Exact simulation of Gaussian boson sampling in polynomial space and exponential time

Nicolás Quesada* and Juan Miguel Arrazola
 Xanadu, 777 Bay Street, Toronto, Canada

(Received 12 September 2019; revised manuscript received 14 January 2020; accepted 5 March 2020; published 3 April 2020)

We introduce an exact classical algorithm for simulating Gaussian boson sampling (GBS). The complexity of the algorithm is exponential in the number of photons detected, which is itself a random variable. For a fixed number of modes, the complexity is in fact equivalent to that of calculating output probabilities, up to constant prefactors. The simulation algorithm can be extended to other models such as GBS with threshold detectors, GBS with displacements, and sampling linear combinations of Gaussian states. In the specific case of encoding non-negative matrices into a GBS device, our method leads to an approximate sampling algorithm with polynomial runtime. We implement the algorithm, making the code publicly available as part of Xanadu's The Walrus library and benchmark its performance on GBS with random Haar interferometers and with encoded Erdős-Renyi graphs.

DOI: [10.1103/PhysRevResearch.2.023005](https://doi.org/10.1103/PhysRevResearch.2.023005)

I. INTRODUCTION

Boson sampling is a model of photonic quantum computing that was introduced to argue that nonuniversal photonic quantum computers cannot be efficiently simulated classically [1]. Since then, significant work has been done to pursue its practical implementation [2–5] and to design variants that are more amenable to existing experimental techniques [6–9]. Recently, Gaussian boson sampling (GBS) has emerged as a new paradigm that addresses some of the major challenges in scaling boson sampling devices [10,11]. Instead of employing indistinguishable single photons as inputs, GBS prepares a multimode Gaussian state that is subsequently measured by using photon detectors [12]. The Gaussian state is typically obtained by sending squeezed light through a linear-optical interferometer, while more general versions employ displacements together with squeezing operations. GBS has raised additional interest due to the discovery of applications to quantum chemistry [13], optimization [14–16], graph similarity [17], and point processes [18]. Initial experimental implementations have also been recently reported [19–21].

The complexity arguments underlying the hardness of classically simulating GBS are only valid asymptotically. In practice, it is necessary to perform a comparison with state-of-the-art simulation algorithms to understand the actual advantages of employing a quantum computer. Considerable progress has been made in developing classical simulation algorithms for the original boson sampling model. Reference [22] reported the first simulation method by describing an approximate

Markov chain Monte Carlo algorithm for boson sampling. These results were improved in Ref. [23], where an exact sampling algorithm was developed with the same asymptotic complexity but better constant prefactors. An algorithm for other models of boson sampling has also been recently reported [24]. As stated in Ref. [25], extending these techniques to GBS has been challenging because these algorithms rely on specific properties of boson sampling that are not present in GBS. Despite these challenges, an exact simulation algorithm has been reported and implemented for the specific case of GBS with threshold detectors [26,27]. This algorithm suffers from the critical drawback that its memory requirement also scales exponentially, limiting the scope of problems that can be simulated.

In this work, we introduce an exact simulation algorithm for GBS with efficient space complexity and applicable to all versions of GBS. The worst-case time complexity of the algorithm is linear in the number of modes and exponential in the number of detected photons. The core strategy is to employ the chain rule of probability to sequentially sample the number of photons in each mode conditioned on the results from previous modes. Conditional sampling is particularly well suited to Gaussian states because it is straightforward to describe the marginalized density matrices of any number of subsystems using covariance matrices and a vector of means. The method presented here can also be applied to several other models, including GBS with threshold detectors, GBS with displacements, and sampling linear combinations of Gaussian states. We also show that in the special case of encoding non-negative matrices into a GBS device, the algorithm results in an approximate sampling procedure with polynomial runtime. Finally, we implement the algorithm and benchmark its performance on random interferometers drawn from the Haar measure. We also benchmark the approximate algorithm for random Erdős-Renyi graphs. The code used to generate these results is freely available as part of Xanadu's The Walrus library [28].

*nicolas@xanadu.ai

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

II. CLASSICAL GAUSSIAN BOSON SAMPLING ALGORITHM

In this section we describe and study the algorithm for GBS simulation. For completeness, we begin with a brief description of GBS. We continue by outlining the algorithm in detail, performing a theoretical analysis of its complexity, and discussing its scope of application.

A. Gaussian boson sampling

The quantum state ϱ of a system of m bosonic modes can be uniquely specified by its Wigner function $W(\mathbf{q}, \mathbf{p})$ [29,30], where $\mathbf{q} \in \mathbb{R}^m$ are the canonical positions and $\mathbf{p} \in \mathbb{R}^m$ are the canonical momenta. Gaussian states are defined simply as the set of states with Gaussian Wigner functions. They can be uniquely described by a covariance matrix $\mathbf{V}$ and a vector of means $\bar{\mathbf{q}}, \bar{\mathbf{p}}$. It is convenient to write the covariance matrix in terms of the complex amplitudes $\boldsymbol{\alpha} = \frac{1}{\sqrt{2\hbar}}(\mathbf{q} + i\mathbf{p}) \in \mathbb{C}^m$. The variables $\boldsymbol{\alpha}$ are complex-normal distributed with mean $\bar{\boldsymbol{\alpha}} = \frac{1}{\sqrt{2\hbar}}(\bar{\mathbf{q}} + i\bar{\mathbf{p}}) \in \mathbb{C}^m$ and covariance matrix $\boldsymbol{\Sigma}$ [31].

