Exact sampling hardness of Ising spin models

B. Fefferman,[1] M. Foss-Feig,[2,3,1] and A. V. Gorshkov[3,1]

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

We study the complexity of classically sampling from the output distribution of an Ising spin model, which can be implemented naturally in a variety of atomic, molecular, and optical systems. In particular, we construct a specific example of an Ising Hamiltonian that—after time evolution starting from a trivial initial state—produces a particular output configuration with probability very nearly proportional to the square of the permanent of a matrix with arbitrary integer entries. In a similar spirit to BosonSampling, the ability to sample classically from the probability distribution induced by time evolution under this Hamiltonian would imply unlikely complexity theoretic consequences, suggesting that the dynamics of such a spin model cannot be efficiently simulated with a classical computer. Physical Ising spin systems capable of achieving problem-size instances (i.e. qubit numbers) large enough so that classical sampling of the output distribution is classically difficult *in practice* may be achievable in the near future. Unlike BosonSampling, our current results only imply hardness of *exact* classical sampling, leaving open the important question of whether a much stronger *approximate*-sampling hardness result holds in this context. As referenced in a recent paper of Bouland, Mancinska, and Zhang [1], our result completes the sampling hardness classification of two-qubit commuting Hamiltonians.

I. INTRODUCTION

It is often taken for granted that quantum computers can efficiently perform certain computational tasks that classical computers cannot. But finding a quantum task that, on the one hand, admits compelling complexity-theoretic arguments against efficient classical simulation, and on the other hand admits experimental demonstration with technology that is feasible in the near future, remains an important and challenging task in the field of quantum information science. An extremely exciting line of work, starting with results of Terhal and DiVincenzo and Bremner, Jozsa, and Shepherd, has shown that quantum computers are capable of sampling from distributions that cannot be sampled exactly by randomized classical algorithms [2, 3]. The BosonSampling protocol [4], proposed by Aaronson and Arkhipov, gives a hardness of sampling result that may be within reach for near-term quantum experiments. The basic idea is to send photons through a network of linear optical devices, arranged in such a way that the probabilities of typical output configurations of the photons are proportional to the squares of permanents of matrices with independent and Gaussian-distributed random entries. Given reasonable assumptions about the hardness of computing permanents of such matrices, the ability to efficiently classically sample from any distribution even close (in total variation distance) to this distribution would imply extremely unlikely complexity theoretic consequences.

A number of proof-of-principle experiments implementing BosonSampling have already been carried out [5–8]. However, a remaining bottleneck to producing an experimentally convincing demonstration of BosonSampling is the technical difficulty of building linear-optical systems that are large enough and clean enough to realize BosonSampling instances for which classical sampling is actually difficult. By comparison, state preparation and readout of individual spins can be done with high fidelity and relative ease, and the ability to massively parallelize spin-spin interactions between large numbers of qubits is reasonably sophisticated; experiments have successfully implemented some simple instances of the Ising model with system sizes ranging from tens [9] to many hundreds of spins [10]. Moreover, recent developments in ion-trapping experiments raise the exciting prospect of implementing *arbitrary* Ising interaction graphs in systems of (potentially) many tens of trapped ions [11]. For this reason, finding results analogous to BosonSampling for simple spin models is highly desirable, and potentially affords a simpler route towards the experimental demonstration of an efficient quantum task that, under extremely plausible assumptions about classical complexity theory, cannot be efficiently performed by a classical system [3, 12].

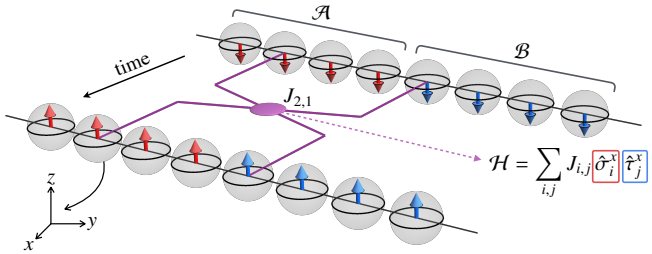

FIG. 1. Schematic of the model: (a) Spins in sublattice $\mathcal{A}$ (red) are coupled to spins in sublattice $\mathcal{B}$ (blue) via Ising couplings $\hat{\sigma}_i^x \hat{\tau}_j^x$ and all of them start off in $|\downarrow\rangle$. To lowest order in time, the matrix element of the time evolution operator between an initial state with all spins intialized in $|\downarrow\rangle$ and a final state with all qubits in $|\uparrow\rangle$ receives contributions in which each spin is flipped precisely once (one such contributing term, between the spin on the second site of $\mathcal{A}$ and the first spin of $\mathcal{B}$, is shown.

