

PHYSICS

Quantum computational advantage via high-dimensional Gaussian boson sampling

Abhinav Deshpande^{1,2,3†}, Arthur Mehta^{4,5†}, Trevor Vincent⁴, Nicolás Quesada^{4,6}, Marcel Hinsche⁷, Marios Ioannou⁷, Lars Madsen⁴, Jonathan Lavoie⁴, Haoyu Qi⁴, Jens Eisert^{7,8,9}, Dominik Hangleiter^{1,7}, Bill Fefferman¹⁰, Ish Dhand^{11*}

Photonics is a promising platform for demonstrating a quantum computational advantage (QCA) by outperforming the most powerful classical supercomputers on a well-defined computational task. Despite this promise, existing proposals and demonstrations face challenges. Experimentally, current implementations of Gaussian boson sampling (GBS) lack programmability or have prohibitive loss rates. Theoretically, there is a comparative lack of rigorous evidence for the classical hardness of GBS. In this work, we make progress in improving both the theoretical evidence and experimental prospects. We provide evidence for the hardness of GBS, comparable to the strongest theoretical proposals for QCA. We also propose a QCA architecture we call high-dimensional GBS, which is programmable and can be implemented with low loss using few optical components. We show that particular algorithms for simulating GBS are outperformed by high-dimensional GBS experiments at modest system sizes. This work thus opens the path to demonstrating QCA with programmable photonic processors.

INTRODUCTION

We are arriving at an exciting era for quantum computing in which quantum experiments are pushing the limits of what is efficiently computable by the most powerful classical supercomputers. The first major goal for this era is the demonstration of a scalable quantum advantage or quantum computational advantage (QCA) (also termed “quantum computational supremacy”) over classical computers. QCA is important as a probe of the foundations of computer science, where it can be seen as an experimental violation of the extended Church-Turing thesis, and it also serves as an important benchmarking tool for comparing near-term experiments on different platforms in a fair and consistent manner. The recent groundbreaking demonstrations of QCA (1, 2) constitute the first notable experimental evidence against the extended Church-Turing thesis.

Notwithstanding, multiple potential loopholes have been pointed out (3–5). QCA will not be marked by a single isolated experiment but rather will be established by gradually improving and scaling up “high complexity” experiments run over the course of many years, which improving classical algorithms will try to simulate. Our confidence that we have arrived in this new era will grow as multiple experiments, performed in different physical architectures, independently reach this conclusion in a comparable fashion. In this way, the goal may be seen as being analogous to Bell inequality

violations, which were originally conducted in landmark experiments starting in the 1970s performed on a variety of different platforms but only much later were loopholes closed.

In the same vein, theoretical results about QCA justify the classical hardness of simulating an experiment in the realm of asymptotically large system sizes. To interpret conclusions from experiments performed at a fixed system size, we should also consider the concrete cost of simulating these finite-size experiments using known algorithms. The two lines of inquiry are complementary to each other and support each other in a claim that any experiment is likely impossible to feasibly simulate with current hardware.

Among different approaches to demonstrating QCA (1, 2, 6, 7), photonics provides a promising path as it enables room temperature operation, fast gate speeds, and notable potential for scalability (8, 9). Arguably, the most feasible approach to demonstrating QCA with photonics is to perform the Gaussian boson sampling (GBS) protocol (10, 11). This protocol is at the heart of the recent QCA demonstration performed by a team from University of Science and Technology of China (USTC) (2), which used a GBS device with 100 modes and an average of around 45 photons. However, GBS has several important limitations. On the experimental side, current implementations of GBS either lack programmability (2) or have high loss rates, which could render the system classically simulable (12, 13). In addition, from a theoretical standpoint, there is a comparative lack of complexity-theoretic evidence for the hardness of GBS (5) and an understanding of the classical runtime of concrete algorithms to simulate GBS instances.

In this work, we aim to address these challenges. We close important theoretical loopholes in the hardness argument for GBS and provide evidence for the hardness of classically simulating GBS even in the presence of loss. We moreover propose a new, programmable architecture for GBS that promises better robustness to loss in a near-term experiment and an asymptotic quantum speedup over classical algorithms. In addition, our proposed architecture is designed so that it is outside of known regimes where current algorithms can simulate finite-size GBS instances in feasible time, as we show through numerical benchmarking.

¹Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park, MD 20742, USA. ²Joint Quantum Institute, NIST/University of Maryland, College Park, MD 20742, USA. ³Institute for Quantum Information and Matter, Caltech, Pasadena, CA 91125, USA. ⁴Xanadu Quantum Technologies, Toronto, ON M5G 2C8, Canada. ⁵Department of Mathematics, University of Toronto, Toronto, ON M5S 1A1, Canada. ⁶Department of Engineering Physics, Ecole Polytechnique de Montréal, Montréal, QC H3T 1JK, Canada. ⁷Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany. ⁸Helmholtz-Zentrum Berlin für Materialien und Energie, 14109 Berlin, Germany. ⁹Department of Mathematics and Computer Science, Freie Universität Berlin, 14195 Berlin, Germany. ¹⁰Department of Computer Science, The University of Chicago, Chicago, IL 60637, USA. ¹¹Institut für Theoretische Physik und Center for Integrated Quantum Science and Technology (IQST), Albert-Einstein-Allee 11, Universität Ulm, 89069 Ulm, Germany. *Corresponding author. Email: ishdhand@gmail.com †These authors contributed equally to this work.

We first address the open theoretical questions about GBS, namely, hardness in the regime with little overall noise in the form of optical loss. More specifically, to provide complexity-theoretic evidence for the hardness of approximately simulating GBS, we prove average-case hardness of computing output probabilities in the noise-free case, formulate the so-called “hiding property” (7) for GBS in terms of a random-matrix theory conjecture, and provide analytical and numerical evidence for this conjecture. These results bring GBS to the level of evidence shared by other QCA proposals such as random circuit sampling (RCS) and conventional boson sampling [see (6, 7, 12)] up to a mild conjecture in random matrix theory. We then show that average-case hardness of computing output probabilities still holds in a regime of high loss rates, building on recent results (13), and discuss the implications of this result on the noise-regimes in which one may still expect GBS to be hard to simulate on a classical computer. These results bolster the evidence for QCA in the USTC experiment and also any future GBS experiments.

Given these theoretical results, we then address the programmability versus low-loss trade-off in current architectures. To this end, we introduce a new architecture, high-dimensional GBS, using a time-domain approach. This architecture can be implemented programmably with low overall loss while at the same time being hard to simulate for the known classical simulation algorithms. The hardness of this architecture is borne out by the hardness of computing output probabilities for the lossy, high-dimensional GBS setup. These results provide evidence of classical hardness for asymptotic system sizes. In the realm of finite system sizes, we take care to avoid regimes where the experiment can be tractably simulated (12, 16), such as when the linear-optical network has limited connectivity [such as one-dimensional (1D) network topology] or when the system is too lossy. Our proposed high-dimensional GBS architecture voids these algorithms by taking advantage of the enhanced connectivity available in higher dimensions than one.

To this end, we perform benchmarking simulations to estimate the cost of high-dimensional GBS against state-of-the-art algorithms for simulating GBS and for simulating high-dimensional quantum many-body systems (17, 18). These simulations give evidence that classically intractable instances of high-dimensional GBS can be built in the laboratory with a small number of optical components. These advantages make high-dimensional GBS an ideal near-term architecture for demonstrating QCA with a programmable photonic device. Thus, by addressing the abovementioned shortcomings of GBS from the theoretical and experimental perspectives and understanding the limits of its classical simulability through both asymptotic analysis and finite-size benchmarking, this work paves the way toward more “loophole-free” demonstrations of QCA with a programmable photonic quantum device.

Hardness of approximate GBS

We begin by reviewing and strengthening the hardness argument for the task of simulating GBS as introduced in (10, 11). We first introduce the model of GBS and then examine the evidence for the hardness of approximate boson sampling. Two properties are required for establishing complexity-theoretic hardness of sampling using the standard QCA arguments, namely, hiding and average-case hardness of approximating probabilities. Here, we strengthen the results in (10, 11) by providing strong evidence for these properties in GBS. Specifically, we reduce the hiding property to a highly plausible conjecture in random matrix theory, for which we provide analytical

and numerical evidence. In addition, we provide evidence for approximate average-case hardness by proving approximate worst-case hardness and near-exact average-case hardness of computing the output probabilities. Thereby, up to a random-matrix theory conjecture, we bring the hardness argument for GBS to the same standard as that of boson sampling. We then extend the latter results to the case of computing output probabilities of noisy GBS, which can be well motivated when the noise model describing the experimental data is trusted. These results show that the evidence of a quantum “signal” remains in the output distribution even in the presence of noise. Last, we discuss the implications of these results on the complexity of simulating GBS in the presence of noise.

Recap: GBS

GBS is the computational task of sampling the photon number statistics of a Gaussian state. Obtaining a sample from a typical GBS experiment involves the following steps. First, a general Gaussian state is prepared at the input, often taken to be M single-mode-squeezed vacuum states. These states are then interfered on an M -mode linear/optical interferometer containing beam splitters and phase shifters. Last, the Gaussian state at the output of the interferometer is impinged on M photon number-resolving (PNR) detectors. The resulting pattern of photon number outcomes from the detectors is the required sample. Because single-mode-squeezed states can be generated and interfered deterministically at room temperatures with high rates, GBS is experimentally feasible on large scales already today, as evidenced by the recent experiment from USTC (2).

In more detail, a typical GBS experiment involves interfering M single-mode-squeezed vacuum states with squeezing parameters $\{r_i\}_{i=1}^M$ at an interferometer specified by an $M \times M$ linear-optical unitary matrix U . Note that some of the modes can be optionally prepared in the vacuum state, and these can be specified by setting their squeezing parameter to zero.

The probability of detecting n_1 photons in the first mode, n_2 in the second, and so on, denoted by $\mathbf{n} = (n_1, \dots, n_M)$, is

$$\Pr(\mathbf{n}) = \frac{|Haf(A_{\mathbf{n},\mathbf{n}})|^2}{\prod_{j=1}^M n_j! \cosh r_j}$$

Here, $A = A^T = U(\oplus_{i=1}^M \tanh(r_i))U^T$ is the so-called adjacency matrix of the (pure, zero displacement) Gaussian state (10), and $A_{\mathbf{n},\mathbf{n}}$ is the symmetric matrix of size $N = \sum_{i=1}^M n_i$ (i.e., the total photon number) obtained by repeating the i^{th} column and row of A a total of n_i times. In particular, if $n_i = 0$, then the corresponding row and column is deleted. Last, the Hafnian $Haf(\cdot)$ of a symmetric $N \times N$ matrix B is given by

$$Haf(B) = \sum_{\mu \in \text{PMP}(N)} \prod_{(i,j) \in \mu} B_{ij}$$

where $\text{PMP}(N)$ is the set of perfect matching permutations of N elements for even N , i.e., permutations $\mu: [N] \rightarrow [N]$ satisfying $\mu(2k-1) < \mu(2k)$, $\mu(2k-1) < \mu(2k+1)$. Equivalently, this is the set of all $N!/(2^{N/2}(N/2)!) = (N-1)!!$ ways of partitioning the set $\{1, 2, \dots, N\}$ into $N/2$ subsets of size 2. The Hafnian of a 0×0 matrix is defined to be 1, and that of an odd-size matrix is defined to be 0, which is a manifestation of the fact that squeezed states are supported on even photon number states only. By allowing for arbitrary

linear-optical unitaries and arbitrary squeezing parameters on each squeezer, an arbitrary symmetric matrix A can be encoded (up to scaling prefactors) into a Gaussian state. For generic instances, the best-known algorithms to calculate Hafnians have a runtime scaling as $N^3 2^{N/2}$, where N is the size of the matrix (14).

