device and would directly certify that the classically defined sampling problem was solved.

In this work, we rigorously prove for a broad range of sampling problems, specifically for boson sampling [14], universal random circuit sampling [15, 17], IQP circuit sampling [16, 24], and sampling from post-selected-universal 2-designs [20–22, 28, 29] that they cannot be efficiently certified from classical samples and a description of the target probability distribution. Ironically, it turns out that the same property of a distribution that allows to prove the known approximate-hardness results also forbids their non-interactive sample-efficient device independent certification, to the effect that with the known proof methods both properties cannot be achieved simultaneously in such schemes. We directly bound the sample complexity of certification, which means that we automatically also lower bound the computational complexity and that our results cannot be circumvented by increasing the classical computational power of the certifier.

The specific question of certification we focus on here is

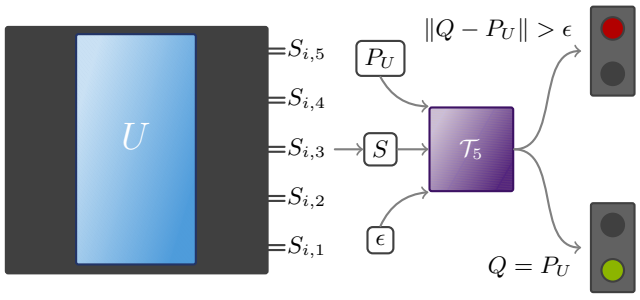

Figure 1. We consider the problem of certifying probability distributions of the form $P_U(S) = |\langle S|U|S_0\rangle|^2$ with an input state $|S_0\rangle = |0\rangle^{\otimes n}$ and a unitary $U \sim \mu_n$ drawn from some measure μ_n. Given $\epsilon > 0$ and access to an arbitrary-precision description of the target distribution P_U, the test $\mathcal{T}_n$ treats the sampler as a black box and receives a sequence $\mathcal{S} = (S_i)_{i=1}^s \sim Q$ of s samples from an unknown distribution Q. Given S the test is asked to output "Accept" if $Q = P_U$ and "Reject" if $\|Q - P_U\|_1 > \epsilon$ with high probability.

[1] Acknowledging the recent debate, we use the term "quantum (computational) supremacy" strictly in its established technical meaning [1].

(see Figure 1): Given unlimited computational power and a full description of the target distribution, how many samples from an unknown distribution are required to guarantee that this distribution is either identical to the target distribution or at least some preset distance away from it? This problem of distinguishing one (target) distribution from all sufficiently different alternatives is known as *identity testing* [30] in the property testing literature. Identity testing is an easier task than its robust version in which the certifier is moreover required to accept a constant-size region around the target distribution [26, 31]. At the same time, it is much harder than mere *state-discrimination*, where the task is to differentiate between two fixed distributions.

Lower bounds on the sample complexity of restricted state-discrimination scenarios [10] prompted the development of schemes [25] that allow to corroborate and build trust in experiments [6, 7, 32]. This helped spark interest in the problem of device-independent certification — on which there had not been much progress since [24]. In contrast to previous work [10], here, the certifier is given a full description of the target distribution[2] and unlimited computational power.

Our proof, detailed in the Supplementary Material [33], makes use of a key property for the proof of hardness of approximate sampling, namely an upper bound on the second moments of the output probabilities with respect to the choice of a random unitary specifying the instance of the sampling problem. The bound on the second moments implies that the probabilities are concentrated around the uniform distribution and hence an anti-concentration property[3]. This anti-concentration allows lifting results on the hardness of approximate sampling up to relative errors to ones for additive errors — provided relative-error approximation of the output probabilities is hard *on average*. It is thus a key property to prove hardness in the physically relevant case of approximate sampling that prevents a purely classical non-interactive certification of the output distribution, see Figure 2.

