Dynamical Phase Transitions in Sampling Complexity

Abhinav Deshpande,[1,2] Bill Fefferman,[1,3] Minh C. Tran,[1,2] Michael Foss-Feig,[4,1,2] and Alexey V. Gorshkov[1,2]

[1]*Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park, Maryland 20742, USA*
[2]*Joint Quantum Institute, NIST/University of Maryland, College Park, Maryland 20742, USA*
[3]*Electrical Engineering and Computer Sciences, University of California, Berkeley, California 94720, USA*
[4]*United States Army Research Laboratory, Adelphi, Maryland 20783, USA*

We make the case for studying the complexity of approximately simulating (sampling) quantum systems for reasons beyond that of quantum computational supremacy, such as diagnosing phase transitions. We consider the sampling complexity as a function of time t due to evolution generated by spatially local quadratic bosonic Hamiltonians. We obtain an upper bound on the scaling of t with the number of bosons n for which approximate sampling is classically efficient. We also obtain a lower bound on the scaling of t with n for which any instance of the boson sampling problem reduces to this problem and hence implies that the problem is hard, assuming the conjectures of Aaronson and Arkhipov [*Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (ACM Press, New York, New York, USA, 2011), p. 333]. This establishes a dynamical phase transition in sampling complexity. Further, we show that systems in the Anderson-localized phase are always easy to sample from at arbitrarily long times. We view these results in light of classifying phases of physical systems based on parameters in the Hamiltonian. In doing so, we combine ideas from mathematical physics and computational complexity to gain insight into the behavior of condensed matter, atomic, molecular, and optical systems.

In the quest towards building scalable and fault-tolerant quantum computers, demonstration of a quantum speedup over the best possible classical computers is an important milestone and is termed *quantum computational supremacy* [1]. There are several candidates for tasks where such a speedup could be demonstrated [2–13], where the problem is to simulate a quantum system in the sense of approximate sampling. However, there has also been some debate about the required system size before one can claim quantum computational supremacy, due to improved simulation techniques and algorithms [14–16]. In this Letter, we consider the impact of the field of quantum computational supremacy on other areas of physics and show that studying the complexity of simulating physical systems is useful for understanding phase transitions.

Here we consider the classical complexity of approximate sampling, referred to as "sampling complexity." This is the task of producing samples from a distribution close to the probability distribution occurring in a quantum system upon measurement in a standard basis. This task is a good notion of what it means to simulate physics on a classical computer since it captures how well a computer can mimic an experiment in which one can measure the output at several sites. When we consider sampling complexity as a function of system parameters, the system can be classified as easy in some regimes and hard in some others. Since the designations "easy" and "hard" are exhaustive and there is no smooth way to go from one regime to another, we posit

that this transition from easy to hard happens abruptly, a phenomenon very reminiscent of phase transitions. Just like order parameters are zero on one side of a phase transition and nonzero on another, sampling complexity is different on either side of the transition, and can be used to draw phase boundaries as a function of time and other system parameters. These boundaries can be different from those drawn by more conventional order parameters, signifying new physics in otherwise well-studied systems. Indeed, phase transitions in average-case complexity have been studied both in the classical [17] and quantum regime [18].

In this Letter, we show that transitions in sampling complexity [19] are indicative of physical transitions. We consider a system of n bosons hopping from one site to another on a lattice of m sites and study the sampling complexity as a function of time for an initial product state. We show that it goes from easy to hard as the scaling of evolution time t with the number of bosons n increases, thereby exhibiting a dynamical phase transition [20,21]. We find that the timescale at which the complexity changes is the timescale when interference effects start becoming relevant. We conjecture that in general, this is linked to the Ehrenfest timescale at which quantum effects in a system become considerable [22]. We also show that systems in the Anderson-localized phase are always easy to simulate. We use the Lieb-Robinson bound [23] as an ingredient in our proof of easiness.