Recap: Approximate sampling hardness of boson sampling

Before we state our technical results, we review the main steps of the hardness argument for conventional boson sampling as given by Aaronson and Arkhipov (7). These steps provide context for the hardness results of GBS that we present below.

In a standard boson sampling experiment, instead of interfering single-mode-squeezed states at an interferometer as done in GBS, an N -photon M -mode Fock state is prepared and evolved under a linear-optical unitary and then measured in the photon-number basis. The boson sampling task is to, given a linear-optical unitary as an input, output samples from the output distribution of a corresponding boson sampling experiment.

Aaronson and Arkhipov showed that it is not possible for a classical computer to efficiently do this task unless certain complexity-theoretic conjectures are false. In particular, they reduced the task of approximating the probabilities of outputs to the task of efficient sampling, making use of an approximate counting algorithm due to Stockmeyer (15). This probability estimation can in turn be related to approximating the permanent of a certain submatrix of the linear-optical unitary, which is probably hard for a class known as #P (16). While the Stockmeyer reduction is not efficient, the existence of a classical efficient sampling algorithm would imply that #P-hard problems could be solved using fewer computational resources than expected, amounting to an argument by contradiction.

The main difficulty in the hardness argument for boson sampling arises when extending it to the setting of approximate sampling. Here, the task is to sample from any distribution that is within constant-size total variation distance from a given ideal boson sampling distribution. This additional constraint takes into account that actual devices are bound to achieve only some finite and typically additive precision. In this setting, one may therefore argue for a separation of computational power between quantum and classical devices.

Given this constraint, the hardness argument for the task of approximate sampling must take into account that the constant error budget on the distribution can be distributed arbitrarily across all outcome probabilities. In particular, this means that any specific outcome probability of the actually sampled distribution might have a large (constant-size) error when compared to the ideal distribution, which would imply that the sampler cannot be used to estimate the true outcome probabilities. To get around this issue, the argument is extended to random problem instances: Via a property of the distribution over problem instances called hiding, one can then translate typical outcomes of fixed instances to fixed outcomes of random instances. This enforces that with high probability, the overall constant error budget for the entire distribution is manifest in small errors on the individual probabilities that are proportional to the inverse size of the sample space, that is, $\propto 1/\binom{M}{N}$. Technically, in standard boson sampling, showing the hiding property boils down to showing that the distribution of any small enough submatrix of a Haar-random unitary is approximately (in total variation distance) an entry-wise complex normal distribution. This implies that all collision-free outcomes are (approximately) equally distributed. In particular, Aaronson and Arkhipov (7) show that when $M \in \omega(N^5)$,

we can “hide” a random Gaussian matrix in a small enough submatrix of the large Haar-random unitary by an appropriate procedure because all of these submatrices are indistinguishable from random Gaussian matrices.

For the approximate sampling task to remain computationally intractable, it remains to show that estimating the outcome probabilities up to inverse exponentially small error is #P-hard for any large-enough fraction of the problem instances—a property called approximate average-case hardness. More precisely, given a random problem instance, approximating the probability of a given outcome must be #P-hard with high probability. As evidence toward this property, it has been shown that exactly computing those output probabilities is, in fact, #P-hard on average (and this was a motivation for boson sampling in the first place), and it is known that estimating them to the required robustness level is worst-case hard. However, the hardness of computing those probabilities to a sufficiently large robustness level on average is still unknown. We now state our results concerning the hardness of general GBS, followed by our proposal for an architecture to perform high-dimensional GBS.

RESULTS

Hiding for arbitrarily many squeezers in GBS

As mentioned above, the property of hiding in boson sampling can be translated into a property of the distribution of submatrices of random linear-optical unitaries chosen from some distribution. We will now show that a similar property about the distributions of submatrices occurring in the evaluation of outcome probabilities also holds in GBS, provided that a plausible random-matrix theory conjecture holds. We focus on the paradigmatic setting in which the linear-optical unitary is drawn from the Haar measure, and we fix the input state to be such that the first K out of M modes are prepared in single-mode-squeezed states with identical squeezing parameter r , and the remaining $M - K$ modes are prepared in the vacuum state. Furthermore, we restrict to collision-free outcomes \mathbf{n} for which $n_i \in \{0, 1\}$, giving rise to a total photon number $N = \sum_{j=1}^M n_j$. The probability of obtaining such an outcome \mathbf{n} can be written as

$$\Pr(\mathbf{n}) = \frac{\tanh^N(r)}{\cosh^K(r)} |\text{Haf}[(UI_K U^T)_{\mathbf{n}, \mathbf{n}}]|^2$$

Here, $I_K = \mathbb{1}_K \oplus 0_{M-K}$ denotes the matrix where $\mathbb{1}_K$ is a K -dimensional identity matrix, 0_{M-K} is an $M-K$ -dimensional all-zero matrix, and as before, the notation $A_{\mathbf{n}}$ stands for the submatrix of A corresponding to the entries of \mathbf{n} [see below Eq. (1)]. The task of estimating output probabilities of GBS, hence, corresponds to estimating $|\text{Haf}[(UI_K U^T)_{\mathbf{n}, \mathbf{n}}]|^2$.

To show the GBS hiding property, we need to characterize the distribution of matrices $(UI_K U^T)_{\mathbf{n}, \mathbf{n}}$ —of which the Hafnian is computed—as induced by the Haar-random choice of U and depending on the scaling relations between K , N , and M . To ensure that for every choice of K , we can restrict to collision-free outcomes, we choose the squeezing parameter r such that the average photon number $\mathbb{E}[N] = K \cdot \sinh^2 r \in o(\sqrt{M})$ (7). This condition ensures that the collision-free outcomes dominate the probability weight.

Here, we formulate the hiding property in GBS in terms of random matrix theory and provide strong numerical and analytical evidence that it holds regardless of the fraction of squeezed input

modes so long as the collision-free condition is satisfied. Observe that the matrix $(UI_K U^T)_{\mathbf{n},\mathbf{n}} = U_{\mathbf{n},1_K} U_{\mathbf{n},1_K}^T$ can be expressed in terms of the submatrix $U_{\mathbf{n},1_K}$ of U obtained by choosing rows according to \mathbf{n} and the first K columns. To show the hiding property, we need to relate this distribution over matrices to the distribution of the symmetric product XX^T of a complex Gaussian $N \times K$ matrix X with mean 0 and variance $1/M$, denoted as $X \sim \mathcal{G}_{N,K}(0,1/M)$. We provide analytical and numerical evidence for the conjecture that these distributions are indistinguishable for any number of squeezers K satisfying $N \leq K \leq M$.

Conjecture 1 [hiding in GBS (informal)]

For any K such that $N \leq K \leq M$ and $N \in o(\sqrt{M})$, the distribution of the symmetric product $U_{\mathbf{n},1_K} U_{\mathbf{n},1_K}^T$ of submatrices of a Haar-random $U \in U(M)$ closely approximates the distribution of the symmetric product XX^T of a Gaussian matrix $X \sim \mathcal{G}_{N,K}(0,1/M)$ in total variation distance.

We provide a formal statement of the conjecture in the Supplementary Materials. There, we also discuss regimes in which the conjecture is known to be partially true (7, 17) and provide numerical evidence for it. Proving this conjecture is an open research problem in random matrix theory.

Conjecture 1 characterizes the distribution of the symmetric product of $N \times K$ submatrices of Haar-random unitaries. In turn, the Hafnian of such symmetric products determines the output distribution of GBS. While in standard boson sampling, the hiding property amounts to hiding a small $N \times N$ Gaussian matrix in a large $M \times M$ Haar-random unitary matrix, in GBS, it amounts to hiding a small $N \times N$ symmetric Gaussian matrix XX^T in a large symmetric unitary matrix $UI_K U^T$ for any $K \geq N$. This means that any particular submatrix cannot be distinguished from any other such submatrix of the same size, enforcing the constant error budget to be roughly equally distributed across all outcomes.

In particular, the conjecture implies that the hiding property can be achieved with any number K of input squeezers as long as the average total photon number is sufficiently small. In turn, the average total photon number is determined by the total amount of squeezing across all input squeezers. Intuitively, this is due to the fact that the output of a Haar-random unitary does not depend on any fixed input state. The average output state is a product of identical thermal states whose average photon number is determined by the total input squeezing. However, the number K is still crucial for the estimation task as it determines the rank of the matrix $(UI_K U^T)_{\mathbf{n},\mathbf{n}}$. Since the complexity of computing the Hafnian of a matrix depends on the rank of that matrix (14), K should be chosen such that it is at least N . Note that the USTC experiment (2) used $K = M/2$ many squeezers, so our results are directly applicable there, strengthening the arguments for their QCA demonstration.

More generally, we consider three regimes of interest and provide evidence for Conjecture 1 in the Supplementary Materials. First, the highly sparse regime in which the total number of modes scales as $M = \omega(K^5)$, and the number of photons is equal to the number of squeezers, $N = K$, features provable hiding results due to (7). Realistic experiments and proposals today operate in the regime $K = cM$, meaning that a constant fraction c of the input modes is squeezed. In this regime, the result in (17) provides analytical evidence for hiding in the asymptotic limit as long as the input squeezing is such that $N \in o(\sqrt{M}/\log M)$. Last, we also consider the intermediate regime of how M scales with K between these two extremes and give numerical evidence for hiding in this general case.

Let us note that we do not expect Conjecture 1 to hold for large $N \in \omega(\sqrt{M})$. In this case, it is known that hiding fails for standard boson sampling (18, 19).

Average-case hardness of computing GBS probabilities

As outlined earlier, the question of hardness of approximate sampling boils down to whether it is #P-hard to approximate most output probabilities. We now show the average-case hardness of this task when the allowed additive approximation error is exponentially small, using techniques from (13). We have established that the output probabilities of GBS are given in terms of $|\text{Haf}((UI_K U^T)_{\mathbf{n},\mathbf{n}})|^2$. By virtue of the previous discussion and more precisely, Conjecture 1, the distribution over the $N \times N$ matrices $(UI_K U^T)_{\mathbf{n},\mathbf{n}}$ for Haar random U is well approximated by complex, symmetric Gaussian matrices XX^T . Hence, to show the average-case hardness of computing output probabilities of GBS, it suffices to consider the following problem

(δ, ϵ) -SQUARED-HAFNIANS-OF-GAUSSIANS.

Input A matrix XX^T with $X \sim \mathcal{G}_{N,K}(0,1/M)$.

Output $|\text{Haf}(XX^T)|^2$ to additive error ϵ , with probability $\geq \delta$ over the distribution $\mathcal{G}_{N,K}(0,1/M)$.

To complete the argument that an efficient classical approximate sampling algorithm for GBS cannot exist, it remains to prove the #P-hardness of (δ, ϵ) -SQUARED-HAFNIANS-OF-GAUSSIANS as formalized by the following approximate average-case hardness conjecture.

Conjecture 2

The (δ, ϵ) -SQUARED-HAFNIANS-OF-GAUSSIANS problem is #P-hard for any $\epsilon = O(N! \tanh^N(r)/(\cosh^K(r) M^N))$ and any constant $\delta > 3/4$.