A central ingredient to our proof is a recent result by Valiant and Valiant [36] specifying the optimal sample complexity of certifying a known target distribution P. It can be stated as follows. Fix a preset distance $\epsilon > 0$ up to which we want to certify. Now, suppose we receive samples from a device that samples from an unknown probability distribution Q. Then — for some constants c_1, c_2 — it requires at least

$$c_1 \cdot \max\left\{\frac{1}{\epsilon}, \frac{1}{\epsilon^2}\|P^{-\max}_{-2\epsilon}\|_{2/3}\right\} \tag{1}$$

and at most

$$c_2 \cdot \max\left\{\frac{1}{\epsilon}, \frac{1}{\epsilon^2}\|P^{-\max}_{-\epsilon/16}\|_{2/3}\right\} \tag{2}$$

many samples to distinguish the case $P = Q$ from the case $\|P - Q\|_1 \geq \epsilon$ with high probability. Here $\|\cdot\|_1$ denotes the ℓ_1-norm reflecting the total-variation distance. The central quantity determining the sample complexity of certification is thus the quasi-norm $\|P^{-\max}_{-\epsilon}\|_{2/3}$ which is defined as follows. First, find the truncated distribution $P^{-\max}_{-\epsilon}$ by removing the tail of the target distribution P with weight at most ϵ as well as its largest entry, see Figure 3. Then, take the $\ell_{2/3}$-norm as given by $\|x\|_{2/3} = (\sum_i |x_i|^{2/3})^{3/2}$ for a vector x with entries x_i.

We now proceed in two steps. First, we show lower and upper bounds on the quantity $\|P^{-\max}_{-\epsilon}\|_{2/3}$ in terms of the largest probability p_0 occurring in P and its support $\|P^{-\max}_{-\epsilon}\|_0$ as given by

$$\begin{aligned} p_0^{-\frac{1}{2}}\,(1-\epsilon-p_0)^{3/2} &\leq \|P^{-\max}_{-\epsilon}\|_{2/3} \\ &\leq (1-p_0)\,\|P^{-\max}_{-\epsilon}\|_0^{\frac{1}{2}}. \end{aligned} \tag{3}$$

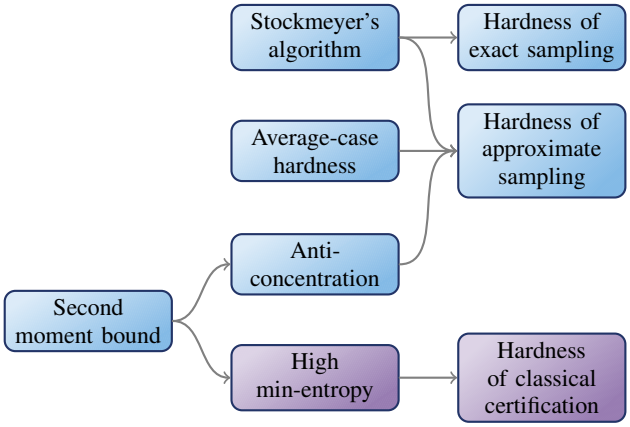

Figure 2. A high level overview of the approximate sampling "quantum supremacy" proofs of Refs. [14–16, 18–20, 34] using Stockmeyer's algorithm [23]. Invoking a worst-case hardness result for the calculation of the output probabilities of some circuit family, Stockmeyer's algorithm can be used to prove the hardness of exact sampling. The key properties of the output probabilities that allows to prove hardness of *approximate* sampling are that computing these probabilities is even hard *on average* and that the distribution anti-concentrates. We show that the same property that is essential to arrive at a hardness result for approximate sampling via anti-concentration also makes it hard to certify from classical samples and a complete description of the target distribution, even with unbounded *computational power*.

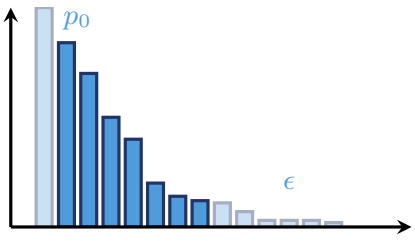

Figure 3. The vector $P^{-\max}_{-\epsilon}$ is obtained from P by removing the largest element p_0 of P as well as the smallest probabilities that accumulate to a total weight bounded by ϵ.

[2] In particular, the certifier is given the value of all target probabilities to arbitrary precision.