Setup.—The model consists of free bosons hopping on a lattice in d dimensions (denoted d-D) with sites labeled by indices i, j. Our results in this Letter can be applied to linear optics as well, with the bosonic sites being replaced by photonic modes. The Hamiltonian is given by $H = \sum_{i,j=1}^{m} J_{ij}(t) a_i^\dagger a_j$, where $a_i^\dagger$ is the creation operator of a boson at the ith site. $J(t)$, which can be time dependent in general, is an $m \times m$ Hermitian matrix that encodes the connectivity of the lattice. One way to show hardness in the boson sampling proposal by Aaronson and Arkhipov (AA) [5] is to generate any linear optical unitary U acting on the bosonic sites, in particular, Haar-random unitaries. Any U can be generated through a free boson Hamiltonian by taking $H = i \log U$ and evolving it for unit time. However, this Hamiltonian can require arbitrarily long-range hops on the lattice in general. Since such long-ranged Hamiltonians may not be realistic, we consider the sampling complexity of a bosonic Hamiltonian with nearest-neighbor hops. J_{ij} is nonzero only if $i = j$ or i and j label adjacent sites. We further restrict J_{ij} to satisfy $|J_{ij}| \leq 1$ in order to set an energy scale. Our proofs remain valid even if we allow the diagonal terms J_{ii} to be unbounded.

One can efficiently solve the equations of motion $i \dot{a}_i^\dagger(t) = [a_i^\dagger(t), H(t)]$ on a classical computer to obtain $a_i^\dagger(t) = \sum_k a_k^\dagger(0) R_{ki}(t)$ for some transformation matrix R. From here onward, we shall take $R(t)$ to be the input to the problem, since it can be determined from the input Hamiltonian H and time t in time $\text{poly}(m, \log t)$.

The m sites in the problem are numbered from 1 to m, and together with n ancilla sites, are arranged in a lattice of side length $(m + n)^{1/d}$ in d dimensions. The initial state has n bosons equally spaced in the lattice as shown in Fig. 1. We take $m = c_1 n^\beta$, where β controls the sparsity of occupied sites in the lattice and can be set to 5 as required for the hardness of boson sampling [5]. The minimum spacing between any two bosons in the initial state is $2L = [(m + n)/n]^{1/d} > c_1^{1/d} n^{(\beta-1)/d}$. The quantity L is an
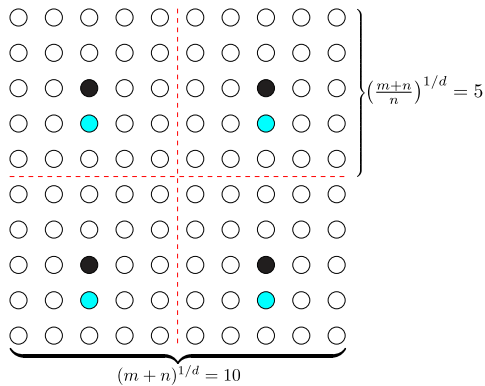

FIG. 1. An example of the initial state in $d = 2$ dimensions. Here $m = 96$, $n = 4$, $\beta = 3$, and $c_1 = 3/2$. The black circles represent sites with a single boson. The cyan circles represent the ancillae.

important length scale in the problem. The ancillae in the lattice, marked in cyan, are not counted as part of the m sites and are present in order to accelerate the time required to construct an arbitrary unitary, which is useful for the hardness result. Their presence does not change the scaling of quantities like L with n.

The input states are described by vectors of the form $r = (r_1, \ldots, r_m)$, specifying the number of bosons on each site, so that $r_1 + \ldots + r_m = n$. Measurement in the boson number basis defines a distribution $\mathcal{D}_U$, which we aim to sample from. The probability of finding an output state $s = (s_1, s_2, \ldots, s_m)$ is given by

$$\text{Pr}_{\mathcal{D}_U}[s] = \frac{1}{r! s!} |\text{Per}(A)|^2, \tag{1}$$

where $r! := r_1! \cdots r_m!$ (with $s!$ defined similarly), $A_{n \times n}$ is a matrix formed by taking s_i copies of the ith column and r_j copies of the jth row of R in any order, and $\text{Per}(A)$ denotes the permanent of A (see Refs. [24,25] for details).

For the particular choice of initial states described in Fig. 1, the task is to sample from a distribution that is close to $\mathcal{D}_U$ in variation distance when given a description of the unitary $R(t)$. We now formalize the notion of efficient sampling.

Definition 1.—Efficient sampler: An efficient sampler is a classical randomized algorithm that takes as input the unitary R_{ij} and outputs a sample s from a distribution $\mathcal{D}_\mathcal{O}$ such that the variation distance between the distributions $\epsilon = \|\mathcal{D}_\mathcal{O} - \mathcal{D}_U\| \leq O(1/\text{poly}(n))$, in runtime $\text{poly}(n)$ (see Refs. [28,29]).

