Simulability of partially distinguishable superposition and Gaussian boson sampling

Jelmer J. Renema

Mesa+ Institute for Nanotechnology, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

We study the hardness of classically simulating boson sampling with superposition and Gaussian input states at nonzero photon indistinguishability. We find that, similar to regular boson sampling, distinguishability causes exponential attenuation of the many-photon interference terms in both these boson sampling variants. For superposition sampling, we find that it is not simulable with out method at zero indistinguishability, which is evidence for the computational hardness of this problem, and we find that it is simulable at any level of particle distinguishability, similar to regular boson sampling. If an efficient classical algorithm to approximate a given sum over permanents is found, this approach also leads to an efficient classical algorithm to simulate Gaussian boson sampling in the presence of distinguishability.

The next milestone in experimental photonic quantum information processing is the demonstration of a quantum advantage: a well-defined computational task at which a quantum device outperforms a classical computer [1–4]. A photonic implementation of this concept is boson sampling [5], which consists of sending single photons through a linear optical network followed by photodetection, a task which is strongly believed to be hard for a classical computer to simulate. The best known algorithm to classically simulate a boson sampler is that of Clifford and Clifford [6], which generates a sample in $m2^m$ steps, where m is the number of photons. Given this scaling and the speed of modern supercomputers, it is believed that a boson sampler will outperform a classical simulation on a supercomputer around $m \approx 50$ photons [7,8]. Experimentally demonstrating boson sampling is the subject of experimental efforts worldwide [9–17].

A major problem in boson sampling theory is understanding the degree to which quantum photonic interference is susceptible to imperfections [18–26]. These imperfections can take the form of any experimental noise which degrades the quantum nature of the observed interference pattern. The two imperfections most strongly present in experiments are photon loss and distinguishability. Distinguishability is the imperfect overlap of the wave functions of the photons in the other degrees of freedom besides their position (e.g., polarization, frequency, time). Loss is when not all photons produced by the sources are ultimately detected. Other imperfections include dark counts (detection events not associated with the arrival of a photon) and fluctuations in the settings of the interferometer.

In a series of recent works [27–32], it was shown that boson samplers that have any of these imperfections at a sufficiently strong level can be efficiently classically simulated, i.e., in polynomial time, thereby negating any potential quantum advantage. The strategy behind these simulations is to split the interference pattern into a series of j-photon interference terms, where j runs from 0 to the number of photons m. Imperfections degrade the higher order interference terms more strongly than the lower orders, which means that for sufficiently strong imperfections, the series over j can be truncated at some value k. The resulting algorithm is polynomial in m.

The strength of the imperfections determines the maximum level of interference k, thereby functioning as a criterion for simulability of k-photon interference. The physical picture arising from these results is that imperfect boson sampling can be thought of as interference between all possible groups of k photons and classical (i.e., single-particle) transmission of the remaining $m - k$ photons.

However, these simulability criteria do not automatically extend to boson sampling variants. Variants on boson sampling exist as workarounds for the fact that high-efficiency, deterministic single-photon sources are not available experimentally. The two main kinds of photon sources currently available are those based on single-photon emission from quantum dots, where a π Rabi flip is used to excite a two-level system [33–35], which then spontaneously emits a photon, and those based on spontaneous parametric down-conversion (PDC) in nonlinear optical media such as β-barium borate and potassium titanyl phosphate [36].

While quantum dot sources have seen rapid progress in recent years [37], PDC sources still hold the record for photon collection efficiency (which is the relevant quantity for complexity analysis) and indistinguishability. However, their main weakness is that they do not emit single photons but two-mode squeezed states, which are nonclassical states consisting of a thermal distribution of pairs of photons, including zero pairs. The presence of a zero-photon component in the state means that these sources cannot be made to function deterministically. To circumvent the fact that these sources are nondeterministic, the usual strategy is heralding: By detecting one photon, the existence of the other can be inferred, and it can then be used for experiments [38]. Such heralded nondeterministic sources can be used to sample over a classical probability distribution of input states by recording which sources produced a photon in any given run of the experiment. This problem, known as scattershot boson sampling, directly derives its hardness from the original boson sampling problem [39,40].

A further step in this line of thought is to dispense with the heralds altogether, and feed both photons into the linear network [41]. This protocol is known as Gaussian boson

sampling (GBS), after the fact that a squeezed state is a Gaussian state. For this situation, the reduction to a known hardness argument is via the case of very weak squeezing (so that multiphoton contributions can be neglected), and where the linear optical transformation effected by the interferometer is separable into one acting on the signal modes and one acting on the herald modes [42], or equivalently where the action on the herald modes is the identity. In that case, the problem reduces to scattershot boson sampling [39,40].

It is largely unknown how imperfections affect Gaussian boson sampling in a complexity-theoretical sense [43]. However, it is evident that Gaussian boson sampling suffers from many of the same types of experimental imperfections as regular boson sampling, including photon loss in the sources, interferometers, and detectors, as well as partial distinguishability induced by spectral correlations in the parametric downconversion sources which are typically used.

In this work, we will investigate the simulability of Gaussian boson sampling (GBS) and an auxilliary model, which we call superposition boson sampling (SBS), which is inspired by recent work on emission of coherent superposition states from quantum dots [44]. We will consider these two systems both with and without noise. We will show that without noise, both systems are not simulable by the method outlined above. For superposition sampling, this is evidence (although circumstantial) of computational hardness for this system. When we introduce noise, we see that the same physical picture arises for both systems as for regular boson sampling: The interference pattern falls apart into smaller interference processes of size $k < m$, where k depends only on the strength of the imperfections. For SBS, we are able to give an explicit algorithm to efficiently classically simulate imperfect boson sampling. For GBS, we reduce the problem of finding such an algorithm to an open problem in the theory of matrix permanents.