Our goal in this manuscript is to show that the dynamics of

an experimentally implementable commuting spin model—the Ising model with no transverse field—can induce an output distribution over the spin states that is hard to sample from classically. The general strategy, which will be elaborated on below, is to divide a set of Ising spins into two mutually interacting registers, each having N spins (see Fig. 1). The N spins in the first and second register can be placed in correspondence with the N row and column labels, respectively, of an $N\times N$ matrix J; each of the N^2 pairwise Ising couplings $J_{i,j}$ between a spin (i) in one register and a spin (j) in the other is a matrix element of J. By initializing the system in a spatially homogeneous product state and then letting it evolve under Ising interactions for a short time, it can be shown that a single probability of the output distribution induced by measurement is proportional to the square of the permanent of J, plus an $o(1)$ correction. This is enough, using a tool known as "Stockmeyer counting" [13], to imply a hardness of "exact sampling" result: no efficient classical randomized algorithm can sample from *exactly* this distribution, under a ubiquitous hardness assumption (namely, that the Polynomial-time Hierarchy does not collapse). Note that in a recent paper [12], BosonSampling was directly generalized to the context of spin Hamiltonians. However, our work encounters the permanent in a fundamentally different way; an important difference is that our results do not rely on a "diluteness criterion", and thus N is set by—as opposed to much less than—the number of physical qubits. Much like other "exact sampling" results, our result also demonstrates hardness to classically sample from any distribution in which all probabilities are within a constant multiplicative factor of the ideal quantum distribution. However, unlike BosonSampling, a recent proposal of Bremner, Montanaro and Shepherd (sometimes called "IQP" sampling), and Quantum Fourier Sampling, it is not yet clear whether the distributions we consider can be used to show an "approximate-sampling" hardness result [3, 4, 14]. This would show something far stronger: there is no classical algorithm that can sample from any distribution inverse polynomial in total variation distance from the ideal quantum distribution.

II. THE MODEL

The model we consider consists of $2N$ spin-1/2 particles, which we divide into two sublattices of N spins each, denoted $\mathcal{A}$ and $\mathcal{B}$ (blue and red spins in Fig. 1). We consider quench dynamics under an Ising Hamiltonian with exclusively two-body *inter-sublattice* interactions (but no interactions within either sublattice), which can take arbitrary integer values,

$$\mathcal{H} = \sum_{i,j} J_{i,j}\hat{\sigma}_i^x\hat{\tau}_j^x. \tag{1}$$

Here, Pauli operators $\hat{\sigma}$ act on the spins of sublattice $\mathcal{A}$, while Pauli operators $\hat{\tau}$ act on the spins of sublattice $\mathcal{B}$. These spins could be, for example, two subsets of ions in a Paul trap, where the $|\downarrow\rangle$ and $|\uparrow\rangle$ are, respectively, the electronic ground state and some long-lived metastable state (in general either an excited hyperfine level of the electronic ground-state manifold or a dipole-forbidden optical excitation). The Ising interactions can then be implemented via a spatially-structured Molmer-Sørensen interaction [11, 15, 16].

We consider a quantum quench in which the system is initialized at time $t = 0$ with all of the spins (in both registers) in the spin-down state along the z-direction,

$$|\psi(0)\rangle = \bigotimes_{i\in\mathcal{A}} |\downarrow\rangle_i \bigotimes_{j\in\mathcal{B}} |\downarrow\rangle_j . \tag{2}$$

We then allow the system to evolve under the Hamiltonian in Eq. (1) for a time t.