We call the sampling problem *easy* if there exists an efficient sampler for the problem in the stated regime. Conversely, the problem is hard if there cannot be an efficient sampler. Since a negative statement such as the inexistence of an algorithm is difficult to prove, our practical definition of hardness of a sampling problem is if it is at least as hard as boson sampling. This enables us to use the results of AA [5] to claim the hardness of sampling in some regimes. In doing so, our hardness results ultimately rely on the truth of AA's conjectures. One of these conjectures concerns the hardness of additively approximating $|\text{Per}(G)|^2$ for Gaussian-random G with high probability (for which AA give reasonable evidence). The second, more widely believed conjecture is that the polynomial hierarchy, an infinite tower of complexity classes, does not "collapse," i.e., is truly infinite.

We restrict our attention to two special cases where we can show the existence of an efficient sampling algorithm: (i) when the system evolves for a time smaller than the system timescale L/v [where v is the Lieb-Robinson velocity of information spreading in the lattice, defined more precisely in Eq. (2)], and (ii) when there is Anderson localization in the system [30]. These two cases correspond to a promise on the input unitary R. We now state our main results.

Theorem 1.A (easiness of simulation at short times). For $\beta > 1$ and for all dimensions d, the sampling problem is easy for all $t \leq 0.9L/v$, i.e., $\forall\, t \leq c_2 n^{(\beta-1)/d}$ for some constant c_2.

The intuition behind this theorem is that when the time is smaller than the Lieb-Robinson timescale of particle interference L/v, the dynamics is approximately classical (in the sense that the particles are distinguishable).

An important technical achievement of this Letter is to prove rigorously that this intuition is correct. This is done in Lemma 3 by showing that the approximation works.

Theorem 1.B (hardness of simulation at longer times, based on Theorem 3 of Ref. [5]). When $t = \Omega(n^{1+\beta/d})$, the sampling problem is hard in general.

For this result, we show that we can apply any unitary R after the stated time. Therefore, if an efficient sampler exists for this problem, then we have an efficient boson sampler, which cannot exist by our assumption.

The result for the case of Anderson localization comes out as a corollary from Theorem 1.A:

Corollary 2 (easiness of Anderson-localized systems). For Anderson-localized systems in any dimension d, the sampling problem is easy for all times.

The easiness of sampling for Anderson-localized systems is analogous to results showing efficient simulation of localized systems according to various definitions [31–35].

Easiness at short times.—In this section, we prove Theorem 1.A and Corollary 2. First, let us examine the promise we have on the unitary R in both cases. We use the Lieb-Robinson bound [23] on the speed of information propagation in a system. Applying the bound to our Hamiltonian, we get

$$|[a_i(t), a_j^\dagger(0)]| = |R_{ij}(t)| \leq \min\left(1, \exp\left(\frac{vt - \ell_{ij}}{\xi}\right)\right), \quad (2)$$

where ℓ_{ij} is the distance between two sites i and j, v is the upper bound to the velocity of information propagation called the Lieb-Robinson velocity, and ξ is a length scale. Note that the results from Ref. [36] do not apply since we have free bosons here and we work in the single-particle subspace. The Lieb-Robinson velocity is at most $4(1 + 2de)$ [37] when $|J_{ij}| \leq 1$ and $\xi = 1$.

When the Hamiltonian is Anderson localized, the unitary R satisfies the following promise at all times [38]:

$$|R_{ij}| \leq \exp\left(\frac{-\ell_{ij}}{\xi}\right). \quad (3)$$

Here, ξ is the maximum localization length among all eigenvectors. Equation (3) can be viewed as a consequence of a Lieb-Robinson bound with zero velocity [39]. On account of the zero-velocity Lieb-Robinson bound, all results for the time-dependent case can be ported to the Anderson localized case, setting $v = 0$.

We give an algorithm that efficiently samples from the output distribution for short times $t < (0.9L/v) = O(n^{(\beta-1)/d})$, given the promise in Eq. (2). The algorithm outputs a sample from a distribution $\mathcal{D}_{\mathrm{DP}}$, obtained by assuming that the bosons are distinguishable particles, ignoring the effects of interference. The algorithm is described in more detail in Ref. [25].