We restrict ourselves to studying the effect of one particular imperfection, namely photon distinguishability. We note, however, that for regular (i.e., Fock-state) boson sampling, all imperfections considered in the literature so far can be analyzed using the strategy which we will use. This restriction is therefore not as strong as it might at first appear. We hope to study the effect of other imperfections on Gaussian boson sampling in later work.

The paper is organized into six sections as follows: In Sec. I, we will begin by recalling a method to compute the output of a boson sampler fed with arbitrary input states. Then in Sec. II, we will rederive the simulability criteria for regular boson sampling. Our results begin in Sec. III, where we will show how to combine these two into a criterion for simulability of boson sampling with arbitrary input states. Next, we will apply these results to the special case of superposition boson sampling in Sec. IV. Finally, we will show in Sec. V that under weak pumping conditions, Gaussian boson sampling reduces to superposition sampling. The paper ends with concluding remarks in Sec. VI.

I. BOSON SAMPLING WITH ARBITRARY INPUT STATES

The theory for boson sampling with arbitrary input states was derived in Ref. [45]. Here, we provide a derivation of the

expressions for a detection probability at the output of a linear optical network with arbitrary states at the input and finite distinguishability for reasons of exposition and to establish the framework used in the present work.

Consider an arbitrary multimode photonic quantum state $|\psi\rangle$ impinging on a linear optical network U with detectors in the Fock basis at the output of the network. We wish to compute the probability of an arbitrary pattern of detection events $|\phi_{pd}\rangle$ and to study the hardness of computing that outcome. The probability is given by

$$P(\phi_{pd}) = |\langle\psi|U|\phi_{pd}\rangle|^2. \quad (1)$$

We can then insert a resolution of the identity in the Fock basis: $\hat{I}_{\text{fs}} = \sum_p |\xi_p\rangle\langle\xi_p|$, where $|\xi\rangle = \prod_{i=1}^n |m_i\rangle$ is a product of Fock states (and the index p runs over all such possible Fock state products), and m_i are the mode occupation numbers:

$$P(\phi_{pd}) = \left|\sum_p \langle\psi|\xi_p\rangle\langle\xi_p|U|\phi_{pd}\rangle\right|^2. \quad (2)$$

Since the interferometer U is photon number preserving, all terms in the sum over ξ which do not contain the same number of photon numbers as $|\phi_{pd}\rangle$ drop out, i.e., $\langle\xi_p|U|\phi_{pd}\rangle = 0$ unless $|\langle\xi_p|\hat{N}|\xi_p\rangle|^2 = |\langle\phi_{pd}|\hat{N}|\phi_{pd}\rangle|^2$, where $\hat{N}$ is the multimode photon number operator. Hence we relabel the sum over p to contain only those terms which meet this condition. Expanding, we have

$$P(\phi_{pd}) = \sum_p \sum_q \langle\psi|\xi_p\rangle\langle\psi|\xi_q\rangle^\dagger\langle\xi_p|U|\phi_{pd}\rangle\langle\xi_q|U|\phi_{pd}\rangle^\dagger$$

$$= \sum_p \sum_q c_p c_q^\dagger \text{Perm}(M_p)\text{Perm}(M_q)^\dagger. \quad (3)$$

Here we have used the fact that since ξ is a product of Fock states, we can apply the identity $\langle\xi_p|U|\phi_{pd}\rangle = \text{Perm}(M_{\xi_p\phi})/\sqrt{\mu(\xi_p)\mu(\phi_{pd})}$, where Perm is the permanent function $\text{Perm}(M) = \sum_\sigma \prod_i M_{i,\sigma_i}$, and M is the submatrix connecting ξ_p and ϕ_{pd} [46]. Since we are concerned only with a single output ϕ_{pd}, we will suppress this dependence, as well as the subscript ξ and denote the matrix as M_p. The coefficients are given by $c_p = \langle\psi|\xi_p\rangle/\sqrt{\mu(\xi_p)\mu(\phi_{pd})}$. $\mu(\xi)$ is the multiplicity of a particular Fock state configuration: $\mu(\xi) = \prod_i(m_i!)$. Furthermore, in what follows, we will assume $\mu(\phi_{pd}) = 1$, a condition which can be enforced with high probability by making the linear interferometer large enough [5].

Note that Eq. (3) is a natural way to consider interference from a quantum state with an indeterminate photon number in each mode: This equation simply tallies all the ways in which the set of sources could have produced a given number of photons, interfering with all the other ways those sources could produce that number of photons. These interference terms have been observed in experiments [47].

Anticipating our complexity analysis, we reorder terms in Eq. (3):

$$P(\phi_{pd}) = \sum_p \sum_q c_p c_q^\dagger \sum_\sigma \text{Perm}(M_p \circ M_{\sigma(q)}^\dagger), \quad (4)$$

where $\sigma(q)$ denotes a permutation of the matrix elements of M_q, and $\circ$ denotes the elementwise product.