III. OUTPUT DISTRIBUTION

After evolution for a time t under the action of $\mathcal{H}$, measurement in the z basis samples from the induced probability distribution

$$\begin{aligned} &P_t(\sigma_1,\ldots,\sigma_N,\tau_1,\ldots,\tau_N) \\ &\quad = |\langle\sigma_1,\ldots,\sigma_N,\tau_1,\ldots,\tau_N|\exp(-i\mathcal{H}t)|\downarrow,\ldots,\downarrow\rangle|^2, \end{aligned} \tag{3}$$

where $\sigma_j,\tau_j = \downarrow,\uparrow$. We are interested in just one such probability,

$$P_t \equiv P_t(\uparrow,\ldots,\uparrow) = |\langle\uparrow,\ldots,\uparrow|\exp(-it\mathcal{H})|\downarrow,\ldots,\downarrow\rangle|^2 \equiv |M_t|^2,$$

to end in the state with all spins in both registers pointing up. By writing an individual term in the Hamiltonian as

$$\hat{\sigma}_i^x\hat{\tau}_j^x = \hat{\sigma}_i^+\hat{\tau}_j^+ + \hat{\sigma}_i^+\hat{\tau}_j^- + \hat{\sigma}_i^-\hat{\tau}_j^+ + \hat{\sigma}_i^-\hat{\tau}_j^-, \tag{4}$$

it is straightforward to see that repeated applications of $\mathcal{H}$, and thus time evolution, generates population in all possible spin states in the z basis. Expanding $e^{-i\mathcal{H}t}$ as a power series in time, the lowest-order in time non-vanishing contribution to the matrix element $M_t = \langle\uparrow,\ldots,\uparrow|\exp(-it\mathcal{H})|\downarrow,\ldots,\downarrow\rangle$ arises at order t^N, because every spin needs to be flipped at least once. The contributing terms contain exactly N powers of operators $\hat{\sigma}_i^+\hat{\tau}_j^+$, with no repetitions of the indices i and j, so that each qubit gets flipped from $|\downarrow\rangle$ to $|\uparrow\rangle$ exactly one time; see Fig. 2 for an illustration of such a term for $N = 3$. It is straightforward to show that, to order t^N, the matrix element M_t is given by

$$\begin{aligned} M_t &= \frac{(-it)^N}{N!}\times N!\sum_{\sigma}\prod_{j=1}^{N} J_{\sigma(j),j} + O(t^{N+2}) \\ &= (-it)^N\mathrm{Per}(J) + O(t^{N+2}), \end{aligned} \tag{5}$$

where the summation is over all permutations σ of the integers $i = 1,\ldots,N$. As a result, and defining $\mathscr{P} = |\mathrm{Per}(J)|^2$, we have

$$P_t = t^{2N}\left(\mathscr{P} + O(t^2)\right). \tag{6}$$

We next aim to place a constraint on how t must scale with N in order to ensure that the $O(t^2)$ additive error to the permanent is $o(1)$ with respect to the system size N.

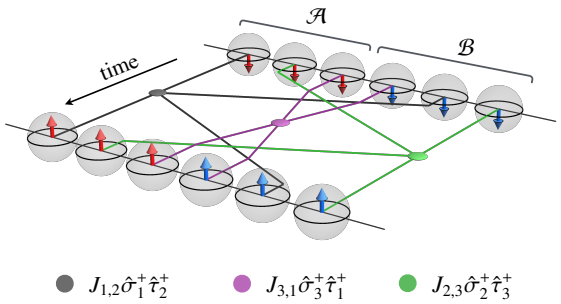

FIG. 2. Example of a single term contributing to the matrix element M_t at lowest order in time (t^N, here with $N = 3$). Here, all spins are flipped from down to up by a particular pairing off of the spins between the $\mathcal{A}$ and $\mathcal{B}$ sublattices. The depicted process contributes a term $(J_{1,2} \times J_{3,1} \times J_{2,3}) \times (t^3/3!)$ to M_t. The set of all possible ways to pair the spins in sublattice $\mathcal{A}$ with the spins in sublattice $\mathcal{B}$ is in one-to-one correspondence with terms in the permanent of the matrix $J_{i,j}$, and thus M_t is proportional to this permanent.