Analysis.—We prove the correctness of the algorithm by showing that the variation distance between the distributions is upper bounded by an inverse exponential in n. When the bosons are distinguishable, their dynamics is given by a Markov process, described by the matrix $\mathcal{P}_{kl} = |R(t)|_{kl}^2$. The probability of getting an outcome s is given by

$$\Pr_{\mathcal{D}_{\mathrm{DP}}}[s] = \sum_{\sigma} \frac{1}{s!} \mathcal{P}_{\mathrm{in}_1,\mathrm{out}_{\sigma(1)}} \mathcal{P}_{\mathrm{in}_2,\mathrm{out}_{\sigma(2)}} \cdots \mathcal{P}_{\mathrm{in}_n,\mathrm{out}_{\sigma(n)}}, \quad (4)$$

where the sum is over all permutations σ mapping the input bosons to the output ones. In the above equations, in_i is the site index of the ith boson in nondecreasing order in the input and out is defined similarly at the output (see Ref. [25] for an example). We now state a result on how close $\mathcal{D}_{\mathrm{DP}}$ is to the true distribution $\mathcal{D}_U$.

Lemma 3. When $\beta > 1$ and $t \leq 0.9L/v$, the variation distance satisfies $\|\mathcal{D}_{\mathrm{DP}} - \mathcal{D}_U\| = O(\exp[2(vt - L)/\xi + 2(d-1)\log L])$.

This lemma makes intuitive sense: because of the Lieb-Robinson bound Eq. (2) and the fact that the initially occupied sites are separated by a minimum distance $\Theta(L)$ from each other, it takes a time $t = \Theta(L/v)$ for the bosons to start interfering considerably. Therefore, the classical and quantum distributions agree exponentially closely in L when $t \leq 0.9L/v$. For a proof of Lemma 3, see Ref. [25]. Assuming this lemma, we now show Theorem 1.A.

Proof of Theorem 1.A.—Lemma 3 shows that the algorithm samples from a distribution with exponentially small error in n, since $L = \Theta(n^{(\beta-1)/d})$. To complete the proof of Theorem 1.A, we need to show that the runtime of the algorithm is polynomial in n. This is true because the corresponding Markov process of n distinguishable bosons walking on m sites for one step is efficiently simulable: for each of the n particles, we select one among the m sites to walk to, based on the matrix elements of $\mathcal{P}$. This takes time $O(n\,\mathrm{poly}(m)) = O(\mathrm{poly}(n))$ to simulate on a classical computer. □

Hardness at longer times.—If we allow the system to evolve for a longer amount of time, we can use the time-dependent control to effect any arbitrary unitary and implement any boson sampling instance in the system. We can perform phase gates on a site k by setting J_{kk} to be nonzero for a particular time, with the hopping terms and other diagonal terms set to zero. We can apply a nontrivial two-site gate between adjacent sites, for example the balanced beam splitter unitary on the sites 1 and 2,

$U = \frac{1}{\sqrt{2}} \left(\begin{smallmatrix} 1 & -1 \\ 1 & 1 \end{smallmatrix} \right)$, by setting $H = -i(a_1^\dagger a_2 - a_1 a_2^\dagger)$ for time $t = (\pi/4)$. One can also apply arbitrary unitaries using arbitrary on-site control $J_{ii}(t)$ and fixed, time-independent nearest-neighbor hopping $J_{ij}(t) = 1$.

Using the constructions in Refs. [40,41], we can effect any arbitrary $m \times m$ unitary on the m sites with $O(m^2)$ depth from a nontrivial beam splitter and arbitrary single-site phase gates. A construction of AA that employs ancillae to obtain the desired final state rather than applying the full unitary on all the sites can be used to achieve a depth $O(nm^{1/d})$. Each of the n columns of the unitary are implemented in time $O(m^{1/d})$, which corresponds to the timescale set by the Lieb-Robinson velocity and the distance between the furthest two sites in the system (see Ref. [42]).

Proof of Theorem 1.B.—From the above, when $t = \Omega(n^{1+(\beta/d)})$, we see that we can effect any arbitrary unitary. This implies that an efficient sampler for this regime can also be used as an efficient boson sampler, which is widely believed not to exist because of AA's results [5].