[3] See the Supplementary Material [33], Sec. S3 A

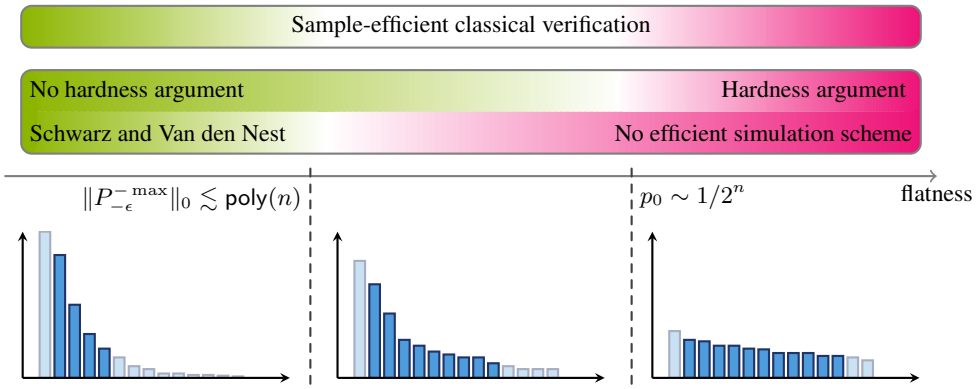

Figure 4. Hardness and certification in terms of the flatness of $P_{-\epsilon}^{-\max}$ for the example of IQP circuits [16, 24] on n qubits as obtained from the present result and the classical simulation algorithm of Schwarz and Van den Nest [35]. There, it is shown that a certain natural family of quantum circuits (including IQP circuits) can be efficiently simulated on a classical computer if the output distribution is essentially concentrated on a polynomial number of outcomes only. In this case, i.e., for $\|P_{-\epsilon}^{-\max}\|_0 \lesssim \mathsf{poly}(n)$, the output distribution is also sample-efficiently certifiable as the bounds (2) and (3) show. Their classical simulation algorithm breaks down if the distribution is essentially spread out over more than polynomially many outcomes, and we even have a rigorous hardness argument by Bremner *et al.* [16] for exponentially flat distributions. Conversely, the number of samples required for certification becomes prohibitively large if the distribution is exponeentially spread out, as measured by the $\ell_{2/3}$-norm (1). Nevertheless, as we illustrate here, there could be "room in the middle" where, for reasonably but not exponentially flat distributions, one may hope to find tasks that are both classically intractable and sample-efficiently certifiable in a device-independent fashion.

Then it follows from Eqs. (1) and (3) that the sample complexity of certifying a distribution P up to a constant total-variation distance ϵ is essentially lower bounded by $1/\sqrt{p_0}$. Hence, if P is exponentially flat in the sense that the largest probability is exponentially small in the problem size (here, the number of particles), ϵ-certification requires exponentially many samples. Conversely, if $P_{-\epsilon/16}^{-\max}$ is supported on polynomially many outcomes only, sample-efficient certification is possible by the converse bound (2).

Second, we connect this result to the output distributions of "quantum supremacy" schemes. In all schemes that rely on the Stockmeyer argument, the problem instances are defined in terms of a unitary U that is *randomly chosen* from some restricted family, e.g., linear optical circuits in the case of boson sampling [14] or random universal circuits [15, 17] in a qubit architecture , captured by a respective measure μ_n on the n-particle unitary group. Specifically, we prove that with high probability over the choice of the random unitary, the distribution over outputs P_U associated with this unitary (induced via $P_U(S) = |\langle S|U|S_0\rangle|^2$, $U \sim \mu_n$) will be exponentially flat in the sense that $H_\infty(P_U) \geq \Omega(n)$. Here, $H_\infty(P) = -\log \max_x p_x$ is the min-entropy of P. We show that, ironically, this follows from an upper bound on the second moments of the output probabilities, which is at the same time a central ingredient in the Stockmeyer hardness argument for *approximate* sampling. Specifically, we prove that with probability at least $1-\delta$ over the choice of U (see Lemma 5 in the Supplementary Material [33])

$$H_\infty(P_U) \geq \frac{1}{2}\left(\log\delta - \log\sum_S \mathbb{E}_{U\sim\mu_n}[P_U(S)^2]\right). \quad (4)$$

Putting everything together we obtain lower bounds on the sample complexity of certification for boson sampling, IQP circuit sampling and random universal circuit sampling with (sufficiently many) n particles. In all of these cases, the sample complexity scales at least as fast as

$$\frac{1}{\epsilon^2}(2^n\delta)^{1/4}, \quad (5)$$

with probability at least $1-\delta$ over the random choice of the unitary.