GBS is a model of photonic quantum computing where a Gaussian state is measured in the Fock basis. A general pure Gaussian state can be prepared by using single-mode squeezing and displacement operations in combination with linear-optical interferometry [30,32–34]. It was shown in Ref. [10] that, when the modes of a Gaussian state with zero mean ($\bar{\boldsymbol{\alpha}} = 0$) are measured, the probability of obtaining a pattern of photons $S = (s_1, \dots, s_m)$, where s_i is the number of photons in mode i , is given by

$$p(S) = \frac{1}{\sqrt{\det(\mathbf{Q})}} \frac{\text{Haf}(\mathbf{A}_S)}{s_1! \cdots s_m!}, \quad (1)$$

where

$$\mathbf{Q} := \boldsymbol{\Sigma} + \mathbb{1}/2, \quad (2)$$

$$\mathbf{A} := \mathbf{X}(\mathbb{1} - \mathbf{Q}^{-1}), \quad (3)$$

$$\mathbf{X} := \begin{bmatrix} 0 & \mathbb{1} \\ \mathbb{1} & 0 \end{bmatrix}, \quad (4)$$

and $\mathbf{A}_S$ is the matrix obtained by repeating rows and columns i and $i + m$ of the $2m \times 2m$ symmetric matrix $\mathbf{A}$. We refer to $\mathbf{A}$ as the *kernel* matrix. If $s_i = 0$, rows and columns i and $i + m$ are deleted from $\mathbf{A}$; if $s_i > 0$, the rows and columns are repeated s_i times. For example, consider $m = 3$ modes and let

$$\mathbf{A} = \begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,3} & a_{1,4} & a_{1,5} & a_{1,6} \\ a_{2,1} & a_{2,2} & a_{2,3} & a_{2,4} & a_{2,5} & a_{2,6} \\ a_{3,1} & a_{3,2} & a_{3,3} & a_{3,4} & a_{3,5} & a_{3,6} \\ a_{4,1} & a_{4,2} & a_{4,3} & a_{4,4} & a_{4,5} & a_{4,6} \\ a_{5,1} & a_{5,2} & a_{5,3} & a_{5,4} & a_{5,5} & a_{5,6} \\ a_{6,1} & a_{6,2} & a_{6,3} & a_{6,4} & a_{6,5} & a_{6,6} \end{bmatrix}. \quad (5)$$

For the photon pattern $S = (3, 0, 1)$, one has

$$\mathbf{A}_S = \begin{bmatrix} a_{1,1} & a_{1,1} & a_{1,1} & a_{1,3} & a_{1,4} & a_{1,4} & a_{1,4} & a_{1,6} \\ a_{1,1} & a_{1,1} & a_{1,1} & a_{1,3} & a_{1,4} & a_{1,4} & a_{1,4} & a_{1,6} \\ a_{1,1} & a_{1,1} & a_{1,1} & a_{1,3} & a_{1,4} & a_{1,4} & a_{1,4} & a_{1,6} \\ a_{3,1} & a_{3,1} & a_{3,1} & a_{3,3} & a_{3,4} & a_{3,4} & a_{3,4} & a_{3,6} \\ a_{4,1} & a_{4,1} & a_{4,1} & a_{4,3} & a_{4,4} & a_{4,4} & a_{4,4} & a_{4,6} \\ a_{4,1} & a_{4,1} & a_{4,1} & a_{4,3} & a_{4,4} & a_{4,4} & a_{4,4} & a_{4,6} \\ a_{4,1} & a_{4,1} & a_{4,1} & a_{4,3} & a_{4,4} & a_{4,4} & a_{4,4} & a_{4,6} \\ a_{6,1} & a_{6,1} & a_{6,1} & a_{6,3} & a_{6,4} & a_{6,4} & a_{6,4} & a_{6,6} \end{bmatrix}. \quad (6)$$

Note that the first and fourth rows and columns of the original matrix are repeated three times, the second and fifth columns disappeared, and the third and sixth column of $\mathbf{A}$ appear only once. Finally, the matrix function $\text{Haf}(\cdot)$ is the hafnian [35], which for a $2m \times 2m$ matrix $\mathbf{A}$ is defined as

$$\text{Haf}(\mathbf{A}) = \sum_{\mu \in \text{PMP}} \prod_{(i,j) \in \mu} a_{i,j}, \quad (7)$$

where $a_{i,j}$ are the entries of $\mathbf{A}$ and PMP is the set of perfect matching permutations, namely, the possible ways of partitioning the set $\{1, \dots, 2m\}$ into disjoint subsets of size two. The hafnian is #P-Hard to approximate for worst-case instances [36] and the runtime of the best known algorithms for computing hafnians scales exponentially with the dimension of the input matrix [37]. The hardness of computing hafnians has been leveraged to show that, under the validity of specific technical conjectures, sampling from the output distribution of a GBS device cannot be performed efficiently using classical computers [1,10].

B. Algorithm

Ideally, a classical sampling algorithm for GBS will have the following properties: (i) sampling from the GBS distribution is *exact*, (ii) space complexity is polynomial, and (iii) time complexity is proportional to the complexity of computing output probabilities. We now describe an algorithm satisfying all these properties. The main strategy is to apply the definition of conditional probability $p(s_k | s_{k-1}, \dots, s_1) = p(s_k, s_{k-1}, \dots, s_1) / p(s_{k-1}, \dots, s_1)$ to sequentially sample each mode conditioned on outcomes from previous modes. When sampling the k th mode, the probability $p(s_k, s_{k-1}, \dots, s_1)$ can be calculated from the reduced state of the first k modes, which can be efficiently computed for Gaussian states. This in turn involves calculating the hafnian of a matrix whose size depends on the number of photons detected so far. The probability $p(s_{k-1}, \dots, s_1)$ is calculated from the previous step.