A proof of Conjectures 1 and 2 would imply that approximate sampling from a random, general GBS instance, is hard on average. Let us see how. Assume that there exists a classically efficient sampler O that samples from an associated distribution whose output probability for outcome i is given by q_i . From the promise that this distribution is ϵ -close in total variation distance to the target distribution, we have $\sum_i |p_i - q_i| \leq 2\epsilon$, where p_i is the corresponding output probability of the target distribution. Choose a photon number N so that Conjecture 1 is satisfied. Among the space of all outcomes with N total photons, for a randomly chosen outcome i , we have

$$\Pr_i \left[|p_i - q_i| \leq \frac{2\epsilon k}{(M+N-1)} \right] \geq 1 - \frac{1}{k} \tag{1}$$

Assuming Conjectures 1 and 2, with probability at least $3/4$, p_i is #P-hard to compute to additive error $\epsilon' = O(\frac{N!}{M^N})$. Therefore, with probability at least $3/4(1 - 1/k)$, it is also #P-hard to compute q_i to within error $\epsilon' + \frac{2\epsilon k}{(M+N-1)} = O(\exp[-N \log N - \Omega(N)])$ assuming $M = \Theta(N^2)$. On the flip side, the Stockmeyer algorithm (15) allows us to compute the output probability of an arbitrary outcome q_i to within inverse-multiplicative polynomial precision. Furthermore, by the Markov inequality, most outcomes q_i cannot be much larger than $1/(M+N-1)$

$$\Pr_i \left[q_i > \frac{l}{(M+N-1)} \right] \leq \frac{\Pr(N)}{l} \leq \frac{1}{l} \tag{2}$$

where the quantity $\Pr(N)$ is the probability of seeing N total photons. This means that with probability at least $1 - 1/l$, q_i can be

computed to additive error $O(l \exp[-N \log N - \Omega(N)])$ using a BPP^{NP^0} machine running the Stockmeyer algorithm. Therefore, setting $l = 4k$, a PH algorithm can solve with high probability a problem that is average-case $\#P$ -hard. This collapses the polynomial hierarchy.

Note that since we have phrased Conjecture 2 in terms of additive error instead of multiplicative error, we do not explicitly need an anticoncentration condition of the form $\Pr_X[p_0 \geq (M + N^{-1})^{-1}] \geq \gamma$ for some constant $\gamma > 0$, as is often conjectured for permanents (7). Nevertheless, it is possible that Conjecture 2 already implies a weak form of anticoncentration. Informally, an anticoncentration condition states that on a large fraction of the instances, the output probabilities are large enough so that a trivial algorithm for computing the probabilities that outputs “0” is not sufficient to solve the (δ, ϵ) -SQUARED-HAFNIANS-OF-GAUSSIANS problem. This is because in order for Conjecture 2 to be true, it is necessary for the trivial algorithm to fail with high probability.

As in all other known proposals for demonstrating QCA, this approximate average-case hardness conjecture remains open. Nonetheless, just like in other proposals, it turns out that one can give evidence for Conjecture 2. Namely, we can prove a weaker version of the conjecture with a smaller robustness level $\epsilon = O(\exp[-6N \log N - \Omega(N)])$ as opposed to $\epsilon = O(\exp[-N \log N - \Omega(N)])$ in Conjecture 2.

Theorem 3

The (δ, ϵ) -SQUARED-HAFNIANS-OF-GAUSSIANS problem is $\#P$ -hard under PH reductions for any $\epsilon \leq O(\exp[-6N \log N - \Omega(N)])$ and any constant $\delta > 3/4$.

We provide a detailed proof of Lemma 3 in the Supplementary Materials. The technique we use in the proof is a worst-to-average case reduction [see (7)]. That is, by assuming access to an oracle for the (δ, ϵ) -SQUARED-HAFNIANS-OF-GAUSSIANS problem, we show that one in fact approximate $\text{Haf}(XX^T)$ for any matrix $X \in \mathbb{C}^{N \times K}$. This latter task is $\#P$ -hard in the worst case as we show in the Supplementary Materials. At a high level, the worst-to-average case reduction relies on the fact that $|\text{Haf}(XX^T)|^2$ is a low degree (of degree $2N$) polynomial over the entries of the matrix X . This allows us to use the oracle to perform polynomial interpolation. Therefore, by combining this observation with the techniques in (7, 13, 20), we obtain a worst-to-average case reduction for exactly computing the output probabilities.

Together, our results on the hiding property and the approximate average-case conjecture in GBS strengthen the evidence for the hardness of approximately simulating GBS in terms of the total variation distance to the ideal output distribution. Given our results, GBS is now on par with the other leading QCA proposals in terms of complexity-theoretic evidence for approximate sampling hardness (6, 7, 12, 13, 20, 21), up to a plausible conjecture in random matrix theory—for which we provided theoretical and numerical evidence. To achieve a demonstration covered by those complexity-theoretic results, however, the loss rate at every element of the linear-optical circuit must scale inversely with the total number of these elements—a daunting challenge from an experimental perspective.

Hardness of computation of output probabilities for noisy GBS

We now go one step further and assess how the complexity-theoretic argument for sampling hardness is affected by more realistic noise levels, particularly, in terms of photon loss. In terms of scaling, any constant loss rate of the individual optical elements can lead to

the output distribution rapidly approaching a classical distribution. We now show that, nonetheless and unexpectedly, an evidence of a quantum signal remains even in the presence of significant loss. We then discuss to what extent and in which regimes such a quantum signal might lead to the hardness of simulating a lossy GBS experiment.

One of our main results is the average-case hardness of computing the noisy output probability of a random GBS instance, which we obtain by using similar arguments to recent work of Bouland *et al.* (13) but now extended to the GBS setting. Our results are valid for any noise model that is local, stochastic, and is error detectable using linear optics. More specifically, we consider a setting where the noise acts locally after every gate and is of the \mathcal{E} form

$$\mathcal{N}_i[\rho] = (1 - \eta_i)\rho + \eta_i \mathcal{E}_i[\rho]$$

where stochasticity requires \mathcal{E}_i to itself be a valid channel (i.e., a completely positive trace preserving map) with no identity component.

Consider the following problem.

(ϵ, η) -NOISYGBS-PROBABILITY.

Input A noisy GBS instance, consisting of the linear-optical unitary U on M modes chosen from the Haar measure \mathcal{H} , the squeezing parameters at the input, a description of the noise channels with parameters η_i , and a description of a collision-free outcome \mathbf{n} with $N = \text{poly}(M)$ total photons. Let $\eta = \max_i \eta_i$.

Output With probability δ over instances, an estimate of the quantity $\Pr(\mathbf{n})$ to additive error ϵ , where $\Pr(\mathbf{n})$ is the probability of obtaining outcome \mathbf{n} .

With probability $1 - \delta$, an arbitrary output.

In the above definition, we take $\delta = 1$ to mean the worst-case problem. We prove the following statement of average-case hardness of computing noisy probabilities.

Theorem 4

There exists a noise threshold η^* and a sufficiently large polynomial such that the problem (ϵ, η) -NOISYGBS-PROBABILITY is $\#P$ -hard under PH reductions for any constant $\delta > 3/4$, $\eta \leq \eta^*$, and $\epsilon \leq 2^{-\text{poly}}$.

There are two parts to the proof. The first part is a proof of worst-case hardness of the problem (when $\delta = 1$) and the second a worst-to-average case equivalence. For worst-case hardness, it turns out that because of a result of Fujii (22), it suffices for the noise channel to be a convex combination of the lossless and lossy channels and to be able to error-detect it. These conditions are both met for optical loss since it is a convex combination of the channels corresponding to no photon loss, single-photon loss, and so on (23). Moreover, optical loss can also be detected and corrected using only linear-optical operations and photodetection with high thresholds (9). In Fujii’s argument, one postselects on the error-free outcome of an error-detection code and obtains noiseless universal gates for the class of postselected quantum computation, postBQP. This argument can apply to the optical case as well since linear optics with postselection is universal for quantum computing (24).

For the worst-to-average case equivalence, all we need is for the polynomial structure in the problem to be preserved. This can be satisfied for any local noise model. Preserving the polynomial structure of the output probability enables us to continue to use the same proof techniques as earlier.

Before moving on, we again remind the reader that we considered the hardness of computing output probabilities. While these

are not tasks that are feasible for any realistic quantum device, our results nevertheless indicate that there is a computationally intractable (but exponentially small) “quantum signal” present in the system.

The complexity of noisy and approximate GBS

We now discuss the implications of the hardness result for computing noisy GBS probabilities on the complexity of sampling from the output distribution of noisy GBS. An immediate implication of this result is that it is classically hard to exactly sample from the noisy distribution of a worst-case GBS experiment. This is because the quantum signal is still present in the distribution, so the argument based on Stockmeyer’s algorithm is valid. Thus, in the idealized situation in which loss is the only source of noise of an experimental system and the exact loss rate is known, simulating a worst-case GBS experiment is classically intractable. Note that loss rates can be inferred from standard optical tomography procedures such as that in (25). Given that this result links the hardness of simulating the noisy experiment to an exponentially small quantum signal in the form of output probabilities, it is crucial that the noise model accurately captures the working of the device.

We remark that an alternative proof establishing the classical hardness of exact sampling could possibly be made using a postselection argument similar to the one outline in section 4.2 of (7). As noted in (7) however, this approach has not been shown to provide a viable path toward the goal of showing hardness of approximate sampling. By establishing the average case hardness of approximating output probabilities, Theorem 4, takes a substantive step toward establishing the hardness of approximate sampling, even in the presence of noise.

We now discuss the more realistic situation in which loss is the predominant, but not the sole, source of noise in a photonic experimental system. What can we say about the hardness of approximate sampling in such a situation? To begin with, let us draw on some intuition from RCS schemes acting on n qubits. Here, the additive error incurred in estimating output probabilities using the Stockmeyer algorithm is $O(2^{-n})$ with high probability (since this is the size of a typical output probability in an RCS experiment). In the presence of uncorrected noise, an error of $O(2^{-n})$ in the noisy output probability can be too large for hardness. For example, there is evidence that with gate-wise depolarizing noise, the probabilities will deviate from uniform by merely $O(2^{-m})$, where m [typically $\omega(n)$] is the total number of gates (6). This means that approximate-sampling hardness cannot be shown using these techniques since it is not hard to approximate the noisy probabilities any more. In this regime, the noisy distribution is exponentially close in total variation distance to the uniform distribution, rendering the approximate sampling task for the noisy distribution classically simulable.

In the case of noisy GBS, the dominant noise model, namely, loss, leads to the vacuum state for a sufficiently deep network, which is again a distribution that is easy to classically sample from (similar to the uniform distribution in qubit RCS schemes). However, if we postselect on a certain minimal number of photons surviving, the distribution need not be easy to simulate. This postselection is efficient when the depth of the circuit scales poly-logarithmically in the number of modes. In this case, the quantum signal will be large enough so that even with an inverse exponential error, deviations from the easy distribution can be detected.

This excludes the simulation algorithm that samples from an easy-to-simulate distribution such as the one uniform on every

photon number sector with every sector sampled according to the ideal photon number distribution. Ruling out trivial algorithms is a necessary condition for approximate average-case hardness to hold. In summary, our results indicate that there might be ‘room in the middle’ in terms of gate depth and noise rates, where hardness of sampling might hold. This intuition lies at the heart of the high-dimensional architecture (presented below). This architecture is designed in such a way that only as few gate applications as necessary for hardness are executed, so that the leeway for noise to ruin the hardness of sampling is minimized. We stress, however, that at the moment, existing proof techniques do not suffice to make a claim of this nature. In certain regimes of noisy GBS, approximate sampling is known to be classically efficient (26).