The upshot is: a key ingredient of the proof of approximate sampling hardness as effected by the random choice of the unitary prohibits sample-efficient certification.

We show that one cannot hope for purely classical, non-interactive, device-independent certification of the proposed quantum sampling problems. This highlights the importance of devising alternative schemes of certification, or improved hardness results for more peaked distributions. We hope to stimulate research in such directions.

A particularly promising avenue of this type of certification has been pioneered by Shepherd and Bremner [24]: By allowing the certifier to choose the classical input to the sampling device rather than drawing it fully at random, it is under some plausible cryptographic assumptions possible to efficiently certify the correct implementation of a quantum sampler from its classical outcomes. This is facilitated by checking a previously hidden bias in the obtained samples and has been achieved for a certain family of IQP circuits [24]. However, in contrast to Ref. [16], there is no approximate sampling hardness result for this family.

Focusing on so-called *relational problems* as opposed to sampling problems, it has been argued via new complexity-theoretic conjectures that the task *HOG* of outputting the *heavy outcomes* of a quantum circuit (those outcomes with probability weight larger than the median of its output distribution) is classically intractable [37]. Clearly, this task

is sample-efficiently checkable via its in-built bias, but still requires exponential classical computation to determine the probabilities of the obtained samples, which are compared to the median.

Taking a pragmatic stance, one can make additional assumptions on the device. In fact, only recently has it been shown [17] that cross-entropy measures [15] provide direct bounds on the total-variation distance provided the entropy of the real distribution is larger than that of the target distribution. One might also be content with weaker notions of certification in total-variation distance such as the certification of a coarse-grained version of the full output distribution [38]. Coarse-graining procedures are practically useful as corroboration schemes when distinguishing against plausible alternative distributions such as the uniform distribution, but of course fail to certify against adversarial distributions on the full sample space. All such approaches yield sample-efficient certificates that require exponential computational effort, rendering them feasible at least for intermediate-scale devices.

Another way to certify a sampling device is the certification of the entire machine from its components. However, such schemes need to make assumptions about the absence of unwanted influences between the components such as crosstalk. In a similar vein, one can make use of implementation details and give the certifier some quantum capabilities such as access to a small quantum computer [39], the ability to manipulate single qubits [40], or to measure the output quantum state in different bases with trusted quantum detectors [27, 41] to devise certificates even in non-iid. settings [42]. In this way, one can partially trade-in the simplicity of sampling schemes for better certifiability.

It is interesting to note the connection of our result with results on classical simulation. Similarly to our findings for the case of certification, Schwarz and Van den Nest [35] find that for certain natural families of quantum circuits (including IQP circuits) classical simulation is possible for highly concentrated distributions, but impossible for flat ones, see Figure 4. This again gives substance to the interesting connection between superior computational power, the flatness of the distribution and the impossibility of an efficient certification.

Curiously, at the same time, the property that prohibits sample-efficient certification is by no means due to the hardness of the distribution. It is merely the flatness of the distribution on an exponential-size sample space as effected by the random choice of the unitary that is required for the approximate hardness argument via Stockmeyer's algorithm and standard conjectures. The uniform distribution on an exponentially large sample space, which is classically efficiently samplable, can also not be sample-efficiently certified.

A further noteworthy connection is that to Shor's algorithm. The output distribution of the quantum part of Shor's algorithm is typically spread out over super-polynomially many outcomes and can hence neither be efficiently simulated via the algorithm of Schwarz and Van den Nest [35], nor certified as we show here. However, after the classical post-processing, the output distribution is strongly concentrated on few outcomes — the factors — from which one can verify the correct working of the algorithm. A certification of the intermediate distribution is simply not necessary to demonstrate a quantum speedup in Shor's algorithm, as its speedup is derived from it solving a problem in NP and not from it sampling close to a hard distribution. This shows that while intermediate steps of a computation might not be certifiable, the final outcome may well be. Whether this is enough to demonstrate a speedup depends on the nature of the hardness argument. In fact, the abovementioned task HOG [37] bears many similarities to factoring and its certifiability from the outcomes of the algorithm.