Formally, let $\mathbf{V}^{(k)}$ denote the reduced covariance matrix of the first k modes; this is simply the submatrix of $\mathbf{V}$ obtained by keeping rows and columns 1 to k and $m+1$ to $m+k$. From the reduced covariance matrix we can construct $\mathbf{Q}^{(k)} = \mathbf{V}^{(k)} + \mathbb{1}/2$, $\mathbf{O}^{(k)} = \mathbb{1} - (\mathbf{Q}^{(k)})^{-1}$, and $\mathbf{A}^{(k)} = \mathbf{X}\mathbf{O}^{(k)}$. The probability of observing a partial pattern $S^{(k)} = (s_1, s_2, \dots, s_k)$ of photons in the first k modes is given by

$$p(S^{(k)}) = \frac{1}{\sqrt{\det(\mathbf{Q}^{(k)})}} \frac{\text{Haf}(\mathbf{A}_S^{(k)})}{s_1! \cdots s_k!}. \quad (8)$$

The algorithm is specified as follows:

(1) Compute $\mathbf{A}^{(1)}$ and sample the number of photons s_1 in the first mode from the distribution

$$p(s_1) = \frac{1}{\sqrt{\det(\mathbf{Q}^{(1)})}} \frac{\text{Haf}(\mathbf{A}_{s_1}^{(1)})}{s_1!}.$$

This is done by computing each probability $p(s_1 = 0), p(s_1 = 1), \dots, p(s_1 = n_{\max})$ up to a maximum photon number $n_{\max}$ and sampling from the resulting distribution. The parameter $n_{\max}$ must be chosen to ensure that the resulting distribution is sufficiently close to normalized, as discussed in Appendix B. Let s_1^* denote the output of this step.

(2) Compute $\mathbf{A}^{(2)}$ and, as before, sample the number of photons s_2 in the second mode from the conditional distribution

$$p(s_2|s_1^*) = \frac{p(s_1^*, s_2)}{p(s_1^*)}, \quad (9)$$

where s_1^* is fixed from the result of the previous step and $p(s_1^*, s_2)$ is defined as in Eq. (8).

(3) Repeat this procedure for all modes. The conditional distribution of the photon number s_k for the k th mode is

$$p(s_k|s_1^*, \dots, s_{k-1}^*) = \frac{p(s_1^*, s_2^*, \dots, s_{k-1}^*, s_k)}{p(s_1^*, \dots, s_{k-1}^*)}. \quad (10)$$

The outputs $(s_1^*, \dots, s_{k-1}^*)$ are fixed from the results of the previous steps, $p(s_1^*, s_2^*, \dots, s_{k-1}^*, s_k)$ is given by Eq. (8), and the probability $p(s_1^*, \dots, s_{k-1}^*)$ that is necessary to sample at step k has already been computed at step $k-1$.

The correctness of the algorithm follows from the chain rule of probability

$$p(s_1, s_2, \dots, s_m) = \prod_{k=1}^m p(s_k|s_1, s_2, \dots, s_{k-1}), \quad (11)$$

which implies that the algorithm performs exact sampling from the GBS distribution.

C. Complexity

The algorithm has $O(m^2)$ space complexity: the largest memory requirement arises from the need to store the $2m \times 2m$ matrix $\mathbf{A}$. Therefore, the algorithm is memory efficient. To determine the time complexity, consider step k where the goal is to sample photons from the k th mode. Assume that n photons have been sampled in the previous $k-1$ modes. To generate a sample, the algorithm computes the $n_{\max}$ conditional probabilities of observing $0, 1, \dots, n_{\max}$ photons in mode k . This in turn requires calculating the hafnian of matrices with dimensions $2n, 2(n+1), \dots, 2(n+n_{\max})$. State-of-the-art algorithms for evaluating hafnians [37] have complexity $O(n^3 2^n)$, where $2n$ is the dimension of the matrix. Therefore, having detected n photons, the time $t(n)$ required to sample one additional mode is

$$t(n) = O\left(\sum_{j=0}^{n_{\max}} (n+j)^3 2^{(n+j)}\right) = O(n^3 2^n). \quad (12)$$

After each new step, the number of photons detected at that stage of the algorithm may either increase if more photons are observed, or stay the same if no more photons are detected.

Let μ_n denote the number of steps for which exactly n photons were detected at that stage of the algorithm. Note that $\mu_n \leq m$. To produce a sample of N photons in total, the algorithm therefore requires time

$$T(N, m) = \sum_{n=0}^N \mu_n t(n) = O(mN^3 2^N). \quad (13)$$

The largest matrix whose hafnian needs to be computed has dimension precisely $2(N + n_{\max})$, so computing its hafnian takes time $(N + n_{\max})^3 2^{N+n_{\max}} = O(N^3 2^N)$. Thus, up to constant prefactors, the complexity of generating one sample from the algorithm is roughly m times that of computing one output probability.

The complexity of the algorithm depends on how many photons have been sampled previously: two samples with the same number of photons can take different amounts of time to be generated. The two extreme cases occur when most of the photons are detected in the first few modes (longest time) and when they are detected in the last few modes (shortest time). When N photons are measured in the first N modes, the algorithm must compute $m - N$ hafnians of matrices with maximum dimension $2(N + n_{\max})$. Conversely, if the N photons are detected in the last N modes, the algorithm calculates hafnians of the smallest possible dimension $2n_{\max}$ in the first $m - N$ modes.

Although asymptotically the best algorithms for computing hafnians of matrices of size $2n \times 2n$ have complexity $O(N^3 2^N)$, the constant factors are sufficiently large that, for all values of N for which hafnians can be computed in practice, we note that it is in fact preferable to employ the algorithm of Ref. [38], whose runtime scales as $O(N^5 2^N)$.

D. Scope of application

The sampling algorithm relies on an important property of Gaussian states: the marginal of a Gaussian state is another Gaussian state whose covariance matrix and vector of means can be efficiently computed from those of the larger state. This is a special feature of Gaussian states; indeed, computing reduced states generally requires exponential time. Using this observation, we generalize the sampling algorithm to several variants of GBS where marginals can be efficiently computed.