Next, we look for cases where the combination of ξ_p, ξ_q, and σ causes a product between the ith row of M and its complex conjugate to appear in the permanent. Such a row of modulus-squared terms implies classical interference of the corresponding photon. We will later see that these terms also correspond to the fixed points j of the partial permutation $\xi_p \rightarrow \sigma(\xi_q)$. We can split off these positive rows using Laplace expansion along those rows:

$$
\begin{aligned}
P(\phi_{pd}) &= \sum_p \sum_q c_p c_q^\dagger \sum_j^m \sum_{\sigma^j} \mathrm{Perm}(M_p \circ M_{\sigma(q)}) \\
&= \sum_p \sum_q c_p c_q^\dagger \sum_j^m \sum_{\sigma^j} \sum_\rho \mathrm{Perm}\big(M_{p,\rho} \circ M_{\sigma_p(q),\rho}^\dagger\big) \\
&\quad \times \mathrm{Perm}(|M_{p,\bar\rho}|^2),
\end{aligned} \tag{5}
$$

where σ_p stands for the part of σ not corresponding to fixed points, ρ is a j partition of the number of detected photons m, the sum runs over all j partitions, and the overbar denotes the complement of ρ on the set ξ_p.

We now determine the effect of distinguishability on this optical system. From Eq. (4), it is clear that for fixed p and q, we are considering all the ways in which photons from configurations ξ_p and ξ_q can give rise to a given detection event. Therefore, when considering distinguishability, we must consider the overlap of the internal degrees of freedom at the output side of a double-sided Feynman diagram [25]. The prefactor to adjust for distinguishability is given by $\prod_i \langle \bar\psi_{p_i} | \bar\psi_{\sigma_i(q)} \rangle$, where $|\bar\psi_{q_i}\rangle$ denotes the internal state (degrees of freedom unaffected by U) for the ith photon in q, and $\sigma_i(q)$ denotes the ith element of the permutation σ acting on q. Hence, we obtain

$$
\begin{aligned}
P(\phi_{pd}) &= \sum_p \sum_q c_p c_q^\dagger \sum_j \sum_{\sigma^j} \left(\prod_i \langle \bar\psi_{p_i} | \bar\psi_{\sigma_i(q)} \rangle \right) \\
&\quad \times \sum_\rho \mathrm{Perm}\big(M_{p,\rho} \circ M_{\sigma_p(q),\rho}\big) \mathrm{Perm}(|M_{p,\bar\rho}|^2).
\end{aligned} \tag{6}
$$

To simplify the analysis without losing any of the essential features, we will assume throughout this paper that all distinguishabile inner products between the pth and qth photons are equal to a constant value x for $p \neq q$. In that case, the distinguishability factor $\prod_i \langle \bar\psi_{p_i} | \bar\psi_{\sigma_i(q)} \rangle$ reduces to x^j. The case of unequal indistinguishability is treated elsewhere [28].

II. NOISE THEORY FOR A FOCK STATE INPUT

To introduce our techniques for complexity analysis, we review the derivation of the simulability criterion for the case where the input state $|\Psi\rangle$ is a product of Fock states. This section is a summary of Refs. [28,29], and the supplemental material included therein.

For the case of a product-of-Fock-states input, the double sum over ξ in Eq. (6) drops out since there is only a single $\xi = |\Psi\rangle$, and we are left with

$$
P(\phi_{pd}) = \sum_{j=0}^m x^j \sum_{\sigma^j} \sum_\rho \mathrm{Perm}(M_{1,\rho} \circ M_{\sigma_p,\rho}^\dagger) \mathrm{Perm}(|M_{1,\bar\rho}|^2). \tag{7}
$$

We wish to construct an algorithm to approximate P by some efficiently computable quasiprobability P'. The strategy which we will follow is to truncate the outer sum of Eq. (7) at some value $k < m$. The intuition for this is that for all $x < 1$, the terms in Eq. (7) are exponentially suppressed by the effect of partial distinguishability. We will show that, averaged over M, the sum over permanents in Eq. (7) is of equal magnitude for all j. This means that $P(\phi_{pd})$ can be interpreted as a polynomial in x with coefficients of order 1, where the higher power terms can be neglected when the polynomial is evaluated at sufficiently small values of x.

The natural way to evaluate the quality of such an approximation is to compute the variational distance $d = \frac{1}{2} \sum_\phi |P(\phi) - P'(\phi)|$, where the sum runs over all possible output configurations ϕ of a boson sampler. The first thing to note is that d is a function of the matrix U associated with a given boson sampler, which is highly inconvenient. Therefore, we look instead at the average $E_U(d)$, where the average is taken over the Haar measure of unitary matrices. Such an average can be related to the value of d of any arbitrary matrix U by a Markov inequality: The probability that a given matrix U has a value $d(U) > c E_U(d)$ (where c is a positive constant) is at most $1/c$. If we are willing to accept some small probability δ of failure in our algorithm, it therefore suffices to compute the average of d over all unitaries.

In order to compute $E_U(d)$, we make the assumption that the dimension N of our matrix U is much larger than the number of photons m. In this limit, the correlations between elements of U due to the unitary constraint can be neglected; the elements approach independent complex Gaussians with mean $\mu = 0$ and standard deviation $\sigma = 1/\sqrt{2N}$ in both real and imaginary parts. This limit is the situation which which the hardness of boson sampling is believed to hold. Furthermore, in this situation, all output configurations are equivalent, the average probability of each individual outcome of the sampler is given by $m!/N^m$, and the probability of collision events (with two or more bosons emerging from the same output port) can be neglected. For these reasons, it suffices to compute the error $E(\Delta P) = E(|P(\phi) - P'(\phi)|)$ of a single outcome, and show that it is of the form $E(\Delta P) = Cm!/N^m$, with C some constant [29]. In that case, the average variational distance is given by $E_U(d) = C$.