We hope that our result will stimulate research into new ways of proving hardness of approximate sampling tasks that are more robust than those based on anti-concentration, as well as into devising alternative verification schemes possibly based on mild and physically reasonable assumptions on the sampling device or the verifier.

ACKNOWLEDGEMENTS

We are grateful to Adam Bouland who pointed us to the literature on property testing and thus provided the missing clue for finishing this project. We would like to thank Ashley Montanaro, Bill Fefferman, Tomoyuki Morimae, Martin Schwarz, and Juan Bermejo-Vega for fruitful discussions and Aram Harrow and Anand Natarajan for sharing an early version of their related work. Finally, we would like to thank two anonymous referees for their thorough proof-checking and interesting questions which helped improve the manuscript.

D. H. and J. E. acknowledge support from the ERC (TAQ), the Templeton Foundation, the DFG (EI 519/14-1, EI 519/9-1, EI 519/7-1, CRC 183), and the European Union's Horizon 2020 research and innovation programme under grant agreement No 817482 (PASQuanS). C. G. acknowledges support by the European Union's Marie Skłodowska-Curie Individual Fellowships (IF-EF) programme under GA: 700140 as well as financial support from ARO under contract W911NF-14-1-0098 (Quantum Characterization, Verification, and Validation), the Spanish Ministry MINECO (National Plan 15 Grant: FISICATEAMO No. FIS2016-79508-P, SEVERO OCHOA No. SEV-2015-0522), Fundació Cellex, Generalitat de Catalunya (Grants No. SGR 874 and No. 875, CERCA Programme, AGAUR Grant No. 2017 SGR 1341 and CERCA/Program), ERC (CoG QITBOX and AdG OSYRIS), EU FETPRO QUIC, EU STREP program EQuaM (FP7/2007–2017, Grant No. 323714), and the National Science Centre, Poland-Symfonia Grant No. 2016/20/W/ST4/00314.

[1] J. Preskill, *Quantum computing and the entanglement frontier*, Bull. Am. Phys. Soc. **58** (2013), arXiv:1203.5813.

[2] J. B. Spring, B. J. Metcalf, P. C. Humphreys, W. S. Kolthamm-

mer, X.-M. Jin, M. Barbieri, A. Datta, N. Thomas-Peter, N. K. Langford, D. Kundys, J. C. Gates, B. J. Smith, P. G. R. Smith, and I. A. Walmsley, *Boson sampling on a photonic chip*, Science **339**, 798 (2013).

[3] M. Tillmann, B. Dakić, R. Heilmann, S. Nolte, A. Szameit, and P. Walther, *Experimental boson sampling*, Nature Photonics **7**, 540 (2013).

[4] M. A. Broome, A. Fedrizzi, S. Rahimi-Keshari, J. Dove, S. Aaronson, T. C. Ralph, and A. G. White, *Photonic boson sampling in a tunable circuit*, Science **339**, 794 (2013).

[5] A. Crespi, R. Osellame, R. Ramponi, D. J. Brod, E. F. Galvão, N. Spagnolo, C. Vitelli, E. Maiorino, P. Mataloni, and F. Sciarrino, *Integrated multimode interferometers with arbitrary designs for photonic boson sampling*, Nature Photonics **7**, 545 (2013).

[6] J. Carolan, J. D. A. Meinecke, P. J. Shadbolt, N. J. Russell, N. Ismail, K. Wörhoff, T. Rudolph, M. G. Thompson, J. L. O'Brien, J. C. F. Matthews, and A. Laing, *On the experimental verification of quantum complexity in linear optics*, Nature Photonics **8**, 621 (2014), arXiv:1311.2913.

[7] N. Spagnolo, C. Vitelli, M. Bentivegna, D. J. Brod, A. Crespi, F. Flamini, S. Giacomini, G. Milani, R. Ramponi, P. Mataloni, R. Osellame, E. F. Galvão, and F. Sciarrino, *Experimental validation of photonic boson sampling*, Nature Photonics **8**, 615 (2014), arXiv:1311.1622.

[8] P. Clifford and R. Clifford, The Classical Complexity of Boson Sampling, in *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2018) pp. 146–155, arXiv:1706.01260.