1. Threshold detectors

The first case we consider is GBS with threshold detectors, where output probabilities are given by

$$p(S^{(k)}) = \frac{\text{Tor}(\mathbf{O}_s^{(k)})}{\sqrt{\det(\mathbf{Q}^{(k)})}}, \quad (14)$$

where as before $\mathbf{O}^{(k)} = \mathbb{1} - (\mathbf{Q}^{(k)})^{-1}$, $\text{Tor}(\cdot)$ is the *torontonians* [26], and the entries of $S^{(k)}$ take the values $s_i = 0$ for no detection and $s_i = 1$ for detection. The best known algorithms for computing torontonians have the same complexity as those for calculating hafnians [26].

2. Displacements

The algorithm can also be applied to GBS with displacements, i.e., when the Gaussian state has nonzero mean $\bar{\alpha} \neq 0$. We first define the following useful quantities:

$$\bar{\alpha}^{(k)} = (\bar{\alpha}^{(k)}, [\bar{\alpha}^{(k)}]^*), \quad (15)$$

$$\bar{\gamma}^{(k)} = [\mathbf{Q}^{(k)}]^{-1} [\bar{\alpha}^{(k)}]^\dagger, \quad (16)$$

$$\mathcal{N}^{(k)} = \frac{\exp\left(-\frac{1}{2}\bar{\alpha}^{(k)}[\mathbf{Q}^{(k)}]^{-1}[\bar{\alpha}^{(k)}]^\dagger\right)}{\sqrt{\det(\mathbf{Q}^{(k)})} s_1! \cdots s_k!}. \quad (17)$$

As shown in Refs. [37,39,40], for GBS with displacement, the output probabilities are given by

$$p(s) = \mathcal{N}^{(k)} \times \text{lhaf}\{\text{filldiag}(\mathbf{A}_s^{(k)}, \bar{\gamma}_s^{(k)})\}, \quad (18)$$

where $\text{lhaf}(\cdot)$ is the *loop hafnian* introduced in Ref. [37], $\bar{\gamma}_s$ is the vector obtained from $\bar{\gamma}$ by repeating the i and $i+m$ entries of $\bar{\gamma}^{(k)}$ a total of s_i times, and the function $\text{filldiag}(\mathbf{A}_s^{(k)}, \bar{\gamma}_s^{(k)})$ replaces the diagonal of $\mathbf{A}_s^{(k)}$ with the vector $\bar{\gamma}_s^{(k)}$. The loop hafnian has the same complexity as the hafnian, so its sampling algorithm has the same complexity as GBS without displacement.

3. Linear combinations of Gaussian states

The algorithm can also be applied to sampling by using threshold detectors for states that can be written as linear combinations of Gaussian states:

$$\varrho = \sum_{i=1}^{\ell} q_i \varrho_i(\bar{\alpha}_i, \Sigma_i), \quad (19)$$

where $\varrho_i(\bar{\alpha}_i, \Sigma_i)$ is a Gaussian state with vector of means $\bar{\alpha}_i$ and covariance matrix Σ_i .

The coefficients q_i can form a probability distribution, in which case ϱ is a mixture of Gaussian states, or they can be (possibly negative) real numbers. This last case occurs, for instance, when describing non-Gaussian states prepared by using heralding schemes [26]. If the q_i form a probability distribution, then one only needs to draw the Gaussian state $\varrho_i(\bar{\alpha}_i, \Sigma_i)$ with probability q_i and directly use the algorithms derived before. If on the other hand some of the q_i are negative, note that the probability of any event is simply the weighted sum of the probabilities of the Gaussian states $\varrho_i(\bar{\alpha}_i, \Sigma_i)$; thus, the machinery of the algorithm carries over with minimal modifications. In particular, the time required to generate a sample simply gets multiplied by ℓ , the number of Gaussian states combined to generate the state ϱ .

4. Non-negative kernel matrices

Besides being generalizable to different states and measurements, the algorithm can have polynomial time complexity whenever output probabilities can be estimated efficiently. An important example is when the covariance matrix $\mathbf{Q}$, or equivalently Σ , is non-negative, i.e., when $\mathbf{Q}_{ij} \geq 0$ for all i, j . Indeed, as we show in Appendix A, if $\mathbf{Q}^{(k)}$ is both non-negative and a proper quantum covariance matrix, then the reduced Kernel matrices $\mathbf{A}_s^{(k)}$ are also non-negative. This implies that the probabilities $p(S^{(k)})$ can be approximated

efficiently since the hafnian of non-negative matrices can be estimated in polynomial time [41,42].

Let $\mathbf{G}$ be a skew-symmetric random matrix whose entries above the main diagonal are drawn from the standard normal distribution $\mathcal{N}(0, 1)$. Defining a matrix $\mathbf{W}$ with entries $W_{ij} = G_{ij}\sqrt{a_{ij}}$, it holds that [41,42]

$$\text{Haf}(\mathbf{A}) = \mathbb{E}[\det(\mathbf{W})], \quad (20)$$

where a_{ij} are the matrix elements of $\mathbf{A}$. To approximate the hafnian of $\mathbf{A}$, M random matrices $\mathbf{W}_1, \mathbf{W}_2, \dots, \mathbf{W}_M$ are drawn and used to compute the estimator

$$\text{Haf}(\mathbf{A}) \approx \frac{1}{M} \sum_{i=1}^M \det(\mathbf{W}_i). \quad (21)$$

This scenario includes the important situation where the kernel matrix $\mathbf{A}$ is the adjacency matrix of an unweighted graph.

Computing the determinant of a $2n \times 2n$ matrix by using the lower-upper (LU) decomposition method takes time $O(n^3)$, and M such determinants have to be calculated to estimate a single hafnian. As shown in Appendix A, for the cases when the kernel matrix is non-negative and block diagonal $\mathbf{A} = \mathbf{B} \oplus \mathbf{B}$, it follows that the covariance matrix $\mathbf{Q}$ is also non-negative and thus all probabilities in the algorithm can be estimated in polynomial time up to subexponential errors [42]. From the analysis of the previous section, this leads to an asymptotic running time scaling as $O(mMN^3)$.