High-dimensional GBS and evidence for hardness

The discussion thus far here and in the literature has focused on the hardness of GBS with unitary transformations drawn randomly from the Haar measure. This requires implementing arbitrary unitary transformations, an onerous requirement experimentally. Reference (2) did not meet this requirement of being able to implement arbitrary unitary transformations as a result of the interferometer being a fixed nonprogrammable device. Furthermore, there is reason to believe that in the absence of error-correction methods for linear optics, scaling arbitrary programmable interferometers to large numbers of modes is infeasible. This is because implementing an arbitrary unitary transformation requires decomposing it into beam splitters and phase shifters, and assuming that they are all applied locally, this leads to a deep optical circuit, whose depth linearly scales with its size. Since photon loss scales exponentially with the circuit depth, these models necessarily become efficiently simulatable classically for sufficiently large numbers of modes (26, 27).

On the other hand, naively reducing the depth without giving up gate locality is not an option for QCA either. This is because shallow 1D circuits comprising local interactions with logarithmically scaling depths can be efficiently simulated classically as these do not generate enough long-range entanglement (26–28).

These results motivate a demonstration of QCA on random optical circuits with shallow depth but with gates that are long-range in 1D, for example on circuits with local interaction in higher than one dimensions. In such a setting, a potentially reduced amount of complexity due to the reduced depth would be compensated by the large long-range entanglement generation thanks to the inclusion of long-range interactions. Therefore, such architecture would suffer less noise buildup but still remains intractable for classical computers. Models with shallow-depth but with long-range (in 1D) interactions provide a natural approach to demonstrating QCA in qubit systems (29).

We address the challenge of the low-loss versus depth trade-off by introducing high-dimension GBS, where programmable nonlocal gates are exploited to generate entanglement between distant modes. We show how high-dimensional GBS can be implemented scalably using optical delay lines. Before presenting the new architecture, let us recall the relevant notation on GBS and discuss its physical implementation.

A programmable architecture for high-dimensional GBS

Now, we are ready to introduce high-dimensional GBS: A sampling task that retains the programmability of the photonic device can reduce decoherence to a level that prevents classical simulability and

in which large amounts of multipartite entanglement can be generated. The last two requirements are to some extent at odds with each other: Specifically, achieving long-range interactions in fixed linear 1D geometries requires finding intermediary quantum systems to mediate interactions between far separated regions, which can lead to information leaking into the environment and require more challenging experimental conditions than all-optical experiments. A way around this challenge is to consider two- or higher-dimensional geometries where quantum systems can interact with each in more than one direction. While the Google QCA experiment (1) involved interactions in 2D, our proposal can leverage photonics to implement distant nonlocal interactions, which can be equivalently considered as interactions in two or even higher than two dimensions.

More specifically, we show how the idea of using local interactions in high-dimensional spaces to generate large amounts of multipartite entanglement can be naturally imported into photonic quantum computing by using optical delay lines and fast, programmable optical switches. Before formally stating the problem of high-dimensional GBS, we provide intuition for how to construct high-dimensional lattices using minimal optical resources. For the sake of concreteness and ease of visualization, we consider the generation of a lattice of size $a = 3$ in $D = 2$ dimensions where the vertices represent modes, and the edges represent two-body gates. A quantum circuit to achieve this connectivity and a representation of the obtained lattice are shown in Fig. 1 (A and B), respectively. Note that when the bosonic modes are represented as wires in a usual quantum circuit diagram, the gates needed to prepare the state are highly nonlocal. This is because circuit diagrams provide a representation where the modes are arranged linearly (in this case, in the vertical direction of the page). To show how optical delays provide a natural way to program short and long ranged interactions, consider first our temporal modes (pulses) prepared in squeezed vacuum states arranged one after the other, traveling along a single spatial mode, as schematically shown in Fig. 1C.

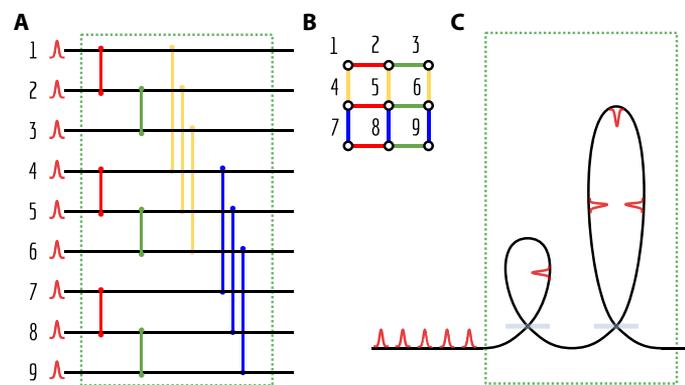


Fig. 1. Different representations of a $D = 2$ D optical delay GBS instance with lattice size $a = 3$. (A) Circuit representation. The vertical lines with dots at the end represent beam splitters. (B) Bidimensional lattice representation. The vertices of the lattice represent the modes, while edges represent beam splitters. (C) Optical circuit representation. The modes are defined by time-bins traveling in a waveguide. The horizontal gray slabs at the bottom of the delays represent the beam splitters. The number of cycles C in a high-dimensional GBS instance corresponds to applying multiple times the gates contained in the green-dotted box in (A). This action physically maps to using concatenating C copies of the delays encircled in the green box in (C). Note that for simplicity, we have not shown the photon-number detectors used to probe the quantum state at the end of the circuit.

We first consider how to achieve nearest-neighbor interactions using a delay line whose length equals the separation between the pulses. As mode i is in the delay line about to exit it, it will interfere with mode $i + 1$, which is about to enter the delay line. The beam splitter mediating the interaction between these two modes can be programmed, allowing us to effect two-mode gates between nearest neighbors. Such programmable and fast (i.e., with less than 50 ns spacing) beam splitters have been demonstrated using electro-optic modulation and have been used in the application of photonic quantum walks in the time domain (33, 34).

Now, consider the second delay line, whose length is $a = 3$ times the separation between the pulses. In this case, as mode i is getting ready to exit the delay line, it will interfere with $i + a$ in the beam-splitter gate keeping the delay line. This configuration allows interactions with range $a = 3$ between the modes in the quantum circuit diagram in Fig. 1. Note that this construction generalizes in a natural way to D dimensions. In particular, nearest-neighbor interactions in a D -dimensional space with a lattice points per dimension (corresponding to gates with range a^{D-1} in a circuit diagram) can be implemented using a circuit with D optical delay lines implementing delays by amounts $\{1, a, a^2, \dots, a^{D-1}\}$. If the light is made to pass through D such multiple delay lines, with C passes, then the effective transformation is composed of C cycles of local interactions in a D -dimensional lattice or equivalently, C cycles of up to a^{D-1} -range gates in a circuit diagram.

Having provided a quantum optical implementation of high-dimensional GBS, we are now ready to formalize it by specifying four quantities: the squeezing parameter r , the lattice dimension D , the lattice size a , and the number of cycles C . An (r, a, D, C) -high-dimensional GBS instance is constructed as follows:

- 1) Prepare $M := a^D$ single-mode-squeezed vacua $|r\rangle^{\otimes M}$.
- 2) For $\tau = 1$, apply a beam-splitter V to mode i and $i + \tau$, where $i \in [0, M - \tau]$.
- 3) Repeat step 2 for $\tau = a^d$ for $d = 0, \dots, D - 1$.
- 4) Repeat steps 2 and 3 for a total of C times.

Having a physical architecture to implement high-dimensional GBS, we can now write down a loss budget to account for the bulk of the decoherence affecting our system. Assume that the photon-number detectors used to probe our quantum state are limited by a rate of ν detections per second, for example, as a result of the detectors dead times. From this time scale, we deduce a length scale $\ell = \nu/\nu$, where ν is the speed of light in the delay lines. We associate with the length scale an energy transmission constant $\eta_{\text{unit-length}}$, which is simply the total energy transmission resulting from a propagation over a total length of ℓ .

Let us first study the case $C = 1$. In this case, every mode will traverse D beam splitters (to access the D different delay lines) and will propagate a total distance of $\ell \times \sum_{i=0}^{D-1} a^i = \ell \times \frac{a^D - 1}{a - 1} \approx \ell a^{D-1}$ if $a \gg 1$. We can approximate the total transmission to scale roughly as

$$\eta = \eta_{\text{BS}}^D \eta_{\text{unit-length}}^{a^{D-1}} = \eta_{\text{BS}}^D \eta_{\text{unit-length}}^{M^{1-1/D}}$$

where η_{BS} is the beam-splitter transmissivity for programmable beam splitters based on electro-optic modulation. Note that in this case, the loss scales subexponentially with the total number of modes. To allow two or more circulations, one can consider $C \geq 2$ copies of the original D delay lines, giving now an updated loss budget in which the modes traverse a length proportional to $Ca^{D-1} = CM^{1-1/D}$ and will pass through CD beam splitters, still leading to subexponential

loss accumulation. An alternative to these C copies of the delay lines is to consider a recirculation loop similar to that proposed in (30), which reroutes the output of the last delay line into the input of the first one. The delay line used to implement the recirculation loop holds any modes that are not interfering inside the delay lines. If the recirculator has a loss per unit length $\eta_{\text{unit-recirc}}$, the net loss scales as $\eta_{\text{unit-recirc}}^L$ where $L = a^D - \sum_{i=1}^{D-1} a^i = \Theta(M)$. Thus, depending on the exact setting, for a fixed C , the losses scale either exponentially (using recirculators) or subexponentially (considering multiple copies of the D loops) with the number of modes.

We note that with current fiber-optic and photon number-resolving technology, $\eta_{\text{unit-length}}$ can be as high as 0.998; η_{BS} values of 0.9 are expected or are observed in state of the art experiments such as in (31). With these, the transmission of an interferometer with parameters ($a = 15, D = 2, C = 2$) can be above 0.70 and above 0.74 for ($a = 6, D = 3, C = 1$). These values promise an order or magnitude or more enhancements in loss values as compared to those expected in fully programmable GBS devices (32). As noted in (33), interferometers implemented using loops will typically have unbalanced losses. The numbers quoted above assume the lossiest interferometer implementable in a loop-based system, which is precisely the one in which each and every mode is fully transmitted into each loop.

From the formal description of high-dimensional GBS, the covariance matrix of the generated Gaussian state can be calculated in the usual manner. In particular, we only need to specify the unitary matrix describing steps (2) to (4) above. This unitary matrix is given by

$$U = \bigotimes_{c=1}^C \bigotimes_{d=0}^{D-1} \bigotimes_{i=0}^{M-a^d} B_{i,i+a^d}(V)$$

where $B_{ij}(V)$ is an $M \times M$ unitary matrix that acts like the locally Haar-random beam-splitter V in the subspace of modes i and j and like the identity elsewhere. We denote by \mathcal{U} the ensemble of linear-optical unitaries applied this way. In Fig. 2, we show heatmaps of the unitary matrices associated with two typical instances from the distribution \mathcal{U} over high-dimensional GBS instances. Note that the structure of circuits considered allows for light from the first mode to be observed in any of the later modes, which leads to a large light cone that is somewhat different from the efficiently simulable circuits

considered recently in (34). From the description of the unitary matrices and the squeezing parameters, we obtain that the complex-valued adjacency matrix (as defined above) of the Gaussian state is dense, full-rank, and given by $A = \tanh(r)UU^T$.