To compute the error ΔP on a approximating a single output probability, we define $c_j = \sum_{\sigma^j} R_\sigma$, with $R_\sigma = \mathrm{Re}[\sum_\rho \mathrm{Perm}(M_{1,\rho} \circ M_{\sigma_p,\rho}^\dagger) \mathrm{Perm}(|M_{1,\bar\rho}|^2)]$, where Re denotes the real part. Hence, we can rewrite Eq. (7) as

$$
P(\phi_{pd}) = \sum_j c_j x^j = \sum_j x^j \sum_{\sigma^j} R_\sigma. \tag{8}
$$

The first thing to note is that in our definition of R_σ, it suffices to consider only the real parts of the terms in Eq. (7). This is due to the fact that $R_\sigma = R_{\sigma^{-1}}^\dagger$, meaning that imaginary contributions to Eq. (8) cancel pairwise. Therefore, we may concern ourselves only with the real parts of the terms.

Next, we compute the variance of c_j:

$$
\mathrm{var}(c_j) = \sum_{\sigma^j} \mathrm{var}(R_{\sigma^j}) + \sum_{\sigma^j, \tau^j} \mathrm{cov}(R_{\sigma^j}, R_{\tau^j}). \tag{9}
$$

Since $E(R_\sigma) = 0$, $\mathrm{cov}(R_\sigma, R_\tau) = E(R_\sigma R_\tau)$. This is only nonzero if the parts of σ and τ which carry phase are each other's inverse. The part of σ and τ which carries phase is the entire permutation apart from the fixed points. We will denote the nonfixed part of a (partial) permutation σ as σ_p, so that the requirement can be stated as $\sigma_p = \tau_p^{-1}$. The reason for this is that when averaging over the Haar measure, each matrix element has an independent, uniformly distributed phase, which means that that matrix element averages out to zero. This is also true for products of matrix elements, unless each phase is canceled pairwise by its conjugate phase, as happens when $\sigma_p = \tau_p^{-1}$ [28]. Note that while the fixed points may in principle be different, in the specific scenario which we are considering now, where there is only one ξ, (i.e., product of Fock states input), the condition $\sigma_p = \tau_p^{-1}$ enforces $\sigma = \tau^{-1}$. Importantly, this will not be true in the general case which we will treat later on.

Continuing our computation of $\mathrm{var}(c_j)$, it turns out that if $\sigma_p = \tau_p^{-1}$, then $1/e < \mathrm{cov}(R_\sigma, R_\tau)/\sqrt{\mathrm{var}(R_\sigma)\mathrm{var}(R_\tau)} < 1$. Therefore, the problem of estimating $\mathrm{var}(c_j)$ consists essentially of estimating $\mathrm{var}(R_\sigma)$, and of counting the number of matches between all possible permutations σ and τ [28]. $\mathrm{var}(R_\sigma)$ is a function of j, which by definition is the size of σ_p. For $\mu(\xi) = 1$, it is given by

$$\mathrm{var}(R_\sigma) = e\binom{m}{j}(m-j)!^2 j!/2N^{2m}. \tag{10}$$

Since for the case of a single ξ, $\sigma_p = \tau_p^{-1}$ implies $\sigma = \tau^{-1}$, in this case we have

$$\mathrm{var}(c_j) < 2\mathrm{var}(R_\sigma)R(m, m-j). \tag{11}$$

where $R(m, m-j)$ is the rencontres number, which counts the number of permutations of size m with $m-j$ fixed points. In the limit of large m, $R(m, m-j) = \binom{m}{j}j!/e$; hence, $\mathrm{var}(c_j) < m!^2/N^{2m}$. Applying the inequality $E(|x|) < \sqrt{\mathrm{var}(x)}$, which holds for any variable x with $E(x) = 0$, we upper bound the deviations of c_j from zero as $E(|c_j|) < m!/N^m$. Substituting this back into the truncated version of Eq. (7) and summing the error terms, we find that for a single outcome, the expected error $E_U(\Delta P) = E_U(|P - P'|)$ is given by $\Delta P = m!/N^m\sqrt{x^{2(k+1)}/(1-x^2)}$, which is of the required form. This implies that for a disinguishability level x, a boson sampler is classically simulable on average with error $E_U(d)$ at an approximation level k given by $E_U(d) = \sqrt{x^{2(k+1)}/(1-x^2)}$. Crucially, in the limit of large m, the level of truncation k does not depend on the number of photons m (apart from the trivial point that k must be smaller than m), but only on the level of distinguishability x. The case of finite k and m is dealt with in Ref. [28], but does not alter the picture substantially.

The next step is to observe that by truncating the sum over j in Eq. (7) at a fixed k, we have produced an approximation which can be computed efficiently. Approximating permanents of matrices of positive numbers can be done efficiently [48], whereas the best known algorithm for computing permanents of arbitrary matrices scales as $n2^n$ with the matrix size n [49,50]. By truncating Eq. (7) at some level k, we have therefore achieved polynomial scaling of our approximate distribtion P' with the number of photons m [28].

Having found a distribution of which individual elements can be computed efficiently, and which is close in variational distance to the output distribution of an imperfect boson sampler, the next task is to convert this into a sampling algorithm. We do this by running a Markov chain Monte Carlo sampler on our approximate distribution [7].

III. NOISE THEORY FOR ARBITRARY INPUT STATE

In Sec. I, we recalled the theory for computing the probability of an outcome of a boson sampler fed with arbitrary quantum states. We observed that for a fixed detection outcome, this results in a double sum over all possible ways in which that quantum state could have given rise to the observed number of photons.