[9] A. Neville, C. Sparrow, R. Clifford, E. Johnston, P. M. Birchall, A. Montanaro, and A. Laing, *Classical boson sampling algorithms with superior performance to near-term experiments*, Nature Physics **13**, 1153 (2017), arXiv:1705.00686.

[10] C. Gogolin, M. Kliesch, L. Aolita, and J. Eisert, *Boson-sampling in the light of sample complexity*, arXiv:1306.3995.

[11] M. J. Bremner, A. Montanaro, and D. J. Shepherd, *Achieving quantum supremacy with sparse and noisy commuting quantum computations*, Quantum **1**, 8 (2017).

[12] M. Oszmaniec and D. J. Brod, *Classical simulation of photonic linear optics with lost particles*, New J. Phys. **20**, 092002 (2018), arXiv:1801.06166.

[13] J. Renema, V. Shchesnovich, and R. Garcia-Patron, *Quantum-to-classical transition in many-body bosonic interference*, (2018), arXiv:1809.01953.

[14] S. Aaronson and A. Arkhipov, *The computational complexity of linear optics*, Th. Comp. **9**, 143 (2013), arXiv:1011.3245.

[15] 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 Physics **14**, 595 (2018), arXiv:1608.00263.

[16] M. J. Bremner, A. Montanaro, and D. J. Shepherd, *Average-case complexity versus approximate simulation of commuting quantum computations*, Phys. Rev. Lett. **117**, 080501 (2016), arXiv:1504.07999.

[17] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, *On the complexity and verification of quantum random circuit sampling*, Nature Physics **15**, 159 (2019), arXiv:1803.04402.

[18] X. Gao, S.-T. Wang, and L.-M. Duan, *Quantum supremacy for simulating a translation-invariant Ising spin model*, Phys. Rev. Lett. **118**, 040502 (2017), arXiv:1607.04947.

[19] J. Bermejo-Vega, D. Hangleiter, M. Schwarz, R. Raussendorf, and J. Eisert, *Architectures for quantum simulation showing a quantum speedup*, Phys. Rev. X **8**, 021010 (2018), arXiv:1703.00466.

[20] D. Hangleiter, J. Bermejo-Vega, M. Schwarz, and J. Eisert, *Anticoncentration theorems for schemes showing a quantum speedup*, Quantum **2**, 65 (2018), arXiv:1706.03786.

[21] A. Bouland, J. F. Fitzsimons, and D. E. Koh, Complexity classification of conjugated Clifford circuits, in *33rd Computational Complexity Conference (CCC 2018)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 102, edited by R. A. Servedio (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2018) pp. 21:1–21:25, arXiv:1709.01805.

[22] M. Yoganathan, R. Jozsa, and S. Strelchuk, *Quantum advantage of unitary Clifford circuits with magic state inputs*, arXiv:1806.03200.

[23] L. Stockmeyer, *On approximation algorithms for #P*, SIAM J. Comput. **14**, 849 (1985).

[24] D. Shepherd and M. J. Bremner, *Temporally unstructured quantum computation*, Proc. Roy. Soc. A **465**, 1413 (2009), arXiv:0809.0847.

[25] S. Aaronson and A. Arkhipov, *BosonSampling is far from uniform*, arXiv:1309.7460.

[26] L. Aolita, C. Gogolin, M. Kliesch, and J. Eisert, *Reliable quantum certification of photonic state preparations*, Nature Communications **6**, 8498 (2015), arXiv:1407.4817.

[27] D. Hangleiter, M. Kliesch, M. Schwarz, and J. Eisert, *Direct certification of a class of quantum simulations*, Quantum Sci. Technol. **2**, 015004 (2017), arXiv:1602.00703.

[28] Y. Nakata, M. Koashi, and M. Murao, *Generating a state t-design by diagonal quantum circuits*, New J. Phys. **16**, 053043 (2014), arXiv:1311.1128.

[29] Y. Nakata, C. Hirche, C. Morgan, and A. Winter, *Unitary 2-designs from random X- and Z-diagonal unitaries*, J. Math. Phys. **58**, 052203 (2017), arXiv:1502.07514.