While implementing a time-domain reconfigurable loop architecture as described above is not a straightforward task, several groups have performed experiments with tens of modes interfering in time-domain multiplexed configurations. These include time-bin (35) and temporal-to-spatial-encoded (36) boson sampling experiments (35) and controllable photonic random walk over multiple time-bins (37, 38). Moreover, recent experiments have shown that it is possible to operate with very high phase stability (39), high quantum-efficiency photon-number detection (40), and very low loss reconfigurable interferometric elements (41).

Last, for the purpose of calculating outcome probabilities, squeezed states can be considered in the Fock basis as qudits that are entangled by the beam-splitter operations. This process, as with any other quantum circuit, can be represented as networks of tensors (42). In more detail, here, the qudits are initially single-index tensors (vectors) that are contracted with four-index tensors representing the beam splitters to build an open tensor network (TN), which can then be contracted to obtain the tensor of the final state. The TN representing the state can be used to calculate probability amplitudes of measurements when the output indices of the TN are contracted with vectors representing measurement outcomes. Similar TN-based techniques have been successful at delineating the QCA frontier in the context of RCS, and together with Hafnian-based methods, these will serve a similar purpose for high-dimensional GBS.

Hardness for computing noisy probabilities in high-dimensional GBS

Here, we now argue for the hardness of computing output probabilities for the noisy, high-dimensional GBS setup. In particular, we show that hardness is present even in shallow depth noisy high-dimensional GBS architectures. This is in contrast to the results discussed earlier, where no restriction is made on the depth.

To do this, we simply observe that the previous argument for worst-case hardness, which depends on the noise being local and error detectable, continues to hold for the limited-depth setup (43).

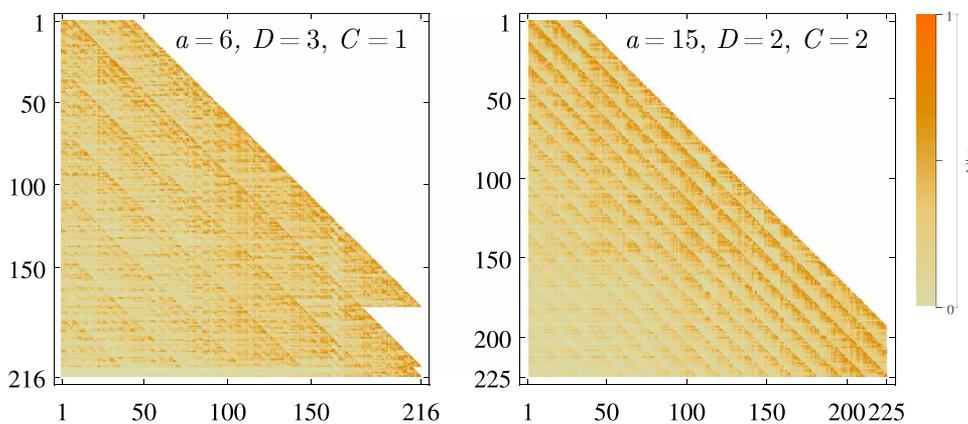


Fig. 2. Absolute values of the entries of the unitary matrices associated with two high-dimensional GBS instances drawn from \mathcal{U} . On the left, we show an ($a = 6, D = 3, C = 1$) instance, and on the right, we show an ($a = 15, D = 2, C = 2$) instance. Note that we explicitly color the zero entries of the unitary white; thus, the color scale is discontinuous at this end.

For average-case hardness of computing noisy probabilities, we again use a worst-to-average case reduction. However, the polynomial interpolation in this case is different since a random instance is not Haar-distributed any more, but rather according to \mathcal{U} , the distribution over random instances of high-dimensional GBS. To explain further, consider the usual interpolation $X(t) = (1-t)X + tY$, where $X(0) = X$ is drawn from \mathcal{U} , and $X(1) = Y$ is the matrix corresponding to a worst-case high-dimensional GBS instance. In this case, there is no guarantee that the interpolated matrices $X(t)$ also correspond to high-dimensional GBS instances of small depth. We get around this issue by choosing a gate-wise interpolation that is similar to that seen in RCS (12, 21).

We first define the problem of computing output probabilities of a restricted-depth high-dimensional GBS architecture.

(ϵ, η) -HIGH-DIMENSIONAL-NOISY-GBS-PROBABILITY.

Input A noisy GBS instance drawn from \mathcal{U} that can be implemented in D dimensions with a constant number of cycles $C = O(1)$ with noise parameter η , and a description of a collision-free outcome \mathbf{n} with $N = \text{poly}(M)$ photons.

Output With probability δ over instances, an estimate of $\Pr(\mathbf{n})$ to additive error ϵ .

With probability $1 - \delta$, an arbitrary output.

Similar to the previous results, we can again obtain an average-case hardness result that we state here and prove in the Supplementary Materials.

Theorem 5

There exists a noise threshold η^* and a sufficiently large polynomial such that the problem (ϵ, η) -HIGH-DIMENSIONAL-NOISY-GBS-PROBABILITY is #P-hard under PH reductions for any constant $\delta > 3/4$, $\epsilon \leq \eta^*$, and $\epsilon \leq 2^{-\text{poly}}$.

QCA frontier for high-dimensional GBS

The evidence presented above for the hardness of high-dimensional GBS comes from complexity-theoretic arguments, which are asymptotic in nature, i.e., they only specify how the hardness of a certain computation scales as the problem size is increased. For a finite-sized device, we now address a complementary but more immediate question: How much actual computational power would a classical adversary need to generate samples similar to those from finite-sized noisy GBS devices?

This question can be addressed with different assumptions about the classical adversary. The experiment can be benchmarked either against simulations that try to match a reasonable model of the experiment (constrained adversary) or against simulations that merely try to spoof a given test (unconstrained adversary). The latter approach would be more rigorous as it requires making fewer assumptions, but coming up with good spoofing methods is a problem beyond the scope of this work and should be seen as an ongoing community effort (5). Similar to the approach of the Google and USTC supremacy experiments (46, 47), we focus on the former approach—with a classical adversary producing samples according to a noisy model distribution—because these samples are likely to perform no worse than the actual device in suitable verification tests. In other words, we assume a specific model of the imperfect GBS device, and we demand that the classical adversary generate samples that have a probability distribution that is sufficiently close in total variation distance to the probability distribution of this model. We note, however, that the chosen model might not have been verified against the actual experiment as this sample-efficient noise-model

verification of QCA experiments is a challenging problem, especially for boson sampling and GBS.

We perform this benchmarking by simulating high-dimensional GBS with state-of-the-art algorithms on the current best supercomputers. In particular, we consider the fastest algorithms based on computing probability amplitudes via Hafnians and via TN contractions. The former, Hafnian-based algorithms have been optimized for simulating GBS and are not restricted to high-dimensional GBS (44). The latter TN algorithms are well-suited for high-dimensional qubit circuits with shallow depth (45). We note that Qi *et al.* (28) also provide a path to simulating lossy GBS if the losses scale exponentially with the system size, but these results are not applicable for high-dimensional GBS, where the losses can scale subexponentially. By benchmarking against these algorithms, we demonstrate that high-dimensional GBS experiments feasible with current optical technology are well beyond the reach of the biggest supercomputers.

DISCUSSION

In this work, we have proposed a new experimental architecture for GBS and provided asymptotic evidence for the hardness of GBS in this specific context, bridging the gap between theory and experiment. We have also benchmarked today's best-known algorithms at simulating such an experiment, obtaining complementary evidence that a reasonably sized setup would outperform classical supercomputers at this task. Still, some theoretical questions are outstanding.

1) We have been able to show that two plausible conjectures in random matrix theory allow us to obtain the hiding property for a noiseless GBS setup, without restrictions on the number of active modes. Can we obtain a similar hiding property for the high-dimensional GBS setup introduced here? Is this also possible in the presence of noise? Answering these questions is crucial for extending the hardness of computing output probabilities to the hardness of approximate sampling from experimentally realizable distributions.

2) Informally, the anticoncentration conjecture for boson sampling (or GBS) states that the output probability of a random instance is unlikely to be very small. If this conjecture was true, then now-standard arguments can show that the output probability corresponding to an approximate sampler is, with high probability, a good multiplicative estimate to the ideal output probability. Proving this conjecture true, in either the case of boson sampling or GBS, would give increased evidence to support the goal of proving QCA via photonics. A proof of such a conjecture is challenging due to the fact that tools of unitary designs (46) are presumably unavailable in the bosonic setting (7).

3) Notwithstanding, it would be insightful to compute the second moments $\mathbb{E}_{X \sim \mathcal{G}(0, 1/M)} |\text{Haf}(XX^T)|^2$ for the distribution we have found to characterize GBS problem instances. These moments thus characterize the so-called collision probability of seeing the same outcome twice in an experiment, which in turn can be related not only to anticoncentration but also the verifiability of approximate GBS from samples, thus shedding some light on the structure of the GBS output distribution.

4) An important task in demonstrating QCA is to verify that the performed experiment indeed contains a nontrivial quantum signal that cannot be efficiently spoofed. The Google QCA demonstration relied on linear cross entropy benchmarking fidelity, and the USTC experiment used a heavy-output generation (HOG) ratio test as an

alternative path to verifiable hardness. Whether the HOG-ratio test can be spoofed efficiently by a classical adversary such as the algorithms considered in (4, 5) is an open problem.

5) The recent result in (34) presents a classical algorithm for the simulation of high-dimensional boson sampling experiments in certain regimes. As described, this algorithm is not applicable to the architecture we propose here. Extending the algorithm to be relevant to the present architecture is an open problem.

6) With current optical technology, loss is the dominant source of noise in any GBS experiment. Consequently, we were motivated to obtain hardness results for computing the output probabilities of a GBS experiment in the presence of significant photon loss. It is natural to investigate whether similar hardness results can be obtained in the presence of other possible sources of experimental noise, such as mode mismatch, multiple Schmidt modes, interferometer phase drift, and detector dark counts.

7) It is a challenge to the community, after all, to relate boson sampling closer to practically important computational tasks and to identify new applications.

In summary, this work brings the demonstration of QCA on a programmable photonics device closer to reality. It addresses previously outstanding theoretical challenges in the field by providing stronger evidence for the hardness of GBS. Crucially, we have presented a novel architecture for high-dimensional GBS using optical delay lines that promises low levels of noise without compromising on its programmability. We benchmarked this architecture against the best available classical simulation algorithms and found that already, experiments involving a moderate number of modes are far beyond reach for those algorithms.

We close by briefly commenting on the experimental prospects of realizing high-dimensional GBS. Since high-dimensional GBS can be implemented in the time domain according to the scheme presented in Fig. 1, only a single squeezer and a single detector are required. If multiple detectors are available, these can be demultiplexed using optical switches to increase the effective repetition rate of the experiment and reduce the length of the delay lines. Especially promising is the case of $D = 3$, $a = 6$, $C = 1$, which can be implemented with only three optical delay lines and three each of reprogrammable beam splitters and phase shifters. Assuming reasonable values of squeezer out-coupling losses, free space to fiber coupling loss, and detector efficiency (39–41, 47), we estimate that such a setup can be built using current optical technology with around 40% transmission, higher than that enabled by the ultralow nonprogrammable loss interferometer in the USTC experiment. Such a setup would enable the largest demonstration of QCA, yet with a mean detected photon number of 80 in a programmable device with 216 total modes. We hope that this work stimulates these developments.