Outlook.—We have defined the sampling problem for local Hamiltonian dynamics and given upper and lower bounds for the scaling of time $t(n)$ with the number of bosons n for which the problem is efficiently simulable or hard to classically simulate, respectively. Our results are captured in Fig. 2 which illustrates the complexity phase diagram of the system. For time-independent systems, we observe that boson sampling is classically easy for all times if the system is Anderson localized. In a future work [43], we show that there is a class of static, local Hamiltonians that generate a hard-to-sample output distribution at some time that is not formally infinity. This means that sampling complexity distinguishes Anderson-localized and delocalized systems, which makes it similar to an order parameter that distinguishes different phases.

The case with nonzero Lieb-Robinson velocity [Eq. (2)] shows two regimes of the scaling of t with n where sampling is provably easy or hard. We have shown that sampling is easy when $t \leq t_{\text{easy}} = \Theta(n^{(\beta-1)/d})$ and hard when $t \geq t_{\text{hard}} = \Theta(n^{1+(\beta/d)})$. Since our definitions of easiness and hardness are exhaustive, we argue that there must exist a constant c such that sampling is efficient for $t < \Theta(n^c)$ and hard otherwise, illustrating a dynamical phase transition. Our proof implies that $c \in [(\beta - 1)/d, (\beta + d)/d]$ and we show
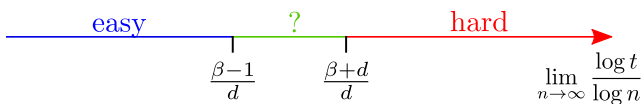

FIG. 2. Complexity phase diagram of free bosons, illustrating that sampling complexity can delineate boundaries of a physical system as a function of system parameters, including time. Our results indicate that $(\beta - 1)/d \leq c \leq (\beta + d)/d$, where c is the transition point for the scaling exponent of t with n.

in future work [43] that the transition is sharp (see Ref. [44]). The transition is between two regimes, one for short times in which the system's dynamics is essentially indistinguishable from classical dynamics; and the other in which quantum mechanical effects dominate to such an extent as to forbid an efficient classical simulation.

This result may be viewed as a generalization of a similar result in Ref. [45], where it was shown that exact boson sampling for depth-4 circuits is hard. The results there are not directly comparable to ours, since Ref. [45] assumes $\beta = 1$, whereas our results need $\beta > 1$ (easiness) and $\beta \geq 2$ or 5 (hardness). The reason we get easiness even after polynomial time is that we deal with approximate sampling, a less stringent notion of simulation. In addition, adapting the hardness proof of Ref. [45] into our setup, we obtain that the system is hard to exactly sample from at timescale $\Omega(8L/v) = \Omega(n^{(\beta-1)/d})$ [43], partially answering what happens in the subregion with the question mark in Fig. 2.

Recent studies [22,46–48] have also studied transitions based on time-dependent features of the out-of-time-ordered correlator. This raises the question of the connection between sampling complexity and scrambling time [49] in quantum many-body systems [50] and fast scramblers like the Sachdev-Ye-Kitaev [46,51–54] models, and black holes [55–58], where one can explore the connection to recent conjectures on complexity in the boundary theory [58–60]. Further, it would be interesting to study sampling complexity in various other physical settings [61], like systems with topological or many-body-localized phases and systems in their ground state to explore the connection with Hamiltonian complexity.

We therefore propose another motivation for considering sampling complexity beyond the field of quantum computational supremacy: complexity provides a natural way to classify phases of matter that is complementary to traditional approaches based on symmetries and topology. This is akin to how the study of entanglement in many-body physics has helped us understand phases of matter [62,63] and characterize thermalization and localization [64,65].

Coming to experimental implementations, platforms such as ultracold atoms in optical lattices [66,67] and superconducting circuits [68,69] are ideal for experimentally studying the transition by comparing the distribution sampled by the algorithm and the experimental distribution. The ingredients required, which have been realized in several groups [70–75], are (1) preparation of the initial state [76,77] of the type shown in Fig. 1 (see Ref. [78]), (2) evolution under a Hamiltonian with either arbitrary time-dependent nearest-neighbor hopping strength or fixed nearest-neighbor hopping strength $J_{ij}(t) = 1$ together with arbitrary time-dependent on-site potential $J_{ii}(t)$ [76,79], and (3) single-site resolved measurement of occupation number of the sites [71,72]. Cold atoms in quantum gas microscopes can be controlled at the single site level,

enabling all three ingredients above. To maintain integrability, we can turn off the Hubbard interaction for the bosonic atoms by tuning to a Feshbach resonance [80–83].