IV. HIGHER ORDERS IN TIME

As discussed above, the lowest-order in time contribution to the matrix element M_t comes at order N. It is not hard to see that all other contributing terms occur at order m such that $m - N$ is a positive *even* integer. In particular, take N_{+-} to be the number of times an operator $\hat{\sigma}_i^+\hat{\tau}_j^-$ occurs inside the matrix element, and similarly for N_{-+}, N_{++}, and N_{--}, such that $N_{++} + N_{--} + N_{+-} + N_{-+} = m$. Since we need to flip the same number of qubits in both registers, we must have $N_{+-} = N_{-+}$. Also, the total number of flipped qubits is equal to $2(N_{++} - N_{--})$, and since all qubits need to be flipped, we have $N_{++} - N_{--} = N$. Now, defining $p(n)$ to be the parity of the integer n, we have

$$\begin{aligned} p(m) &= p(N_{++} + N_{--} + 2N_{+-}) \\ &= p(N_{++} + N_{--}) \\ &= p(N_{++} - N_{--}) \\ &= P(N), \end{aligned} \tag{7}$$

which shows that $m - N$ is an even integer. The matrix element in question can therefore be expanded as

$$M_t = \sum_{\alpha=0}^{\infty} \langle\uparrow,\ldots,\uparrow| \frac{(-it\mathcal{H})^{N+2\alpha}}{(N+2\alpha)!} |\downarrow,\ldots,\downarrow\rangle \equiv \sum_{\alpha=0}^{\infty} M_t^{(\alpha)}, \tag{8}$$

and from above we have

$$M_t^{(0)} = (-it)^N \mathrm{Per}(J). \tag{9}$$

Defining $\delta M_t = \sum_{\alpha=1}^{\infty} M_t^{(\alpha)}$, such that $M_t = M_t^{(0)} + \delta M_t$, we can write

$$\begin{aligned} P_t &= |M_t^{(0)}|^2 + 2\Re[M_t^{(0)}\delta M_t] + |\delta M_t|^2 \\ &= t^{2N}\Big(\mathscr{P} + \eta_t\Big), \end{aligned} \tag{10}$$

where

$$\begin{aligned} \eta_t &\equiv (2\Re[M_t^{(0)}\delta M_t] + |\delta M_t|^2)/t^{2N} \\ &\le |\delta M_t|(2|M_t^{(0)}| + |\delta M_t|)/t^{2N}. \end{aligned} \tag{11}$$

For notational simplicity, here we will assume that the entries of J are drawn from the set $\{-1, 0, 1\}$; note that nothing about our argument would change if arbitrary integers would use, except that the time t would be rescaled in the bounds below by $\max(J_{i,j})$. Using $\langle\uparrow,\ldots,\uparrow|\mathcal{H}^m|\downarrow,\ldots,\downarrow\rangle \le N^{2m}\|\hat{\sigma}^x\|^{2m} = N^{2m}$, $M_t^{(\alpha)}$ can be bounded as $|M_t^{(\alpha)}| \le (N^2t)^{N+2\alpha}/(N+2\alpha)!$. Therefore,

$$|M_t^{(0)}| \le \frac{(N^2t)^N}{N!}, \tag{12}$$

$$|\delta M_t| \le \frac{(N^2t)^N}{N!}\sum_{\alpha=1}^{\infty}(N^4t^2)^\alpha \le 2\frac{(N^2t)^N}{N!}(N^4t^2). \tag{13}$$

The final inequality in Eq. (13) is valid for $t^2 \le 1/(2N^4)$, because $0 \le \sum_{\alpha=1}^{\infty} x^\alpha \le 2x$ whenever $0 \le x \le 1/2$. Plugging Eqs. (12,13) into Eq. (11) leads to

$$\eta_t \le 4N^4t^2\frac{N^{4N}}{(N!)^2}\left(1 + N^4t^2\right) \tag{14}$$

$$\le 6N^4t^2\frac{N^{4N}}{(N!)^2} \le t^2\mathrm{poly(N)}e^{2N(\ln N+1)}, \tag{15}$$

with the final inequality obtained by Stirling's approximation. It follows immediately that $\eta_t = o(1)$ is guaranteed as long as

$$t = o(e^{-2N\ln N}). \tag{16}$$

V. HARDNESS OF SAMPLING

Here we prove our main theorem, establishing a very unlikely complexity theoretic consequence which would arise naturally from the presumed existence of a classical algorithm that samples exactly from the output distribution described in the prior sections. Similar arguments to the one sketched here are implicit in other works on quantum hardness of sampling results starting with the BosonSampling proposal [4].