MATERIALS AND METHODS

Computational task: Sampling from lossy GBS with finite Fock cutoff

Before looking into concrete strategies for the simulation of GBS, we detail the computational task performed by the GBS device and discuss some differences between the task and our simulation. The experimental device samples from a lossy GBS distribution with a finite Fock-basis cutoff, which results from detector limitations. To identify a range of parameters where this task is hard to simulate classically, we benchmark it against classical simulations. The

simulations that we compare are somewhat different from the exact task performed by the experiment but in such a way that is advantageous to the classical simulations, thus providing stronger evidence for the large computational cost of high-dimensional GBS. We now discuss these differences.

The first point of difference is the Fock cutoff, i.e., the number of Fock or photon-number levels considered in each mode. Both Hafnian and TN simulations are performed in Fock basis, and their performance is thus sensitive to the Fock cutoff. This cutoff must be chosen carefully because the squeezed state inputs in GBS have nonzero support on high Fock numbers (which could be infinite in the ideal case) (44). For Hafnian-based simulations, the Fock cutoff c will lead to a constant prefactor 2^c ($2^{c/2}$) in the runtime for calculating mixed-state (pure state) probabilities that would appear in sampling methods. Similarly, for TN simulations, this cutoff sets the qubit dimension in the calculation, which is also the base of the exponential function describing the time and space cost of contracting the TN. Note that squeezed states of light require that we use local Hilbert spaces with at least dimension 3 since truncating a squeezed state to the first two levels of the Fock ladder will project it into the vacuum as $\langle 1 | r \rangle = 0$. Furthermore, using a Fock cutoff of 3 in the beam-splitter gates leads to highly inaccurate simulations as the beam-splitter transformations on a limited Fock subspace no longer preserve photon numbers. In other words, choosing higher Fock cutoffs will lead to more accurate but more expensive simulations. Hence, we use a cutoff of 4 to give a conservative estimate on the computational cost, although this cutoff would lead to inaccurate classical simulations.

A second point of difference is that our simulations deal with the case of simulating pure states with photon numbers equal to the lossy distribution. This is a reasonable simplification since, as shown in (48), simulating pure or mixed state GBS has the same complexity as calculating a number of pure-state probability amplitudes proportional to the number of modes in the system.

Before describing the effect of loss on the two simulation methods, we discuss the effect of loss on the number of detected photons. In Fig. 3, we plot the lossless and lossy (transmission $\eta = 0.5 \approx 3$ dB loss)

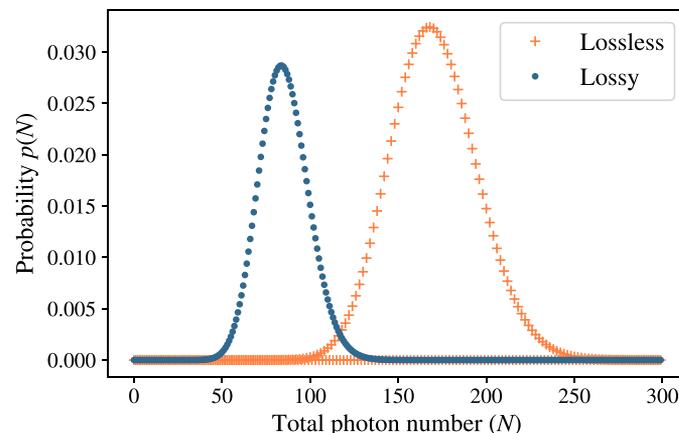


Fig. 3. Distribution of the total photon number for $M = 216$ single mode-squeezed states with squeezing parameter $r = 0.8$. We assume a total transmission of $\eta = 0.5$ (corresponding to roughly 3 dB of loss) for the lossy distribution. Note that the lossless distribution has no support on odd numbers of photons, which explains why visually it looks as if it has more area under the curve.

distribution for $M = 216$ modes and squeezing parameter $r = 0.8$. These parameters have been chosen to correspond with an ($r = 0.8$, $a = 6$, $D = 3$, $C = 1$) high-dimensional GBS instance with experimentally reasonable loss budgets. The squeezing parameter $r = 0.8$ is chosen to be within reach of current sources of single-Schmidt mode degenerate squeezed light (49). Note that the lossy distribution has smaller mean and variance than the lossless one (50), indicating that it becomes easier to simulate a lossy distribution as the transmission η is decreased. For example, the outcome with the highest probability in the lossless distribution

$$n^* = 2 \left\lfloor \left(\frac{M}{2} - 1 \right) \sinh^2 r \right\rfloor = 168 \quad (3)$$

has a probability of 7.28×10^{-8} under the lossy distribution. The leftward shift of this distribution will, in general, be present whenever loss acts on a pure state. For M identical squeezers (with squeezing parameter r) undergoing loss by energy transmission η , the mean and variance contract at least proportionally to η

$$\mathbb{E}(n) = \eta M \sinh^2 r, \text{Var}(n) = \eta M \sinh^2 r (1 + \eta [1 + 2 \sinh^2 r])$$

confirming our intuition and, moreover, showing that the prevailing sources of decoherence in photonic sampling problems behave differently from the ones in RCS implemented in superconducting circuits, where noise makes the output probability distribution become uniform (6).

We now focus on the case of Hafnian-based algorithms. The cost of calculating the relevant probabilities depends only on the number of photons detected. Calculating a photon-number probability $\text{Pr}(\bar{n})$ of a mixed state is roughly quadratically more expensive than calculating a pure state probability of an event with the same number of photons (51). However, the cost of sampling pure and mixed states is similar. This is because lossy GBS states are classical mixtures over a displacement parameter of pure Gaussian states. Therefore, it is possible to sample from a lossy state by sampling from the convex hull parametrized by the displacement parameter and then sampling from the pure state. Thus, sampling lossy GBS states has similar computational cost as sampling pure states with the same number of photons.

Likewise, for the TN-based algorithms, the cost for mixed state calculations would scale at least quadratically worse as compared to pure state calculations. This is because twice as many tensors are involved in a mixed state calculation, analogous to the quadratic overhead of keeping track of the density matrix as compared to a pure state. Note that for noisy RCS of qubits, one can trade fidelity for sampling speed (52). As opposed to GBS, this improvement is possible because in RCS, the amplitudes of the different Feynman-like paths that appear when slicing through two-body gates in the circuit are comparable. Moreover, this improvement is useful as long as the Schmidt-rank of the two-body gates used to generate entanglement is small, which is not the case for the beam splitter. Furthermore, the state vectors associated with two different paths are approximately orthogonal.

A final point of difference between our simulations and the actual experiment is that while our runtime estimates are for the calculation of the GBS probabilities, an actual experiment samples from this distribution. Despite this difference, our simulations allow

a fair benchmarking of the quantum device because current state-of-the-art algorithms have similar complexities of sampling and calculating probabilities. We moreover give the classical adversary an extra advantage in that we allow it to assume that only pure-state output probabilities need to be calculated for sampling, as opposed to the quadratically slower mixed-state output probabilities since, as explained above, mixed-Gaussian states are convex mixtures of pure ones.

In summary, we provide maximal advantage to a classical adversary by choosing a low Fock cutoff, by performing pure state simulations with low photon numbers, and by estimating time for computation rather than sampling (which is at most polynomially slower using currently known methods). This advantage ensures that despite improvements in the classical algorithms, the space of parameters that are hard to simulate classically remains so.

Hafnian-based algorithms

Consider now the probability amplitudes of n -photon events by evaluating the Hafnian. Similar benchmarkings have been performed in the past for the calculation of permanents (53) (relevant to boson sampling) and Torontonians (54) (relevant to GBS with threshold detectors). For either of these two tasks, the time complexity of calculating a probability corresponding to an n -photon event scales like $O(\text{poly}(n)2^n)$, which is quadratically worse than for GBS, which scales as $O(\text{poly}(n)2^{n/2})$. For the case of boson sampling, this difference stems from the fact that any probability amplitude with n photons maps exactly to a GBS instance with $2n$ photons. For the case of threshold detection, it stems from the fact that one cannot assign probability amplitudes to a measurement that is not rank-1, like the positive operator-valued measure representing a “click,” which is a coarse-graining of all the projectors with nonzero photons. In any case, for either of these tasks, benchmarks up to $n = 50$ have been carried out requiring on the order of 2 hours for boson sampling using Tianhe-2 (53) and on the order of 20 hours for GBS with threshold detectors in Sunway TaihuLight (54).

If the matrix has no special property, such as being low rank, nonnegative, banded, or sparse, then the best known algorithms to calculate the Hafnian will scale like $O(n^3 2^{n/2})$ for a matrix of size $n \times n$. The adjacency matrices generated in high-dimensional GBS do not have any of these properties. In Fig. 4, we show the results of our benchmarking by implementing the Hafnian algorithm from (14) using a task-based approach implemented in (55). Even for shared-memory CPU architectures, our new task-based implementation achieves a speed up of about 5× with respect to the current OpenMP implementation described in (56).

On the basis of these benchmarks, we estimate that Fugaku, among the current most powerful supercomputers in the world, would require around 14 hours to compute the Hafnian of a 100×100 matrix. Thus, if the total photon-number distribution of a given GBS setup has significant support past 100 photons, there will be a proportionally significant number of probability amplitudes that will require at least 14 hours in Fugaku to be computed.

We can get an estimate of the average time it would take to generate a sample by averaging the time it takes to generate a sample with n photons over the probability distribution of n photons. Using the same averaging procedure but applied to clicks instead of photons and assuming an overhead of 100 between computing probabilities and generating samples, the authors of (2) estimate that Fugaku would require around 1.9×10^{16} s to generate roughly

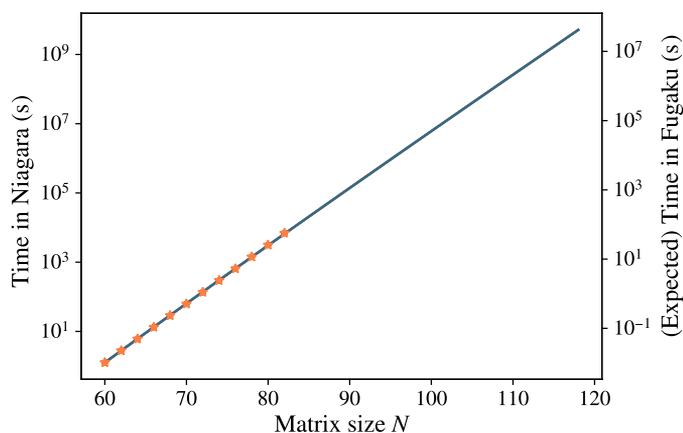


Fig. 4. The time cost of calculating a Hafnian of size n in double precision. The stars indicate actual sizes computed in the Niagara supercomputer (60). The blue line is a fit to $t_{\text{Niagara}}(n) = c_{\text{Niagara}} n^3 2^{n/2}$ with the only fitting parameter $c_{\text{Niagara}} = 5.42 \times 10^{-15}$ s. The standard deviation of fitting parameter c_{Niagara} is 1.2×10^{-16} s, which would give error bands thinner than the width of the line. We find an equivalent expected time in Fugaku, among the most powerful supercomputers, by considering the ratio of their R_{max} scores (maximal LINPACK performance achieved) giving their performance in number of floating point operations per second. The conversion factor between the left scale for Niagara and the right scale for Fugaku is the ratio of R_{max} values of Fugaku and Niagara, or equivalently, $c_{\text{Niagara}}/c_{\text{Fugaku}} = 122.8$. Note that since the computation of Hafnians can be broken into the independent calculation of an exponential number of summands (known as an embarrassingly parallel computation), this scaling is expected to be quite accurate.

the number of samples that their experiment produces in 200 s at megahertz clock speeds.