New architectures like optical tweezers [84] are also promising since they allow for deterministic creation of desired initial states and feature tunable interactions in time [85,86]. Similarly, superconducting circuits have been proposed for quantum simulation of quantum walks [87] and the Bose-Hubbard model [88–90]. The gmon qubit architecture, which was used in Ref. [90], naturally allows for time-dependent variation of coupling strengths [91], and raises the prospect of an experimental study of the transition.

Note added in proof.—Recently, we learned of a related work [92], which shows an analogous result for logarithmic time without constraints on the initial state.

---

[1] J. Preskill, arXiv:1203.5813.
[2] B. M. Terhal and D. P. DiVincenzo, Phys. Rev. A **65**, 032325 (2002).
[3] B. M. Terhal and D. P. DiVincenzo, Quantum Inf. Comput. **4**, 134 (2004).
[4] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Proc. R. Soc. Math. Phys. Eng. Sci. **467**, 459 (2010).
[5] S. Aaronson and A. Arkhipov, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (ACM Press, New York, New York, USA, 2011), p. 333.
[6] R. Jozsa and M. V. den Nest, arXiv:1305.6190.
[7] B. Fefferman and C. Umans, in *11th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2016)*, Leibniz International Proceedings in Informatics (LIPIcs) Vol. 61 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2016), pp. 1:1–1:19.
[8] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Phys. Rev. Lett. **117**, 080501 (2016).
[9] E. Farhi and A. W. Harrow, arXiv:1602.07674.
[10] S. Aaronson and L. Chen, arXiv:1612.05903.
[11] B. Fefferman, M. Foss-Feig, and A. V. Gorshkov, Phys. Rev. A **96**, 032324 (2017).
[12] X. Gao, S.-T. Wang, and L.-M. Duan, Phys. Rev. Lett. **118**, 040502 (2017).
[13] A. W. Harrow and A. Montanaro, Nature (London) **549**, 203 (2017).
[14] A. Neville, C. Sparrow, R. Clifford, E. Johnston, P. M. Birchall, A. Montanaro, and A. Laing, Nat. Phys. **13**, 1153 (2017).
[15] P. Clifford and R. Clifford, 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.
[16] A. M. Dalzell, A. W. Harrow, D. E. Koh, and R. L. La Placa, arXiv:1805.05224.
[17] S. Kirkpatrick and B. Selman, Science **264**, 1297 (1994).
[18] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, Quantum Inf. Comput. **10**, 1 (2010).
[19] K. P. Seshadreesan, J. P. Olson, K. R. Motes, P. P. Rohde, and J. P. Dowling, Phys. Rev. A **91**, 022334 (2015).
[20] M. Heyl, A. Polkovnikov, and S. Kehrein, Phys. Rev. Lett. **110**, 135704 (2013).
[21] M. Heyl, Rep. Prog. Phys. **81**, 054001 (2018).
[22] E. B. Rozenbaum, S. Ganeshan, and V. Galitski, Phys. Rev. Lett. **118**, 086801 (2017).
[23] E. H. Lieb and D. W. Robinson, Commun. Math. Phys. **28**, 251 (1972).
[24] S. Aaronson, Proc. R. Soc. Math. Phys. Eng. Sci. **467**, 3393 (2011).
[25] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.121.030501 for a proof of the main lemma and details on the distributions $\mathcal{D}_U$ and $\mathcal{D}_{DP}$, which includes Refs. [26, 27].
[26] M. Jerrum, A. Sinclair, and E. Vigoda, J. ACM **51**, 671 (2004).
[27] NIST Digital Library of Mathematical Functions, edited by F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, and B. V. Saunders. http://dlmf.nist.gov/, Release 1.0.19 of 2018-06-22.
[28] Our easiness results also apply to a stronger notion of approximate sampling, namely that the algorithm samples from an ε-close distribution in runtime poly$(n, 1/\varepsilon)$.
[29] 1. $f(n) = O(g(n))$ if $\forall n \geq n_0, f(n) \leq cg(n)$. 2. $f(n) = \Omega(g(n)) \Leftrightarrow g(n) = O(f(n))$. 3. $f(n) = \Theta(g(n))$ if $f(n) = \Omega(g(n))$ and $f(n) = O(g(n))$.
[30] P. W. Anderson, Phys. Rev. **109**, 1492 (1958).
[31] M. Friesdorf, A. H. Werner, W. Brown, V. B. Scholz, and J. Eisert, Phys. Rev. Lett. **114**, 170505 (2015).
[32] Y. Huang, arXiv:1508.04756.
[33] F. Pollmann, V. Khemani, J. I. Cirac, and S. L. Sondhi, Phys. Rev. B **94**, 041116 (2016).
[34] X. Yu, D. Pekker, and B. K. Clark, Phys. Rev. Lett. **118**, 017201 (2017).
[35] A. Chapman and A. Miyake, arXiv:1704.04405.
[36] J. Eisert and D. Gross, Phys. Rev. Lett. **102**, 240501 (2009).
[37] M. B. Hastings, arXiv:1008.5137.
[38] J. Fröhlich, F. Martinelli, E. Scoppola, and T. Spencer, Commun. Math. Phys. **101**, 21 (1985).
[39] E. Hamza, R. Sims, and G. Stolz, Commun. Math. Phys. **315**, 215 (2012).
[40] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, Phys. Rev. Lett. **73**, 58 (1994).
[41] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walsmley, Optica **3**, 1460 (2016).
[42] AA's construction applied each column of the unitary in $O(\log m)$ depth, whereas we can only apply it in $O(m^{1/d})$ depth because of the spatial locality of the Hamiltonian and the Lieb-Robinson bounds that follow.