In Sec. II, we observed that the way to compute the effect of noise on a boson sampler is to count all pairs of permutations σ and τ of the photons which match a given pattern, namely $\sigma_p = \tau_p^{-1}$.

It will not come as a surprise, then, that the way to compute the effect of noise on a boson sampler fed with arbitrary input states is to compute the number of ways in which a double sum over input states can give rise to permutations satisfying this same rule.

More formally, if we consider applying the scheme described in Sec. II to Eq. (6), we have

$$\mathrm{var}(c_j) = \sum_{p,q,\sigma} \mathrm{c_p c_q^\dagger var}(R(\xi_p, \sigma(\xi_q)))$$
$$+ \cdots \sum_{p,q,r,s,\sigma,\tau} \mathrm{c_p c_q^\dagger c_r c_s^\dagger cov}(R(\xi_p, \sigma(\xi_q))R(\chi_r, \tau(\chi_s))), \tag{12}$$

where $R(\xi_p, \sigma(\xi_q)) = \mathrm{Re}(\mathrm{Perm}(M_p \circ M_{\sigma(q)}^\dagger))$, and where ξ_p, ξ_q, χ_r, and χ_s are all elements of the set of ξ, and σ and τ are permutations. The sextuple sum is over all valid assignments of these variables where at least one variable from $\{\xi_p, \xi_q, \sigma\}$ differs from its counterpart. Following the phase cancellation argument presented above, the condition for nonzero covariance can now be given as an elementwise set of conditions on the six variables, which we call the *pairing rules*:

$$\forall i \in \{1...m\} : ((\xi_{p,i} = \sigma(\xi_q)_i) \wedge (\chi_{r,\rho_j} = \tau(\chi_s)_{\rho_i}))$$
$$\vee ((\xi_{p,i} = \tau(\chi_s)_{\rho_i}) \wedge (\chi_{r,\rho_i} = \sigma(\xi_q)_i)), \tag{13}$$

where ρ is some freely chosen permutation of the indices of χ_r and $\sigma(\chi_s)$, and $\xi_{p,i}$ denotes the ith element of the permutation ξ_p (and similarly for the other variables). Equation (13) is just an elementwise restatement of the requirement that the unfixed points of the permutation must match up, while the fixed points are free. In other words, for each pair of elements of ξ_p and $\sigma(\xi_q)$ [i.e., for each pair of $\xi_{p,i}$ and $\sigma(\xi_q)_i$] we must either have that $\xi_{p,i} = \sigma(\xi_q)_i$, (fixed point) or if this is not the case, we must have that we can uniquely match up this pair to a counterpart, i.e., $\xi_{p,i} = \tau(\chi_r)_{\rho_i}$ and $\chi_{r,\rho_i} = \sigma(\xi_q)$. Note that the requirement that this pairing is unique (no two elements of ξ can be matched to the same χ), combined with the fact that the pairing rules are symmetric in ξ and χ automatically enforces that σ_p and τ_p are the same size.

Furthermore, it should be noted that unlike the case of a single product-of-Fock-states input, we are now dealing with partial permutations instead of complete permutations. A partial permutation is a bijection between two subsets of a given set, whereas a permutation is a bijection on the entire set. The reason we obtain partial permutations is because ξ_p and ξ_q (and equivalently χ_r and χ_s) do not necessarily contain the same elements. This corresponds to the fact that we are considering interference between histories where different sources produced photons.

To illustrate the pairing rules, and to illustrate how the introduction of arbitrary input states changes things compared to Fock state inputs, an example is in order. For our example, will drill down to a single term in the sextuple sum of Eq. (12). In this example, we consider five sources labeled 1 through 5 collectively emitting three photons. The set of allowed ξ is given by $\{\xi\} = \{(1, 2, 3),\ (1, 2, 4),\ (1, 2, 5),\ (1, 3, 5),\ (1, 4, 5),\ (2, 3, 4),\ ...\}$ where the numbers indicate the presence of a photon in a given mode. We now consider an example of two choices of R have covariance with each other. For illustration, we pick a term where $\xi_p \neq \xi_q$, i.e., a cross term in Eq. (5). An example of a cross term is $\xi_p = (1, 2, 4)$, $\xi_q = (2, 3, 4)$, $\sigma = (231)$, meaning $\sigma(\xi_p) = (4, 2, 3)$ and giving rise to the partial permutation (given in two-line notation for clarity) $(\xi_p, \sigma(\xi_q)) = \left(\begin{smallmatrix} 1 & 2 & 4 \\ 4 & 2 & 3 \end{smallmatrix}\right)$. In this case, the partial permutation neglecting fixed points is $\sigma_p = \left(\begin{smallmatrix} 1 & 4 \\ 4 & 3 \end{smallmatrix}\right)$.

When considering which other partial permutations this permutation will have covariance with, we must find one with matching τ_p. For example, taking $\chi_r = (3, 4, 5)$, $\chi_s = (1, 4, 5)$, and $\tau = (213)$ gives $\left(\begin{smallmatrix} 3 & 4 & 5 \\ 4 & 1 & 5 \end{smallmatrix}\right)$ and hence $\tau_p = \left(\begin{smallmatrix} 3 & 4 \\ 4 & 1 \end{smallmatrix}\right)$, which satisfies the covariance rule.