We first begin with a very brief overview of the computational complexity theoretic components necessary to understand this hardness of sampling result. Computing exactly the permanent of an $N \times N$ matrix X with integer entries is as hard as computing the number of satisfying assignments to a boolean formula. We therefore say it is a #P-hard problem, as established by Valiant [17]. When X has nonnegative integer entries this problem is also in #P.

For our purposes, we will be interested in the complexity of computing *multiplicative estimates* to the permanent. We say an algorithm $\mathscr{A}$ efficiently computes a multiplicative estimate to a function f if, given input x, the output of $\mathscr{A}$ is

within a $1 \pm \epsilon$ multiplicative factor of $f(x)$ in time polynomial in N and $1/\epsilon$. A famous result of Jerrum, Sinclair and Vigoda gives an algorithm for efficiently computing a multiplicative estimate to the permanent of a matrix with *nonnegative* entries [18]. On the other hand, it can be shown using a binary search and padding argument that computing such an estimate to the permanent (or even the square of the permanent) of a matrix with general integer entries is in fact #P-hard (see e.g., [4, 19]). Therefore computing these estimates are as hard as computing the permanent exactly. How powerful is #P? We know from Toda's Theorem that any problem in the Polynomial-time hierarchy, or PH, can be solved using the ability to solve a #P-hard problem [20]. Being a bit more formal, Toda's theorem tells us that $\mathsf{PH} \subseteq \mathsf{P}^{\#\mathsf{P}}$.

Now, for any $N \times N$ matrix X define $\mathcal{D}_X$ to be the outcome distribution from Section III that arises from starting in the $|\downarrow, \ldots, \downarrow\rangle$ state, evolving for a particular time t under the action of the Hamiltonian from Eq. (1) with coupling constants $J_{i,j}$ set to the entries of X, and measuring in the z basis. As shown in Sections III and IV, the probability of observing the $|\uparrow, \ldots, \uparrow\rangle$ outcome at time t is proportional to the square of the permanent of X plus an $o(1)$ correction, provided that t is chosen to be $o(e^{-2N \ln N})$. Notice that this probability is exponentially small. Therefore, to get any reasonable estimate by repeated sampling we would need an exponential number of samples. Indeed, *this does not imply* an efficient quantum algorithm for computing the permanent. Nonetheless, we can use the fact that a single exponentially small amplitude is proportional to the permanent to argue about the classical intractability of sampling from this distribution.

Suppose we have an efficient classical sampler which samples from the same distribution. We define this to be an efficient randomized algorithm that takes as input an $N \times N$ integer matrix X and outputs a sample from the distribution $\mathcal{D}_X$. A classic result of Stockmeyer gives an algorithm for computing a multiplicative estimate to the probability of any given outcome of an efficient classical sampler in the third level of the PH, or Σ_3 [13]. Using this result, together with the presumed existence of an efficient classical sampler for our quantum distribution, we can compute a multiplicative estimate to the square of the permanent of an arbitrary integer matrix in the third level of the PH. As mentioned above, this is a #P-hard problem. This tells us we can solve any problem in #P in the third level of the Polynomial-time hierarchy, or formally, that $\mathsf{P}^{\#\mathsf{P}} \subseteq \Sigma_3$. Combining this with Toda's theorem, we have that $\mathsf{PH} \subseteq \mathsf{P}^{\#\mathsf{P}} \subseteq \Sigma_3$, and so the entire Polynomial-time Hierarchy collapses to the third level, as claimed. Therefore, it is very unlikely that an efficient classical sampler for the distribution with probabilities given by Equation 3 exists.

VI. DISCUSSION AND IMPLICATIONS

These results extend several key ideas of BosonSampling to the context of spin dynamics under Ising spin Hamiltonians. Just like non-interacting bosons, the Ising model without a transverse field is often viewed—from the perspective of many-body quantum physics—to be trivial, since it can be trivially diagonalized. However, just as with non-interacting bosons, this point of view stems from a restricted notion of what it means to "simulate" a quantum system. As in the case of non-interacting bosons, it is indeed classically efficient to compute low-order correlation functions of operators in the model we study [21, 22], but sampling from the output distribution is simply a more general (and less trivial) task.