[43] A. Deshpande, B. Fefferman, M. C. Tran, M. Foss-Feig, and A. V. Gorshkov (unpublished).

[44] The transition here is sharp if there is a specific time $t_{samp} = \Theta(n^c)$ such that sampling is easy for all $t < t_{samp}$ and hard otherwise. The transition is coarse if sampling is easy for all $t = \Theta(n^c)$ and hard for all t such that $\lim_{n \to \infty}(t/n^c) \to \infty$.

[45] D. J. Brod, Phys. Rev. A 91, 042316 (2015).

[46] S. Banerjee and E. Altman, Phys. Rev. B 95, 134302 (2017).

[47] R.-Q. He and Z.-Y. Lu, Phys. Rev. B 95, 054201 (2017).

[48] S. V. Syzranov, A. V. Gorshkov, and V. M. Galitski, arXiv:1709.09296.

[49] P. Hosur, X.-L. Qi, D. A. Roberts, and B. Yoshida, J. High Energy Phys. 02 (2016) 004; D. A. Roberts and B. Yoshida, J. High Energy Phys. 04 (2017) 121.

[50] Y. Huang, Y.-L. Zhang, and X. Chen, Ann. Phys. (Berlin) 529, 1600318 (2017); R. Fan, P. Zhang, H. Shen, and H. Zhai, Sci. Bull. 62, 707 (2017); Y. Chen, arXiv:1608.02765; B. Swingle and D. Chowdhury, Phys. Rev. B 95, 060201 (2017).

[51] S. Sachdev and J. Ye, Phys. Rev. Lett. 70, 3339 (1993).

[52] A. Kitaev, A simple model of quantum holography (2015). http://online.kitp.ucsb.edu/online/entangled15/.

[53] J. Maldacena and D. Stanford, Phys. Rev. D 94, 106002 (2016).

[54] J. Maldacena, S. H. Shenker, and D. Stanford, J. High Energy Phys. 08 (2016) 106.

[55] Y. Sekino and L. Susskind, J. High Energy Phys. 10 (2008) 065.

[56] N. Lashkari, D. Stanford, M. Hastings, T. Osborne, and P. Hayden, J. High Energy Phys. 04 (2013) 022.

[57] B. Swingle, XRDS Crossroads ACM Mag. Stud. 23, 52 (2016).

[58] A. R. Brown, D. A. Roberts, L. Susskind, B. Swingle, and Y. Zhao, Phys. Rev. Lett. 116, 191301 (2016); Phys. Rev. D 93, 086006 (2016).

[59] S. Aaronson, arXiv:1607.05256.

[60] A. R. Brown and L. Susskind, Phys. Rev. D 97, 086015 (2018).