[30] O. Goldreich, *Introduction to Property Testing* (Cambridge University Press, Cambridge, 2017).

[31] G. Valiant and P. Valiant, *A CLT and tight lower bounds for estimating entropy*, Tech. Rep. 10-179 (2010) eCCC.

[32] M. Walschaers, J. Kuipers, J.-D. Urbina, K. Mayer, M. C. Tichy, K. Richter, and A. Buchleitner, *Statistical benchmark for BosonSampling*, New J. Phys. **18**, 032001 (2016).

[33] See the Supplemental Material at [Insert URL], which contains the mathematical details and proofs of our main result as well as details on the boson sampling, IQP circuit sampling and universal circuit sampling schemes, and further references [43–54].

[34] J. Miller, S. Sanders, and A. Miyake, *Quantum supremacy in constant-time measurement-based computation: A unified architecture for sampling and verification*, Phys. Rev. A **96**, 062320 (2017), arXiv:1703.11002.

[35] M. Schwarz and M. Van den Nest, *Simulating quantum circuits with sparse output distributions*, arXiv:1310.6749.

[36] G. Valiant and P. Valiant, *An automatic inequality prover and instance optimal identity testing*, SIAM J. Comput. **46**, 429 (2017), eCCC, TR13-111.

[37] S. Aaronson and L. Chen, Complexity-Theoretic Foundations of Quantum Supremacy Experiments, in *32nd Computational Complexity Conference (CCC 2017)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 79, edited by R. O'Donnell (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017) pp. 22:1–22:67, arXiv:1612.05903.

[38] S.-T. Wang and L.-M. Duan, *Certification of Boson Sampling Devices with Coarse-Grained Measurements*, (2016), arXiv:1601.02627.

[39] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, *Hamiltonian learning and certification using quantum resources*, Phys. Rev.

Lett. **112**, 190501 (2014), arXiv:1309.0876.

[40] D. Mills, A. Pappa, T. Kapourniotis, and E. Kashefi, *Information theoretically secure hypothesis test for temporally unstructured quantum computation (extended abstract)*, Electron. Proc. Theor. Comput. Sci. **266**, 209 (2018), arXiv:1704.01998.

[41] C. Bădescu, R. O'Donnell, and J. Wright, *Quantum state certification*, arXiv:1708.06002.

[42] Y. Takeuchi and T. Morimae, *Verification of Many-Qubit States*, Phys. Rev. X **8**, 021060 (2018), arXiv:1709.07575.

[43] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, Applied and Numerical Harmonic Analysis (Springer New York, New York, NY, 2013).

[44] T. Morimae, *Hardness of classically sampling the one-clean-qubit model with constant total variation distance error*, Phys. Rev. A **96**, 040302 (2017), arXiv:1704.03640.

[45] B. M. Terhal and D. P. DiVincenzo, *Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin games*, Quant. Inf. Comp. **4**, 134 (2004), arXiv:quant-ph/0205133.

[46] H. Wilming, M. Goihl, I. Roth, and J. Eisert, *Entanglement-ergodic quantum systems equilibrate exponentially well*, arXiv:1802.02052.

[47] B. Fefferman and C. Umans, *The Power of Quantum Fourier Sampling*, (2015), arXiv:1507.05592 [quant-ph].

[48] 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.

[49] A. W. Harrow and R. A. Low, *Random quantum circuits are approximate 2-designs*, Commun. Math. Phys. **291**, 257 (2009), arXiv:0802.1919.

[50] S. Scheel, *Permanents in linear optical networks*, arXiv:quant-ph/0406127.

[51] L. Valiant, *The complexity of computing the permanent*, Theoretical Computer Science **8**, 189 (1979).

[52] R. Lipton, *New directions in testing*, in *Distributed Computing and Cryprography*, Vol. 2 (AMS, 1991) pp. 191–202.

[53] T. Jiang, *How many entries of a typical orthogonal matrix can be approximated by independent normals?* Ann. Probab. **34**, 1497 (2006), arXiv:math/0601457.

[54] N. Ermolova and S. G. Haggman, Simplified bounds for the complementary error function; application to the performance evaluation of signal-processing systems, in *2004 12th European Signal Processing Conference* (2004) pp. 1087–1090.