This example illustrates two important points regarding the process of computing the covariance. First, in the case of a coherent superposition of input states, there is more freedom in chosing pairs of R which have covariance, since the two permutations may have different fixed points and still contribute covariance. Second, the pairing rules in Eq. (13) do not split into pairwise requirements on either $\xi_{p,q}$ and $\chi_{r,s}$ or on σ and τ. Indeed, in the example we have given above, all of ξ_p, ξ_q, χ_r, and χ_s are distinct, and $\sigma \neq \tau^{-1}$, yet we can construct a situation where $\sigma_p = \tau_p^{-1}$ holds. The reason the covariance criterion [Eq. (13)] does not split into a criterion on $\xi_{p,q}$ and $\chi_{r,s}$ or on σ and τ is that for each entry, the criterion can either be satisfied by a relation between ξ_p and $\sigma(\xi_q)$ or by a relation between ξ_p and $\tau(\chi_s)$ (and the corresponding variables).

In contrast, since the size of σ_p and τ_p are equal, the pairing rules do split into a criterion on the permuted and unpermuted parts of the permutation. We will use this fact later on in computing the number of possible covariance terms for specific input states.

A few facts from the Fock state case do carry over. In particular, the fact that σ_p and τ_p have the same size means that when we group terms by number of fixed points, we automatically group the R by their potential covariance partners. This implies that c_j have zero covariance with each other, and

hence that the strategy of computing $\text{var}(c_j)$ by considering each j separately is still sound, as it was for Fock states.

Finally, we note that unlike the case of a Fock state input, for arbitrary input states, it does not suffice to show that the coefficients c_j are bounded in variance (and hence that the sum over j can be truncated). The reason for this is that for arbitrary inputs, the truncation does not immediately produce an efficient approximation algorithm. This is because the sum over ξ_p and ξ_q, which arises in the arbitrary state case, can contain exponentially many terms. It is not possible to sample over ξ_p and ξ_q, since the sum is not convex (not all terms are positive). This means that in order to approximate an outcome, we must evaluate not just the sum over j, but also the sum over p and q; we cannot sample over those terms. However, we will see that it is sometimes possible to gather terms in such a way that the overall expression can be efficiently sampled from.

In summary, when encountering a new input state for a boson sampling experiment, our tasks are twofold:

(1) Enumerate the set of ξ_p and corresponding c_p. Compute the covariance between the set of R and hence the truncation point k by applying the pairing rules and Eqs. (10) and (12).

(2) Rewrite Eq. (6) to a form that can be computed efficiently at that truncation, given the particular quantum state in question.

IV. SUPERPOSITION BOSON SAMPLING

Having set up the necessary machinery, we are now ready to investigate the simulability properties of some optical systems of interest. We begin with a model which we name *superposition sampling*. In this model, there are n photon sources, each of which emits the quantum state $|\psi\rangle = \cos(\alpha)|0\rangle + \sin(\alpha)|1\rangle$. These sources form a product state at the input of the interferometer: $|\Psi_{\text{sbs}}\rangle = |\psi\rangle^{\otimes n}|0\rangle^{\otimes N-n}$. Since each source can emit at most one photon, the set of ξ is given by the $\binom{n}{m}$ ways of selecting m photons from n sources. Since all configurations are equally likely and all multiplicities are equal to 1, we can normalize the probabilities assuming that exactly m photons are detected, in which case they are given by $c_p = \sqrt{\binom{n}{m}^{-1}}$ for all p.

This problem is motivated by two facts. First, it is of interest in its own right, due to the fact that it can be implemented using quantum dot sources [44]. Second, it will provide a helpful stepping stone to the analysis of Gaussian boson sampling further on.

To analyze the simulability of this model, we compute $\text{var}(c_j)$ by counting the number of ways in which the pairing rules are satisfied for a given j, as a function of m, n, and j. Since, as we observed in Sec. III, the paring rules split into a requirement corresponding to σ_p and one on the fixed points, we can count the number of ways to satisfy the pairing rules by first assigning the size of σ_p, then the number of ways to create σ_p, and then the number of fixed points.

As an illustration of this combinatorics problem and to show how the paring rules work in practice, we will fill in a table of all variables to which we must assign a value as we go along with our computation of the number of possible assignments. An example of a valid assignment of

all variables is given by

$$
\begin{pmatrix} \xi_p \\ \sigma(\xi_q) \\ \chi_r \\ \tau(\chi_s) \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 1 & 4 & 2 & 3 & 6 & 5 \\ 2 & 3 & 4 & 5 & 6 & 7 \\ 3 & 4 & 2 & 6 & 5 & 7 \end{pmatrix}, \quad (14)
$$

which meets the condition that $\sigma_p = \tau_p^{-1}$.

We begin by noting that since the set of ξ are symmetric under permutation of indices, their choice is entirely equivalent. In our table, we can choose $\xi_p = \xi_1 = \{1...m\}$ out of all $\binom{n}{m}$ possibilities without loss of generality:

$$
\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & ... & m \\ \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \\ \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \\ \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \end{pmatrix},
$$

where the empty set symbol $\emptyset$ represents an as-yet-unassigned value.