III. BENCHMARKING

In this section, we test the performance of the algorithm. Tests are based on a Python/C++ implementation, publicly available in the SAMPLES module of Xanadu's The Walrus library [28]. C++ code is used to compute hafnians, while all other operations are performed in Python. As stated in Sec. II C, the runtime of the algorithm scales exponentially with the number of photons and linearly with the number of modes. Since the number of photons is the dominant parameter, we fix the number of modes and test the computational resources required to produce samples with different photon numbers. We consider two scenarios. First, we simulate GBS with equal squeezing levels and an interferometer selected randomly from the Haar measure. Second, we benchmark the approximate sampling algorithm for non-negative kernel matrices built from random Erdős-Renyi graphs.

A. Exact algorithm

When all modes in a GBS device are equally squeezed, the kernel matrix takes the form $\mathbf{A} = \mathbf{B} \oplus \mathbf{B}^*$, with $\mathbf{B}$ a symmetric matrix. The GBS distribution is then

$$p(S) = \frac{1}{\sqrt{\det(\mathbf{Q})}} \frac{|\text{Haf}(\mathbf{B}_S)|^2}{s_1! \cdots s_m!}. \quad (22)$$

For each choice of squeezing parameters and linear interferometer, there exists a corresponding matrix $\mathbf{B}$. Similarly, for every $\mathbf{B}$ there is a corresponding setting of interferometer and squeezing parameters [43].

We simulate this setting for a GBS device with $m = 100$ modes and an interferometer unitary chosen at random from the Haar measure. Squeezing levels in each mode are selected

FIG. 1. Running times of the exact simulation algorithm for a GBS system with 100 modes. The interferometer is Haar-random and the squeezing levels are set uniformly in the interval $[0,1]$ for each mode. The algorithm is implemented in a cluster of 56 Intel(R) Xeon(R) CPUs operating at 2.6 GHz. The top points (circles) indicate the upper bound on the runtime, which occurs when N photons are each observed in the first N modes. The bottom points (stars) show a lower bound on the runtime, when N photons are each observed in the last N modes. Since we fix the detection pattern and also the total number of photons detected, the running time of these extreme cases is independent of the squeezing parameters used for the simulation.

uniformly at random in the interval $[0,1]$. Simulations are performed on a cluster of 56 CPUs. The runtimes for different photon numbers are shown in Fig. 1. As explained in Sec. II C, there are two extreme cases. The fastest runtimes occur when N photons are observed in each of the last N modes. The longest runtimes happen when N photons are each observed in the first N modes. In the simulation, these two situations are hard coded, giving rise to upper and lower bounds on the total runtime. The lower bound indicates that hours are needed to produce a single sample with 24 or more photons. Proportionate reductions in runtime can be obtained by using more powerful processors and parallelizing the computation.

B. Approximate algorithm for non-negative matrices

We benchmark the approximate sampling algorithm when the kernels is the adjacency matrix of a random Erdős-Renyi graph with 100 vertices, where an edge is added with probability $1/2$. When encoded in a GBS device, this corresponds to $m = 100$ modes. In estimating the hafnians by using Eq. (21), $M = 1000$ samples were used. The mean photon number was set to 20. Upper and lower bounds on the runtimes are shown in Fig. 2. As expected, the runtimes are considerably shorter for the approximate algorithm, which can, for example, produce samples with 50 photons in a few minutes. However, this comes at a price: the errors in the approximation are significant.

To monitor the error in the hafnian estimation, we compute the sum of the conditional photon number probabilities when sampling the last mode, i.e., we calculate $\eta = \sum_{i=0}^{n_{\max}} p(s_m = i | s_1, \dots, s_{m-1})$. For $n_{\max}$ sufficiently large, if all probabilities are calculated correctly, it must hold that $\eta = 1$. Therefore

FIG. 2. Running times of the approximate simulation algorithm for non-negative kernel matrices. The kernel is the adjacency matrix of a random Erdős-Renyi graph with 100 vertices and edge probability $1/2$. The algorithm is implemented in a cluster of 56 Intel(R) Xeon(R) CPUs operating at 2.6 GHz. The top points (circles) show the upper bound on the runtime, which occurs when N photons are each detected in the first N modes. The bottom points (stars) indicate a lower bound on the runtime, when N photons are each observed in the last N modes.

the error $\varepsilon = |1 - \eta|$ can be used as a simple method to estimate the reliability of the sampling procedure. We observe an average error of $\varepsilon = 63\%$ across all samples, with a maximum error of $\varepsilon = 228\%$. This can be interpreted as evidence that the algorithm is sampling from a distribution that is far from the target GBS distribution. From standard results in a Monte Carlo estimation, the error in the hafnian approximation formula (21) scales as $1/\sqrt{M}$, meaning that it becomes expensive to decrease the error further.

IV. DISCUSSION

We have described an exact classical algorithm for simulating GBS. The algorithm has minimal space complexity and its runtime is directly proportional to the time required to compute output probabilities. These properties make it the best currently known method for simulating GBS devices. Compared with the state of the art in algorithms for simulating boson sampling, our results place GBS on similar footing, since in all cases the complexity of the sampling algorithms is proportional to the complexity of computing probabilities. The only difference is that for GBS, the complexity also scales linearly with the number of modes. In our algorithm, this occurs because each mode must be sampled individually to account for the fact that photon numbers can vary among samples, a feature that is not present in boson sampling. It remains an open question whether this small dependency on mode number is fundamental or can be removed.

We implemented the algorithm and benchmarked its runtime as a function of the number of photons. It is informative to make a comparison with the sampling rates of physical devices. Sampling rates in GBS are determined by detector dead times, which limit the speed at which photons can be detected. Depending on specific technologies, detector dead times lead to sampling rates in the range of 10^5 – 10^7 Hz [44].