Another interesting motivation for our result comes from the desire to classify all two-qubit commuting Hamiltonians. Suppose we start in a computational basis state of n qubits, and can apply a fixed two-qubit Hamiltonian to any pair of qubits. A recent result of Bouland, Mancinska, and Zhang gave a hardness of sampling classification for this model [1]. They prove, in all cases except the one we consider (in which the two qubit Hamiltonian is $X \otimes X$) that the corresponding sampling task is classically hard, as long as the commuting Hamiltonian is capable of generating entanglement from a computational basis state. Otherwise, the output is in a product state and clearly classically simulable. Thus our hardness result completes the sampling hardness classification of the complete class of two-qubit commuting Hamiltonians (see their paper for additional details [1]).

VII. ACKNOWLEDGMENTS

We thank A. Deshpande and A. Bouland for helpful discussions. We also thank A. Bouland, Laura Mancinska and Xue Zhang for sharing an early version of their results. A.V.G. acknowledge support by ARL CDQI, ARO MURI, NSF QIS, ARO, NSF PFC at JQI, and AFOSR. This material is based upon work supported by, or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number 025989-001.

[1] A. Bouland, L. Mancinska, and X. Zhang, in *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan* (2016) pp. 28:1–28:33.
[2] B. M. Terhal and D. P. DiVincenzo, Quantum Information and Computation **4**, 134 (2004).
[3] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **467**, 459 (2010).
[4] S. Aaronson and A. Arkhipov, Theory of Computing **9**, 143 (2013).
[5] M. A. Broome, A. Fedrizzi, S. Rahimi-Keshari, J. Dove, S. Aaronson, T. C. Ralph, and A. G. White, Science **339**, 794 (2013).
[6] J. B. Spring, B. J. Metcalf, P. C. Humphreys, W. S. Kolthammer, 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, Science **339**, 798 (2013).

[7] M. Tillmann, B. Dakic, R. Heilmann, S. Nolte, A. Szameit, and P. Walther, Nat. Photon **7**, 540 (2013).

[8] A. Crespi, R. Osellame, R. Ramponi, D. J. Brod, E. F. Galvao, N. Spagnolo, C. Vitelli, E. Maiorino, P. Mataloni, and F. Sciarrino, Nat. Photon **7**, 545 (2013).

[9] P. Richerme, Z.-X. Gong, A. Lee, C. Senko, J. Smith, M. Foss-Feig, S. Michalakis, A. V. Gorshkov, and C. Monroe, Nature **511**, 198 (2014).

[10] J. G. Bohnet, B. C. Sawyer, J. W. Britton, M. L. Wall, A. M. Rey, M. Foss-Feig, and J. J. Bollinger, Science **352**, 1297 (2016).

[11] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Nature **536**, 63 (2016).

[12] B. Peropadre, A. Aspuru-Guzik, and J. J. Garci-Ripoll, arXiv:1509.02703 (2015).

[13] L. J. Stockmeyer, SIAM J. Comput. **14**, 849 (1985).

[14] 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, edited by A. Broadbent (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2016) pp. 1:1–1:19.

[15] K. Mølmer and A. Sørensen, Phys. Rev. Lett. **82**, 1835 (1999).

[16] S. Korenblit, D. Kafri, W. C. Campbell, R. Islam, E. E. Edwards, Z.-X. Gong, G.-D. Lin, L.-M. Duan, J. Kim, K. Kim, and C. Monroe, New J. Phys. **14**, 095024 (2012).

[17] L. G. Valiant, Theoretical Computer Science **8**, 189 (1979).

[18] M. Jerrum, A. Sinclair, and E. Vigoda, J. ACM **51**, 671 (2004).

[19] S. Aaronson, in *Proc. R. Soc. A*, Vol. 467 (The Royal Society, 2011) pp. 3393–3405.

[20] S. Toda, SIAM J. Comput. **20**, 865 (1991).

[21] M. van den Worm, B. C. Sawyer, J. J. Bollinger, and M. Kastner, New J. Phys. **15**, 083007 (2013).

[22] M. Foss-Feig, K. R. A. Hazzard, J. J. Bollinger, and A. M. Rey, Phys. Rev. A **87**, 042101 (2013).