Next, we choose the $m - j$ fixed points of σ, i.e., the number of times we will satisfy the first clause of Eq. (13). For $m - j$ fixed points, this can be done in $\binom{m}{j}$ ways (choosing the positions that will not be fixed points), giving a total of $\binom{n}{m}\binom{m}{j}$ ways of reaching this step, and an example table of

$$
\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & ... & n \\ 1 & \emptyset & 3 & \emptyset & \emptyset & \emptyset & \emptyset \\ \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \\ \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \end{pmatrix},
$$

where we have picked some arbitrary assignment of fixed points for illustration. Next, we must assign the unpermuted part of σ. We have already assigned $m - j$ elements, hence we have $\binom{n-m+j}{j}$ ways of choosing the remaining elements, and we have $j!$ ways of arranging these elements. Note that we overcount here: We have neglected to exclude those permutations which introduce additional fixed points. However, in this we overcount by at most a constant factor: The probability that a full permutation includes a fixed point is $1/e$, and it approaches zero for highly partial permutations, i.e., when $m \ll n$. Therefore, we have $\binom{n}{m}\binom{m}{j}\binom{n-m+j}{j}j!$ ways of reaching the following table (again with arbitrary numbers for illustration):

$$
\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & ... & n \\ 1 & 4 & 3 & 7 & 6 & ... & 2 \\ \emptyset & 4 & \emptyset & 7 & 6 & ... & 2 \\ \emptyset & 2 & \emptyset & 4 & 5 & ... & n \end{pmatrix}.
$$

Observe that by fixing σ_p, we have fixed the corresponding elements of τ_p as well. What remains to be done is to pick the remaining fixed points in τ. This can be done independently of the fixed points in σ since the fixed points need not match between σ and τ (as noted above), but not independently from τ_p, since we cannot reuse elements which we have used there. Note that any choice of fixed points produces a valid χ_r since any choice of m distinct elements from $\{1...n\}$ is a valid χ_r, and by construction our method always arrives at such a choice. In fixing σ_p, we have used at least j elements, we have to choose the remaining $n - j$ fixed points from $m - j$ possible choices, and hence we have $\binom{n}{m}\binom{m}{j}\binom{n-m+j}{j}\binom{n-j}{m-j}j!$

ways of filling out the entire table. Therefore, for the case of superposition sampling, the covariance term in Eq. (12) is upper bounded by

$$
\sum_{p,q,r,s,\sigma,\tau} \mathrm{cov}(R(\xi_p, \sigma(\xi_q)), R(\chi_r, \tau(\chi_s)))
$$
$$
< \binom{n}{m}\binom{m}{j}\binom{n-m+j}{j}\binom{n-j}{m-j}j!\,\mathrm{var}(R). \quad (15)
$$

Note that since $\mathrm{cov}(R, R') < \mathrm{var}(R)$, we can simply ignore the variance term from Eq. (12) when obtaining an upper bound, if we allow the covariance term to run over all assignments of p, q, r, s, σ, and τ.

It is instructive to consider the two extremal cases. If $j = 0$ (i.e., the classical interference term), then σ must consist entirely of fixed points. This can only be achieved by setting $\xi_q = \xi_p$ and $\sigma = I$, leaving only the $\binom{n}{m}$ choices of ξ_p free. Similarly, we are free to choose χ_r, but then the pairing rules enforce that $\tau = I$ and $\chi_s = \chi_r$. For this reason, there are $\binom{n}{m}^2$ covariance terms for $j = 0$. In that case, $\mathrm{var}(c_j) < m!^2/N^{2m}$, since the factor $\binom{n}{m}^2$ drops out against the factor $c_p^4 = \binom{n}{m}^{-2}$, which arises from the normalization.

For the other extremal case, we have $j = m$. In this case, we again have $\binom{n}{m}$ ways of choosing ξ_p, but now all choices of ξ_q and σ are good (since we are upper bounding, we neglect those cases where our choice of σ introduces further fixed points). Therefore, we can pick ξ_p, ξ_q, and σ without any constraints, and we have $\binom{n}{m}^2 m!$ ways of doing so. However, we now have no freedom left at all in choosing τ, χ_r, and χ_s: We must chose $\tau = \sigma^{-1}$, $\xi_p = \chi_s$, and $\xi_q = \chi_r$. Therefore, there are $m!$ more permutations than in the case $j = 0$, but this drops out against the fact that these permutations are each a factor $m!$ smaller in value as well, as given by Eq. (10). Hence, for $j = m$ we also have $\mathrm{var}(c_j) = m!^2/N^{2m}$.

To discuss the simulability of superposition sampling with imperfections, we first observe that the ratio $\mathrm{var}(c_j)/\mathrm{var}(c_0)$ is unbounded. If we take the limit $m \ll n$ and $j \ll m$, then the variance reduces to $\mathrm{var}(c_j) = \binom{n-j}{m-j}\binom{n-m+j}{j}\binom{n}{m}^{-2}m!^2/N^{2m} \approx \binom{m}{j}m!^2/N^{2m}$, which becomes arbitrarily large around $j = m/2$ when m is increased and where the last equality holds at large n. This leaves open the possibility that superposition sampling is tolerant to imperfections.

However, numerical simulations suggest this is not the case: When converting from $\mathrm{var}(c_j)$ to $E(|c_j|)$, the inequality $\sqrt{\mathrm{var}(c_j)} > E(|c_j|)$ becomes less tight as n grows, resulting in coefficients $\mathrm{var}(c_j) \approx \mathrm{var}(c_0)$ for all j. We have observed this effect in simulations from $m = 2$ to $m = 6$, and from $n = 2$ to $n = 25$ (but not all combinations of both), averaging over 1000 Haar-random matrices in each case. In those same simulations, we also computed $\mathrm{var}(c_j)$ for each n, m, and j, and we observed that (accounting for some finite-size effects which give small corrections), our computations of $\mathrm{var}(c_j)$ as described above match our simulations. For large n, we find $\mathrm{var}(c_j) > \mathrm{var}(c_0)$, as suggested by Eq. (15). This shows that the limitations of our method are really in the conversion from variance to absolute moment. We leave this aspect of the problem for future study.