On the other hand, even for the approximate algorithm using only $M = 1000$ samples for hafnian estimation, sampling rates are on the order of 10^{-2} Hz for outputs with a few dozen photons. Even if these rates were boosted by using supercomputing power, the sampling rate of a physical device is likely to be orders-of-magnitude larger than what can be obtained with classical simulators.

Finally, beyond its fundamental importance in determining the practical complexity of GBS simulation, the algorithm can be a valuable tool for researchers working with GBS. The algorithm can be used as a method to test and benchmark new applications, to verify physical implementations, and to model the role of experimental imperfections.

ACKNOWLEDGMENTS

We thank L. Banchi, T.R. Bromley, B. Gupta, J. Izaac, N. Killoran, A. Mari, H. Qi, and M. Schuld for fruitful discussions.

APPENDIX A: NON-NEGATIVE MATRICES AND GAUSSIAN BOSON SAMPLING

Assume Σ is a valid quantum complex-normal covariance matrix. This implies that [45]

$$\Sigma + \frac{1}{2}Z \geq 0 \implies \Sigma > 0, \quad (\text{A1})$$

where $Z = \begin{bmatrix} \mathbb{1} & 0 \\ 0 & -\mathbb{1} \end{bmatrix}$. We furthermore assume that Σ is non-negative,

$$\Sigma_{ij} \geq 0 \forall i, j. \quad (\text{A2})$$

This implies that $\mathbf{Q} = \Sigma + \mathbb{1}/2$ is also non-negative. We want to show that the kernel matrix $\mathbf{A} = \mathbf{X}(\mathbb{1} - \mathbf{Q}^{-1})$ is non-negative. Note that $\mathbf{X}$ is a permutation matrix and thus $\mathbf{A}$ is non-negative if and only if

$$\mathbf{Q}^{-1} = \mathbb{1} - \mathbf{O} \quad (\text{A3})$$

is non-negative. Since $\mathbf{Q}$ has non-negative entries and is positive definite, its inverse $\mathbf{Q}^{-1}$ is a so-called M -matrix [46]. If $\mathbf{M}$ is an M matrix, it can be shown that

$$M_{ii} > 0 \text{ and } M_{ij} \leq 0 \text{ for } i \neq j. \quad (\text{A4})$$

Furthermore, if $\mathbf{M}$ is an M matrix, then there exists a scalar λ_0 and a matrix $\mathbf{N}$ such that [46]

$$\mathbf{M} = \lambda_0 \mathbb{1} - \mathbf{N}, \quad (\text{A5})$$

where $\lambda_0 \geq \max\{\text{eigvals}(\mathbf{N})\}$ and $N_{ij} \geq 0$. Therefore, since $\mathbf{Q}^{-1}$ is an M matrix, we can write

$$\mathbf{Q}^{-1} = \lambda_0 \mathbb{1} - \mathbf{N}. \quad (\text{A6})$$

Comparing Eqs. (A3) and (A6), we can identify

$$\lambda_0 = 1, \quad (\text{A7})$$

$$\mathbf{N} = \mathbf{O}. \quad (\text{A8})$$

However, for this identification to be correct, it remains to show that $\lambda_0 = 1 \geq \max\{\text{eigvals}(\mathbf{O})\}$. This is easily seen from

the following chain of inequalities:

$$\begin{aligned} 0 < \Sigma &\implies \mathbb{1}/2 < \mathbf{Q} \implies \\ 0 < \mathbf{Q}^{-1} < 2\mathbb{1} &\implies -\mathbb{1} < \underbrace{\mathbb{1} - \mathbf{Q}^{-1}}_{=\mathbf{O}} < \mathbb{1}. \end{aligned} \quad (\text{A9})$$

Therefore, we conclude that $\mathbf{O}$ has non-negative entries. Note that $\mathbf{Q}$ can also stand for the covariance matrix of a subset of the modes, which in the main text is labeled $\mathbf{Q}^{(k)}$. This concludes the proof that the kernel matrix is non-negative if the covariance matrix is non-negative.

Now we consider the inverse problem. We study certain kernel matrices that always have non-negative covariance matrices. Consider pure Gaussian states for which the covariance matrix has the following simple parametrization [10,11]:

$$\mathbf{Q} = (\mathbb{1} - \mathbf{X}\mathbf{A})^{-1}, \quad (\text{A10})$$

$$\mathbf{A} = \begin{bmatrix} \mathbf{B} & 0 \\ 0 & \mathbf{B}^* \end{bmatrix}, \quad (\text{A11})$$

$$\mathbf{X} = \begin{bmatrix} 0 & \mathbb{1} \\ \mathbb{1} & 0 \end{bmatrix}, \quad (\text{A12})$$

where $\mathbf{B} = \mathbf{B}^T$ has singular values that satisfy

$$0 \leq \text{singvals}(\mathbf{B}) < 1. \quad (\text{A13})$$

As before, let us now specialize to the case where $\mathbf{B}$ is real and non-negative. The last bound can be restated in terms of the eigenvalues of $\mathbf{B}$ as

$$-1 < \text{eigvals}(\mathbf{B}) < 1. \quad (\text{A14})$$

This last fact implies that the spectrum of $\mathbf{A}$ is also in $(-1, 1)$. Using this observation we can rewrite

$$\mathbf{Q} = f(\mathbf{X}\mathbf{A}), \text{ where} \quad (\text{A15})$$

$$f(y) = \frac{1}{1-y} = \sum_{n=0}^{\infty} y^n. \quad (\text{A16})$$

In the last equation we used the bound on the eigenvalues of $\mathbf{X}\mathbf{A}$ to expand $f(y)$ as a power series in y . The bound in Eq. (A14) guarantees that the series converges. Now note that $\mathbf{X}\mathbf{A}$ has non-negative entries, therefore any of its powers will also have non-negative entries. Furthermore, the sum of two matrices with non-negative entries is also non-negative, thus showing that $\mathbf{Q}$ is non-negative.