For the lossy instance considered in Fig. 3, we find that, on average, Fugaku would require $F_{\text{amplitudes/samples}} \sum_{n=0}^{n_{\text{max}}} p_{\text{lossy}}(n) c_{\text{Fugaku}} n^3 2^{n/2} \approx 4 \times 10^7$ s to generate one sample. In this estimate, we do not extend the sum to all possible photon numbers but only up to those that have a chance of more than 10^{-7} to occur, which happens at $n_{\text{max}} = 166$ and, moreover, assumes a reasonable overhead of $F_{\text{amplitudes/samples}} = 100$ for the calculation of probability amplitudes versus samples. As noted earlier, the complexity of generating a sample for a mixed or pure Gaussian state is proportional to that of calculating a probability amplitude (48) and the number of modes (in our case, 216), thus, using a factor of 100 is likely an underestimate.

To match the number of samples generated in seconds in a quantum device operating at 10 KHz would require 6.8×10^{15} s. Thus, the computational cost of an ($r = 0.8, a = 6, D = 3$)-high-dimensional GBS instance, with 3 dB of loss is on par with the expected classical complexity of the USTC experiment with the added advantage of being programmable and much closer to the collision-free regime: The expected classical complexity of an experiment, such as the one just described, is similar to the expected time complexity of the USTC experiment (2). However, besides the obvious disadvantage of programmability, their experiment is much farther away from the collision-free regime, in which computational complexity theoretic results guarantee the intractability of GBS.

For example, if the USTC experiment had been performed with PNR detectors, we would find that their photon number distribution has mean and standard distribution 83.3 ± 20.1 over 100 modes [where we assume the squeezing parameters quoted in (2) and a net transmission of $\eta = 0.3$]. Note that even within the first SD, one is

already beyond the total number of modes. This should be contrasted with a distribution such as the one in Fig. 3, for which we find 85.2 ± 13.9 over 216 modes.

TN methods

Another promising method to calculate the probability amplitude of high-dimensional GBS is using TN contractions. This has been the strategy of choice for classical adversaries to superconducting circuits performing RCS (52, 57).

In this section, we find that TN algorithms can simulate 2D lossy GBS experiments on 200 modes in a reasonable amount of time. This motivates going to a higher dimension, $D = 3$. We find that after making several allowances to the classical algorithm and accounting for tremendous improvements in classical hardware, one of the fastest supercomputers in the world, Fugaku, would take $\sim 10^{20}$ s to simulate a 3D experiment on 216 modes running for 200 s.

Any given quantum circuit can be written as a network of tensors such that each input quantum state is a rank-1 tensor, each gate acting on ℓ components is a rank- 2ℓ tensor, and each measurement operator is a rank-1 tensor (45). The probability amplitude for the quantum circuit can then be calculated by contracting the TN, i.e., by summing over all the indices of the TN. However, there are multiple different orderings (paths) in which the different indices of a TN can be contracted, which influence the contraction runtime. For some of the first classical benchmarking proposals of random circuits, the contraction paths were handpicked by the researchers (58). More recently, excellent randomized algorithms have been introduced to find contraction paths that have been shown to improve on previous results (45).

A second important practical consideration for TN contraction is that there is a trade-off between space and time complexity. That is, one can speed up substantially the contraction of a TN at the expense of assuming access to large amounts of memory. A systematic way to reduce the memory footprint of a TN contraction (at the expense of decreasing the speed of the computation) is to use a technique known as slicing, also known as variable projection or bond cutting (58).

Unlike for Hafnian methods where one does not need to specify much of the structure of the circuit, this information is vital in understanding the performance and limitations of TN simulations. As before, we fix the squeezing parameter $r = 0.8$ and assume net end-to-end transmission of $\eta = 0.5$. With these parameters and first assuming $D = 2$, we need at least $a = 14$ lattices sites per dimension to get to a mean photon number at the detectors (i.e., after loss) of $\mathbb{E}(n) \sim 80$. For a single-cycle $C = 1$, we use a TN contraction algorithm called cotengra (45) together with Fugaku's LINPACK benchmark to find that this supercomputer would require less than 100 μ s to contract the TN. Thus, for 2D instances up to this size, it is necessary to consider more than one cycle, implying the construction of either D extra delay lines or adding a circulator, both of which will adversely affect the net transmission.

This motivates considering the next dimension, $D = 3$. For this case, and fixing the number of cycles to $C = 1$, we find that we need at least $a = 6$ to have a mean photon number on the order of 80 at the detectors, which would provide a nontrivial support on photon numbers that are beyond the reach of the Hafnian algorithms described above. In Table 1, we show the time it would take Fugaku to contract different 3D GBS circuits for different lattice sizes.

Table 1. Benchmarks for a $D = 3$ high-dimensional GBS instance with minimal Fock space cutoff $c = 4$. The first column gives the number of lattice points, from which the number of modes follows $M = a^3$. The second column is the expected run time in Fugaku. This time is obtained by estimating the number of floating point operations required to contract the tensor using cotengra (45) and converting this into a time by using the R_{\max} floating point operation per second score for Fugaku. Note that cotengra implements randomized algorithms; thus, for each problem size, we run it 200 times and confirm that, after the first 100 runs, there is no significant variation in the best score found. The last column gives the number of elements of the largest tensor ever needed to be stored in memory during the contraction. Note that this places restrictions on the RAM available in each of the nodes of a supercomputer. In particular, the nodes in Fugaku have up to 32 gigabyte of RAM, allowing to store on the order of 4×10^9 64-bit floating point numbers; thus, an $a = 6$ instance will far exceed the required capacity of a single node requiring distributed storage and thus subsequent hit in efficiency due to communication complexity.

Number of lattice points (a)	Expected time in Fugaku (s)	Size of the largest tensor
4	1.65×10^{-1}	4.39×10^{12}
5	4.56×10^2	4.61×10^{18}
6	2.11×10^{14}	7.92×10^{28}

Note that even allowing for a hypothetical scenario in which the random access memory (RAM) of each of its nodes has been expanded by about 19 orders of magnitude, it would take Fugaku on the order 2.11×10^{14} s to calculate a contraction with a minimal (and highly inaccurate) cutoff of 4. In reality, it is infeasible to fit the computation in the memory or even the hard disks of individual nodes, so slicing would be required, which can lead to astronomical overheads over this idealized estimate. Even without this overhead and assuming that generating a sample is as expensive as calculating a probability, simulating a 200-s 10-kHz experiment would require over 4×10^{20} s. Of course, we remind the reader once more that a direct calculation of output probabilities is not what the experiment does but only what one model of the experiment, and there may be more efficient methods for simulating a verifiable experiment.

On the basis of the evidence presented above, a high-dimensional GBS instance with squeezing parameter $r = 0.8$, in $D = 3$ dimensions, with $a = 6$ modes per dimension or a total of 216 modes and a single cycle $C = 1$ is well beyond the capabilities of current simulation methods based either on Hafnian calculations or TN contractions, even when losses of around 3 dB ($\eta \sim 0.5$) are present. This immense computational gap is present even after the fact that we allow the classical computer to ignore substantial overheads in terms of cutoff, number of modes, and samples-to-amplitudes conversion. These experimental parameters we propose are within the reach of current photonics technology, and their implementation using time-domain multiplexing can be achieved with a significantly reduced number of components. Note that after this submission, we became aware of a recent work (59) on an upgraded version of the experiment done in (2).

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <https://science.org/doi/10.1126/sciadv.abi7894>

REFERENCES AND NOTES

1. 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. M. Ewen, 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, J. M. Martinis, Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
2. 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, J.-W. Pan, Quantum computational advantage using photons. *Science* **370**, 1460–1463 (2020).
3. E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, R. Wisnieff, Leveraging secondary storage to simulate deep 54-qubit sycamore circuits. arXiv:1910.09534 (2019).
4. G. Kalai, G. Kindler, Gaussian noise sensitivity and BosonSampling. arXiv:1409.3093 (2014).
5. J. J. Renema, Marginal probabilities in boson samplers with arbitrary input states. arXiv:2012.14917 (2020).
6. S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, H. Neven, Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **14**, 595–600 (2018).
7. S. Aaronson, A. Arkhipov, The computational complexity of linear optics. *Theory Comput.* **9**, 143–252 (2013).
8. J. E. Bourassa, R. N. Alexander, M. Vasmer, A. Patil, I. Tzitrin, T. Matsuura, D. Su, B. Q. Baragiola, S. Guha, G. Dauphinais, K. K. Sabapathy, N. C. Menicucci, I. Dhand, Blueprint for a scalable photonic fault-tolerant quantum computer. *Quantum* **5**, 392 (2021).
9. S. Bartolucci, P. Birchall, H. Bombin, H. Cable, C. Dawson, M. Gimeno-Segovia, E. Johnston, K. Kieling, N. Nickerson, M. Pant, F. Pastawski, T. Rudolph, C. Sparrow, Fusion-based quantum computation. arXiv:2101.09310 (2021).
10. C. S. Hamilton, R. Kruse, L. Sansoni, S. Barkhofen, C. Silberhorn, I. Jex, Gaussian boson sampling. *Phys. Rev. Lett.* **119**, 170501 (2017).
11. R. Kruse, C. S. Hamilton, L. Sansoni, S. Barkhofen, C. Silberhorn, I. Jex, detailed study of Gaussian boson sampling. *Phys. Rev. A* **100**, 032326 (2019).
12. A. Bouland, B. Fefferman, C. Nirkhe, U. Vazirani, On the complexity and verification of quantum random circuit sampling. *Nat. Phys.* **15**, 159–163 (2019).
13. A. Bouland, B. Fefferman, Z. Landau, Y. Liu, Noise and the frontier of quantum supremacy. arXiv:2102.01738 (2021).
14. A. Björklund, B. Gupt, N. Quesada, A faster hafnian formula for complex matrices and its benchmarking on a supercomputer. *ACM J. Exp. Algorithmics* **24**, 1–17 (2019).
15. L. Stockmeyer, in *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing* (ACM, 1983), pp. 118–126.
16. L. Valiant, The complexity of computing the permanent. *Theor. Comput. Sci.* **8**, 189–201 (1979).
17. T. Jiang, The entries of circular orthogonal ensembles. *J. Math. Phys.* **50**, 063302 (2009).
18. T. Jiang, How many entries of a typical orthogonal matrix can be approximated by independent normals? *Ann. Probab.* **34**, 1497–1529 (2006).
19. T. Jiang, Y. Ma, Distances between random orthogonal matrices and independent normals. *Trans. Am. Math. Soc.* **372**, 1509–1553 (2017).
20. Y. Kondo, R. Mori, R. Movassagh, Fine-grained analysis and improved robustness of quantum supremacy for random circuit sampling. arXiv:2102.01960 (2021).
21. R. Movassagh, Quantum supremacy and random circuits. arXiv:1909.06210 (2019).
22. K. Fujii, Noise threshold of quantum supremacy. arXiv:2102.01738 (2016).
23. M. Ozmaniec, D. J. Brod, Classical simulation of photonic linear optics with lost particles. *New J. Phys.* **20**, 092002 (2018).
24. E. Knill, Quantum gates using linear optics and postselection. *Phys. Rev. A* **66**, 052306 (2002).
25. S. Rahimi-Keshari, A. Scherer, A. Mann, A. T. Rezakhanlou, A. I. Lvovsky, B. C. Sanders, Quantum process tomography with coherent states. *New J. Phys.* **13**, 013006 (2011).
26. H. Qi, D. J. Brod, N. Quesada, R. García-Patrón, Regimes of classical simulability for noisy Gaussian boson sampling. *Phys. Rev. Lett.* **124**, 100502 (2020).
27. R. García-Patrón, J. J. Renema, V. Shchesnovich, Simulating boson sampling in lossy architectures. *Quantum* **3**, 169 (2019).
28. H. Qi, D. Cifuentes, K. Brádler, R. Israel, T. Kalajdzievski, N. Quesada, Efficient sampling from shallow gaussian quantum-optical circuits with local interactions. arXiv:2009.11824 (2020).
29. J. Bermejo-Vega, D. Hangleiter, M. Schwarz, R. Raussendorf, J. Eisert, Architectures for quantum simulation showing a quantum speedup. *Phys. Rev. X* **8**, 021010 (2018).