[61] U. Marzolino and T. Prosen, Phys. Rev. A 93, 032306 (2016).

[62] L. Amico, R. Fazio, A. Osterloh, and V. Vedral, Rev. Mod. Phys. 80, 517 (2008).

[63] B. Zeng, X. Chen, D.-L. Zhou, and X.-G. Wen, arXiv:1508.02595.

[64] J. H. Bardarson, F. Pollmann, and J. E. Moore, Phys. Rev. Lett. 109, 017202 (2012).

[65] Z.-C. Yang, A. Hamma, S. M. Giampaolo, E. R. Mucciolo, and C. Chamon, Phys. Rev. B 96, 020408 (2017).

[66] O. Morsch and M. Oberthaler, Rev. Mod. Phys. 78, 179 (2006).

[67] I. Bloch, J. Dalibard, and S. Nascimbène, Nat. Phys. 8, 267 (2012).

[68] I. Buluta and F. Nori, Science 326, 108 (2009).

[69] M. H. Devoret and R. J. Schoelkopf, Science 339, 1169 (2013).

[70] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, Nature (London) 462, 74 (2009).

[71] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Foelling, L. Pollet, and M. Greiner, Science 329, 547 (2010).

[72] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, Nature (London) 467, 68 (2010).

[73] L. Mazza, D. Rossini, R. Fazio, and M. Endres, New J. Phys. 17, 013015 (2015).

[74] M. Miranda, R. Inoue, Y. Okuyama, A. Nakamoto, and M. Kozuma, Phys. Rev. A 91, 063414 (2015).

[75] R. Yamamoto, J. Kobayashi, K. Kato, T. Kuno, Y. Sakura, and Y. Takahashi, Phys. Rev. A 96, 033610 (2017).

[76] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr, Nature (London) 471, 319 (2011).

[77] Y. Wang, X. Zhang, T. A. Corcovilos, A. Kumar, and D. S. Weiss, Phys. Rev. Lett. 115, 043003 (2015).

[78] All of our arguments go through even if the initial state is not regular as in Fig. 1 but still has the property that the initially occupied sites are separated by a minimum distance of $\Theta(L)$ from each other.

[79] T. Fukuhara, A. Kantian, M. Endres, M. Cheneau, P. Schauß, S. Hild, D. Bellem, U. Schollwöck, T. Giamarchi, C. Gross, I. Bloch, and S. Kuhr, Nat. Phys. 9, 235 (2013).

[80] S. Paul, P. R. Johnson, and E. Tiesinga, Phys. Rev. A 93, 043616 (2016).

[81] K. Xu, T. Mukaiyama, J. R. Abo-Shaeer, J. K. Chin, D. E. Miller, and W. Ketterle, Phys. Rev. Lett. 91, 210402 (2003).

[82] J. Herbig, Science 301, 1510 (2003).

[83] S. Dürr, T. Volz, A. Marte, and G. Rempe, Phys. Rev. Lett. 92, 020406 (2004).

[84] A. M. Kaufman, B. J. Lester, C. M. Reynolds, M. L. Wall, M. Foss-Feig, K. R. A. Hazzard, A. M. Rey, and C. A. Regal, Science 345, 306 (2014).

[85] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, Science 354, 1024 (2016).

[86] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Nature (London) 551, 579 (2017).

[87] E. Flurin, V. V. Ramasesh, S. Hacohen-Gourgy, L. S. Martin, N. Y. Yao, and I. Siddiqi, Phys. Rev. X 7, 031023 (2017).

[88] S. Hacohen-Gourgy, V. V. Ramasesh, C. De Grandi, I. Siddiqi, and S. M. Girvin, Phys. Rev. Lett. 115, 240501 (2015).

[89] X.-H. Deng, C.-Y. Lai, and C.-C. Chien, Phys. Rev. B 93, 054116 (2016).

[90] P. Roushan, C. Neill, J. Tangpanitanon, V. M. Bastidas, A. Megrant, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Fowler, B. Foxen, M. Giustina, E. Jeffrey, J. Kelly, E. Lucero, J. Mutus, M. Neeley, C. Quintana, D. Sank et al., Science 358, 1175 (2017).

[91] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White et al., Phys. Rev. Lett. 113, 220502 (2014).

[92] G. Muraleedharan, A. Miyake, and I. H. Deutsch, arXiv:1805.01858.