APPENDIX B: BOUNDING THE EFFECT OF FINITE PHOTON-NUMBER RESOLUTION

Let S denote the random variable describing the outcomes of an ideal GBS setup where the photon-number detectors can resolve any number of photons with perfect accuracy. The support of this random variable is $\mathbb{N}_0^m$, where $\mathbb{N}_0$ is the set of natural numbers including zero and m is the number of modes. Consider the probability distribution $q(S)$ obtained from the ideal distribution $p(S)$ by taking samples S in which more than $n_{\max}$ photons are detected in any of the modes, and mapping them to the event $\emptyset$. For example, for $m = 4$ and $n_{\max} = 8$, the sample $(0, 1, 2, 10)$ is mapped to $\emptyset$, whereas the event $(0, 1, 2, 5)$ is mapped to itself.

FIG. 3. (a) Bound on the total variation distance for a GBS setup as a function of the maximum number of photons k that can be resolved in a detector. We assume 40 squeezed states with input mean photon number $\bar{n} = 1$ are sent into a 100×100 interferometer. The black dotted lines consider the lossless case and the gray lines consider applying 3 dB of loss in each of the modes. Note the exponential decay in the bound on the total variation distance. The results are averaged over 1000 unitaries drawn from the Haar measure and the bars denote one standard deviation. (b) We show the distribution of the total photon number $P(n)$ for pure and lossy GBS with the same parameters used in the top figure.

Any sample in which no more than $n_{\max}$ photons are detected in each mode has the same probability under the two distributions. Thus, they only differ in events in which more than $n_{\max}$ photons are measured. This allows us to write their

total variation distance as follows:

$$\begin{aligned} \delta &:= \frac{1}{2} \sum_{S \in \mathbb{N}_0^m \cup \emptyset} |p(S) - q(S)| = \frac{1}{2} \left(q(\emptyset) + \sum_{S > n_{\max}} p(S) \right) \\ &= \sum_{S > n_{\max}} p(S). \end{aligned} \quad (\text{B1})$$

In the last equation we used the notation $S > n_{\max}$ to indicate any S for which at least one mode had more than $n_{\max}$ photons. The support of the distributions is extended so that $p(\emptyset) = 0$ and $q(S) = 0$ if $S > n_{\max}$. Finally, we used the fact that $q(\emptyset) = \sum_{S > n_{\max}} p(S)$.

It then holds that

$$\delta = \sum_{S > n_{\max}} p(S) < \sum_{k=1}^m \sum_{n=n_{\max}+1}^{\infty} P_k(n) = \sum_{k=1}^m \sum_{n=0}^{n_{\max}} [1 - P_k(n)], \quad (\text{B2})$$

where $P_k(n)$ is the *marginal* probability distribution of the number of photons n in mode k . The quantity in the right-hand side of Eq. (B2) is a bound and not an equality because events like $(0,1,10,10)$ are counted twice; once in the marginal for the third mode and once in the marginal for the fourth mode (again assuming $m = 4$ and $k = 8$).

The single-mode marginal probabilities $P_k(n)$ in Eq. (B2) can be calculated in cubic time in n . This functionality is implemented in the QUANTUM module of The Walrus [28], and therefore the bound on the total variation distance can be computed in polynomial time for any $n_{\max}$. In Fig. 3(a) we plot the values of the bound derived for a GBS setup where 40 squeezed states with input mean photon number $\bar{n} = 1$ are sent into a 100×100 interferometer. For average instances of the interferometer matrix, the bound decays exponentially with the cutoff $n_{\max}$.

Finally, in the case where single-mode Gaussian states are sent into an interferometer one can also calculate the distribution $Q(n)$ of the total photon number by simply taking the convolution of the single-mode photon-number distributions before the interferometer $Q_k(n)$. The distribution of the total photon number is not changed by applying an interferometer, thus $Q(n) = P(n)$ where $P(n)$ is the total photon number distribution after the interferometer. The $k - 1$ convolutions of the single-mode distributions each require $O(n_{\max} \log(n_{\max}))$ steps using fast Fourier transforms. This observation allows us to obtain in polynomial time the total photon number distributions shown in Fig. 3(b) for 40 pure (gray-full lines) or 3 dB-lossy (black-dashed lines) squeezed states sent into a 100×100 interferometer. Note how loss reduces the mean and variance of the total photon number distribution.

- [1] S. Aaronson and A. Arkhipov, in *Proceedings of the Forty-Third Annual ACM Symposium on the Theory of Computing* (ACM, New York, NY, 2011), pp. 333–342.
- [2] M. Tillmann, B. Dakić, R. Heilmann, S. Nolte, A. Szameit, and P. Walther, *Nat. Photonics* **7**, 540 (2013).

- [3] J. B. Spring, B. J. Metcalf, P. C. Humphreys, W. S. Kolthammer, X.-M. Jin, M. Barbieri, A. Datta, N. Thomas-Peter, N. K. Langford, D. Kundys *et al.*, *Science* **339**, 798 (2013).
- [4] M. A. Broome, A. Fedrizzi, S. Rahimi-Keshari, J. Dove, S. Aaronson, T. C. Ralph, and A. G. White, *Science* **339**, 794 (2013).