30. K. R. Motes, A. Gilchrist, J. P. Dowling, P. P. Rohde, Scalable boson sampling with time-bin encoding using a loop-based architecture. *Phys. Rev. Lett.* **113**, 120501 (2014).
31. T. Nitsche, S. Barkhofen, R. Kruse, L. Sansoni, M. Štefaňák, A. Gábris, V. Potoček, T. Kiss, I. Jex, C. Silberhorn, Probing measurement-induced effects in quantum walks via recurrence. *Sci. Adv.* **4**, eaar6444 (2018).
32. C. Tabillione, R. van der Meer, H. J. Snijders, P. Hooijschuur, J. P. Epping, M. de Goede, B. Kassenberg, P. Venderbosch, C. Toebes, H. van den Vlekkert, P. W. H. Pinkse, J. J. Renema, A 12-mode universal photonic processor for quantum information processing. *Mater. Quantum. Technol.* **1**, 035002 (2020).
33. K. R. Motes, J. P. Dowling, A. Gilchrist, P. P. Rohde, Implementing boson sampling with time-bin encoding: Analysis of loss, mode mismatch, and time jitter. *Phys. Rev. A* **92**, 052319 (2015).
34. C. Oh, Y. Lim, B. Fefferman, L. Jiang, Classical simulation of bosonic linear-optical random circuits beyond linear light cone. arXiv:2102.10083 (2021).
35. Y. He, X. Ding, Z.-E. Su, H.-L. Huang, J. Qin, C. Wang, S. Unsleber, C. Chen, H. Wang, Y.-M. He, X.-L. Wang, W.-J. Zhang, S.-J. Chen, C. Schneider, M. Kamp, L.-X. You, Z. Wang, S. Höfling, C.-Y. Lu, J.-W. Pan, Time-bin-encoded boson sampling with a single-photon device. *Phys. Rev. Lett.* **118**, 190501 (2017).
36. H. Wang, J. Qin, X. Ding, M.-C. Chen, S. Chen, X. You, Y.-M. He, X. Jiang, L. You, Z. Wang, C. Schneider, J.-J. Renema, S. Höfling, C.-Y. Lu, J.-W. Pan, Boson sampling with 20 input photons and a 60-mode interferometer in a 10^{14} -Dimensional hilbert space. *Phys. Rev. Lett.* **123**, 250503 (2019).
37. A. Schreiber, A. Gábris, P. P. Rohde, K. Laiho, M. Štefaňák, V. Potoček, C. Hamilton, I. Jex, C. Silberhorn, A 2d quantum walk simulation of two-particle dynamics. *Science* **336**, 55–58 (2012).
38. L. Lorz, E. Meyer-Scott, T. Nitsche, V. Potoček, A. Gábris, S. Barkhofen, I. Jex, C. Silberhorn, Photonic quantum walks with four-dimensional coins. *Phys. Rev. Res.* **1**, 033036 (2019).
39. M. V. Larsen, X. Guo, C. R. Breum, J. S. Neergaard-Nielsen, U. L. Andersen, Deterministic multi-mode gates on a scalable photonic quantum computing platform. *Nat. Phys.* **17**, 1018–1023 (2021).
40. J. M. Arrazola, V. Bergholm, K. Brádler, T. R. Bromley, M. J. Collins, I. Dhand, A. Fumagalli, T. Gerrits, A. Goussev, L. G. Helt, J. Hundal, T. Isacsson, R. B. Israel, J. Izaac, S. Jahangiri, R. Janik, N. Killoran, S. P. Kumar, J. Lavoie, A. E. Lita, D. H. Mahler, M. Menotti, B. Morrison, S. W. Nam, L. Neuhaus, H. Y. Qi, N. Quesada, A. Reppingon, K. K. Sabapathy, M. Schuld, D. Su, J. Swinarton, A. Száva, K. Tan, P. Tan, V. D. Vaidya, Z. Vernon, Z. Zabaneh, Y. Zhang, Quantum circuits with many photons on a programmable nanophotonic chip. *Nature* **591**, 54–60 (2021).
41. S. Takeda, K. Takase, A. Furusawa, On-demand photonic entanglement synthesizer. *Sci. Adv.* **5**, eaaw4530 (2019).
42. I. Dhand, M. Engelkemeier, L. Sansoni, S. Barkhofen, C. Silberhorn, M. B. Plenio, Proposal for quantum simulation via all-optically-generated tensor network states. *Phys. Rev. Lett.* **120**, 130501 (2018).
43. D. J. Brod, Complexity of simulating constant-depth Boson Sampling. *Phys. Rev. A* **91**, 042316 (2015).
44. N. Quesada, J. M. Arrazola, Exact simulation of gaussian boson sampling in polynomial space and exponential time. *Phys. Rev. Res.* **2**, 023005 (2020).
45. J. Gray, S. Kourtis, Hyper-optimized tensor network contraction. *Quantum* **5**, 410 (2021).
46. D. Hangleiter, J. Bermejo-Vega, M. Schwarz, J. Eisert, Anti-concentration theorems for schemes showing a quantum computational speedup. *Quantum* **2**, 65 (2018).
47. M. V. Larsen, X. Guo, C. R. Breum, J. S. Neergaard-Nielsen, U. L. Andersen, Deterministic generation of a two-dimensional cluster state. *Science* **366**, 369–372 (2019).
48. N. Quesada, R. S. Chadwick, B. A. Bell, J. M. Arrazola, T. Vincent, H. Qi, R. García-Patrón, Quadratic speedup for simulating gaussian boson sampling. arXiv:2010.15595 (2020).
49. Z. Vernon, N. Quesada, M. Liscidini, B. Morrison, M. Menotti, K. Tan, J. Sipe, Scalable squeezed-light source for continuous-variable quantum sampling. *Phys. Rev. Appl.* **12**, 064024 (2019).
50. V. Dodonov, O. Man'ko, V. Man'ko, Photon distribution for one-mode mixed light with a generic gaussian wigner function. *Phys. Rev. A* **49**, 2993–3001 (1994).
51. N. Quesada, L. G. Helt, J. Izaac, J. M. Arrazola, R. Shahrokhshahi, C. R. Myers, K. K. Sabapathy, Simulating realistic non-gaussian state preparation. *Phys. Rev. A* **100**, 022341 (2019).
52. I. L. Markov, A. Fatima, S. V. Isakov, S. Boixo, Quantum supremacy is both closer and farther than it appears. arXiv:1807.10749 (2018).
53. J. Wu, Y. Liu, B. Zhang, X. Jin, Y. Wang, H. Wang, X. Yang, A benchmark test of boson sampling on Tianhe-2 supercomputer. *Natl. Sci. Rev.* **5**, 715–720 (2018).
54. Y. Li, M. Chen, Y. Chen, H. Lu, L. Gan, C. Lu, J. Pan, H. Fu, G. Yang, Benchmarking 50-photon gaussian boson sampling on the sunway TaihuLight. arXiv:2009.01177 (2020).
55. T.-W. Huang, D.-L. Lin, Y. Lin, C.-X. Lin, Cpp-taskflow: A general-purpose parallel and heterogeneous task programming system at scale. arXiv:2004.10908 (2020).
56. B. Gupt, J. Izaac, N. Quesada, The walrus: A library for the calculation of hafnians, hermite polynomials and gaussian boson sampling. *J. Open Source Softw.* **4**, 1705 (2019).
57. B. Villalonga, D. Lyakh, S. Boixo, H. Neven, T. S. Humble, R. Biswas, E. G. Rieffel, A. Ho, S. Mandrà, Establishing the quantum supremacy frontier with a 281 pflop/s simulation. *Quantum Sci. Technol.* **5**, 034003 (2020).
58. B. Villalonga, S. Boixo, B. Nelson, C. Henze, E. Rieffel, R. Biswas, S. Mandrà, A flexible high-performance simulator for verifying and benchmarking quantum circuits implemented on real hardware. *nph Quant. Inf.* **5**, 86 (2019).
59. H.-S. Zhong, Y.-H. Deng, J. Qin, H. Wang, M.-C. Chen, L.-C. Peng, Y.-H. Luo, D. Wu, S.-Q. Gong, H. Su, 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, J. Renema, C.-Y. Lu, J.-W. Pan, Phase-programmable gaussian boson sampling using stimulated squeezed light. *Phys. Rev. Lett.* **237**, 180502 (2021).
60. M. Ponce, R. van Zon, S. Northrup, D. Gruner, J. Chen, F. Ertinaz, A. Fedoseev, L. Groer, F. Mao, B. C. Mundim, et al., Deploying a top-100 supercomputer for large parallel workloads: The Niagara supercomputer. *Proceedings of the practice and experience in advanced research computing on rise of the machines (learning)* (2019), pp. 1–8.

Acknowledgments: We thank J. M. Arrazola, L. G. Helt, M. Collins, F. Laudenbach, I. Tzitrin, and Z. Vernon for helpful discussions. We also thank Bouland and colleagues (13) for sharing an early version of their manuscript. M.H., M.L., and D.H. thank K. Zyczkowski for enlightening discussions regarding the distribution of COE submatrices. **Funding:** A.M. is funded by Mitacs Accelerate Program. M.H., M.L., J.E., and D.H. are funded by the DFG (EI 519/21-1, EI 519/9-1, EI 519/14-1, and CRC 183), the MATH+ Cluster of Excellence, the BMBF (HYBRID), the Einstein Research Foundation (Einstein Research Unit on near-term quantum devices), BMBF (QPIC), BMBF (PhoQuant), and the European Union's Horizon 2020 research and innovation programme under grant agreement no. 817482 (PASQuanS). The authors thank SOSCIP and SciNet for computational resources. Computations have been performed on the Niagara and the Mist supercomputers at the SciNet-SOSCIP HPC Consortium. SciNet is funded by The Canada Foundation for Innovation, the government of Ontario; Ontario Research Fund–Research Excellence; and the University of Toronto. SOSCIP is funded by the Federal Economic Development Agency of Southern Ontario, the province of Ontario, IBM Canada Ltd., Ontario Centres of Excellence, Mitacs, and Ontario academic member institutions. **Author contributions:** Theory work was completed by A.D., A.M., M.H., M.L., H.Q., J.E., D.H., and B. F. T.V., N.Q., L.M., and J.L. completed the experimental design and benchmarking work. I.D. managed the project and made contributions to both the theory and benchmarking work. All authors discussed the results and contributed to writing the manuscript. **Competing interests:** B.F. and A.D. acted as paid consultants for Xanadu Quantum Technologies while parts of this work were being performed. The authors declare that they have no other competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials.

Submitted 31 March 2021
Accepted 12 November 2021
Published 5 January 2022
10.1126/sciadv.abi7894