- [5] M. Bentivegna, N. Spagnolo, C. Vitelli, F. Flamini, N. Viggianiello, L. Latmiral, P. Mataloni, D. J. Brod, E. F. Galvão, A. Crespi *et al.*, *Sci. Adv.* **1**, e1400255 (2015).
- [6] S. Barkhofen, T. J. Bartley, L. Sansoni, R. Kruse, C. S. Hamilton, I. Jex, and C. Silberhorn, *Phys. Rev. Lett.* **118**, 020502 (2017).
- [7] A. P. Lund, A. Laing, S. Rahimi-Keshari, T. Rudolph, J. L. O'Brien, and T. C. Ralph, *Phys. Rev. Lett.* **113**, 100502 (2014).
- [8] A. P. Lund, S. Rahimi-Keshari, and T. C. Ralph, *Phys. Rev. A* **96**, 022301 (2017).
- [9] L. Chakhmakchyan and N. J. Cerf, *Phys. Rev. A* **96**, 032326 (2017).
- [10] C. S. Hamilton, R. Kruse, L. Sansoni, S. Barkhofen, C. Silberhorn, and I. Jex, *Phys. Rev. Lett.* **119**, 170501 (2017).
- [11] R. Kruse, C. S. Hamilton, L. Sansoni, S. Barkhofen, C. Silberhorn, and I. Jex, *Phys. Rev. A* **100**, 032326 (2019).
- [12] D. Rosenberg, A. E. Lita, A. J. Miller, and S. W. Nam, *Phys. Rev. A* **71**, 061803(R) (2005).
- [13] J. Huh, G. G. Guerreschi, B. Peropadre, J. R. McClean, and A. Aspuru-Guzik, *Nat. Photonics* **9**, 615 (2015).
- [14] J. M. Arrazola and T. R. Bromley, *Phys. Rev. Lett.* **121**, 030503 (2018).
- [15] J. M. Arrazola, T. R. Bromley, and P. Rebentrost, *Phys. Rev. A* **98**, 012322 (2018).
- [16] L. Banchi, B. Fingerhuth, T. Babej, and J. M. Arrazola, *arXiv:1902.00462*.
- [17] M. Schuld, K. Brádler, R. Israel, D. Su, and B. Gupt, *Phys. Rev. A* **101**, 032314 (2020).
- [18] S. Jahangiri, J. M. Arrazola, N. Quesada, and N. Killoran, *Phys. Rev. E*, **101**, 022134 (2020).
- [19] W. R. Clements, J. J. Renema, A. Eckstein, A. A. Valido, A. Lita, T. Gerrits, S. W. Nam, W. S. Kolthammer, J. Huh, and I. A. Walmsley, *J. Phys. B: At., Mol. Opt. Phys.* **51**, 245503 (2018).
- [20] S. Paesani, Y. Ding, R. Santagati, L. Chakhmakchyan, C. Vigliar, K. Rottwitt, L. K. Oxenløwe, J. Wang, M. G. Thompson, and A. Laing, *Nat. Phys.* **15**, 925 (2019).
- [21] H.-S. Zhong, L.-C. Peng, Y. Li, Y. Hu, W. Li, J. Qin, D. Wu, W. Zhang, H. Li, L. Zhang *et al.*, *Sci. Bull.* **64**, 511 (2019).
- [22] A. Neville, C. Sparrow, R. Clifford, E. Johnston, P. M. Birchall, A. Montanaro, and A. Laing, *Nat. Phys.* **13**, 1153 (2017).
- [23] P. Clifford and R. Clifford, in *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia, PA, 2018), pp. 146–155.
- [24] J. J. Renema, *arXiv:1911.10112*.
- [25] A. E. Moylett, R. García-Patrón, J. J. Renema, and P. Turner, *Quantum Sci. Technol.* **5**, 015001 (2019).
- [26] N. Quesada, J. M. Arrazola, and N. Killoran, *Phys. Rev. A* **98**, 062322 (2018).
- [27] B. Gupt, J. M. Arrazola, N. Quesada, and T. R. Bromley, *arXiv:1810.00900*.
- [28] B. Gupt, J. Izaac, and N. Quesada, *J. Open Source Software* **4**, 1705 (2019).
- [29] C. Weedbrook, S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd, *Rev. Mod. Phys.* **84**, 621 (2012).
- [30] A. Serafini, *Quantum Continuous Variables: A Primer of Theoretical Methods* (CRC Press, Boca Raton, FL, 2017).
- [31] B. Picinbono, *IEEE Trans. Signal Process.* **44**, 2637 (1996).
- [32] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walmsley, *Optica* **3**, 1460 (2016).
- [33] H. de Guise, O. Di Matteo, and L. L. Sánchez-Soto, *Phys. Rev. A* **97**, 022328 (2018).
- [34] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, *Phys. Rev. Lett.* **73**, 58 (1994).
- [35] E. R. Caianiello, *Nuovo Cimento* **10**, 1634 (1953).
- [36] A. Barvinok, *Combinatorics and Complexity of Partition Functions* (Springer, Cham, Switzerland, 2016), Vol. 276.
- [37] A. Björklund, B. Gupt, and N. Quesada, *J. Exp. Algor.* **24**, 1 (2019).
- [38] A. Björklund, in *Proceedings of the Twenty-Third ACM-SIAM Symposium on Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia, PA, 2012), pp. 914–921.
- [39] N. Quesada, *J. Chem. Phys.* **150**, 164113 (2019).
- [40] N. Quesada, L. G. Helt, J. Izaac, J. M. Arrazola, R. Shahrokhshahi, C. R. Myers, and K. K. Sabapathy, *Phys. Rev. A* **100**, 022341 (2019).
- [41] A. Barvinok, *Random Struct. Algor.* **14**, 29 (1999).
- [42] M. Rudelson, A. Samorodnitsky, O. Zeitouni *et al.*, *Ann. Probab.* **44**, 2858 (2016).
- [43] T. R. Bromley, J. M. Arrazola, S. Jahangiri, J. Izaac, N. Quesada, A. D. Gran, M. Schuld, J. Swinarton, Z. Zabaneh, and N. Killoran, *arXiv:1912.07634*.
- [44] R. H. Hadfield, *Nat. Photonics* **3**, 696 (2009).
- [45] R. Simon, N. Mukunda, and B. Dutta, *Phys. Rev. A* **49**, 1567 (1994).
- [46] G. Poole and T. Boullion, *SIAM Rev.* **16**, 419 (1974).