Next, we must rewrite Eq. (6) to a form which can be efficiently sampled from. In the form in which it is given, there are exponentially many terms in the sum over ξ_p and ξ_q, which means that a truncation at j will not lead to efficient sampling. This can be remedied by the following observation: For a given σ_p, we know that we will encounter every possible assignment of the $\binom{n}{m-j}$ remaining possible fixed points, and with equal weight (since all c_p are equal). However, each assignment occurs in a different pair of ξ_p and ξ_q. Grouping terms, we have

$$P = \sum_{j=0}^{m} \sum_{\sigma_p^j} \sum_{\rho} \mathrm{Perm}(M_{I_{\sigma_p}, \rho} \circ M_{\sigma_p, \rho}^{\dagger})$$
$$\times \cdots \mathrm{Perm}(M_{\mathrm{fp}}(\sigma_p, \rho))/(n - m - j)!, \qquad (16)$$

where I_{σ_p} denotes the elements of σ_p in sorted order (i.e., the partial identity permutation of the elements of σ_p), where ρ is a j partition of m, and where σ_p^j denotes all possible partial permutations of size j without fixed points. The matrix $M_{\mathrm{fp}}(\sigma_p, \rho)$ is an $n - 2j$ by $n - 2j$ matrix constructed by the following method: first, construct an m by n matrix of all norms squared, i.e., $M_{\mathrm{fp}, ij} = |M_{ij}|^2$, where M is defined according to Sec. I. Then, delete the j rows corresponding to ρ, and the $2j$ columns corresponding to σ_j^p. A matrix of size $m - j$ by $n - 2j$ remains. Next, pad out this matrix with $n - m - j$ rows of ones to make a square matrix. Laplace expansion along these rows shows that this produces precisely every way of picking fixed points that is complementary with a given choice of σ_j^p and ρ, and the factor $1/(n - m - j)!$ compensates for the spurious permanents of ones introduced by this procedure.

Sampling from Eq. (16) can be done efficiently: Since there are $\binom{n}{2j}$ ways of choosing σ_p, when truncating at a fixed $k = j$, this goes approximately as n^{2k}. Furthermore, the permanent of M_{fp} is a permanent of a positive matrix of size $n - j$, so it can be evaluated at a cost polynomial in n. Since the typical expected number of photons is proportional to n via $\bar{m} = \sin^2(\alpha)n$, this means that the scheme is polynomial in m.

Building on our numerical observations regarding the behavior of $E(|c_j|)$, we can compute the error level of truncation of the outer sum of Eq. (16) at a given k. This computation can be reduced to a known case if we note that we did not move any terms between different j when going from Eq. (6) to Eq. (16). Introducing a distinguishability factor x^j and truncating and summing the series, we find that superposition boson sampling is as loss tolerant as regular boson sampling, with an expected error of $E(d) = \sqrt{x^{2(k+1)}/(1 - x^2)}$, as reported previously for regular boson sampling.

We observe an interesting parallel between superposition boson sampling and Fock state boson sampling with loss: In lossy boson sampling, the input density matrix is of the form $\rho = (|0\rangle\langle 0| + |1\rangle\langle 1|)^{\otimes n}(|0\rangle\langle 0|^{\otimes N-n})$. In that case, there is a similar sum over $\binom{n}{m}$ sources as in superposition boson sampling, representing the $\binom{n}{m}$ ways in which the n sources could have given rise to the m observed photons. However, in this case it is a single sum, since the summation sum is incoherent (i.e., there are no interference terms). When comparing this case to superposition boson sampling, this amounts to considering only terms $\xi_p = \xi_q$. Interestingly enough, lossy Fock state boson sampling is known to be classically simulable because the c_j decrease exponentially with j [29]. Therefore, the following physical picture arises: If we think of superposition boson sampling as a combination of loss terms (that is, terms which would show up in lossy Fock state boson sampling, where $\xi_p = \xi_q$) and cross terms (where $\xi_p \neq \xi_q$), then the loss terms are skewed to low photon number interference, while the cross terms are skewed to high photon number interference, thereby precisely canceling out the deleterious effects of the loss terms.

We conclude this section with a few remarks on the viability of superposition boson sampling as a quantum advantage demonstration. The observation that the coefficients of perfect superposition boson sampling do not decay to zero shows that it is not simulable by our method in the ideal case and that the state $|\Psi_{\mathrm{sbs}}\rangle$ might therefore be a valid state to use for a demonstration of a quantum advantage. However, the fact that this state is no more loss tolerant than a regular Fock state means that there is no obvious reason to use this state to achieve a quantum advantage. Rather, the interest lies in quantum simulation: It is an open problem of which physical systems can be simulated with photonics, and any new input state which can be shown to have computationally interesting properties is a potential addition to our simulation arsenal. The potential arrival of another class of states to sample from is valuable, since this could broaden the range of problems which can be simulated in photonics. The question of whether there are natural problems which can be simulated using superposition states is of high interest.

We stress that these results do not constitute a hardness proof; however, the fact that this state passes the most stringent classical simulation criteria known is some cause for optimism. It is hoped that these results will spur interest in a hardness proof for superposition sampling.

V. GAUSSIAN BOSON SAMPLING

Finally, we turn our attention to Gaussian boson sampling (GBS). In GBS, the input quantum state is no longer factorizable across modes, but is instead factorizable over pairs of optical modes: $|\psi\rangle = \cosh(r)^{-1} \sum_{j=0}^{\infty} [-e^{i\phi} \tanh(r)]^j |j, j\rangle$, and $|\Psi\rangle = |\psi\rangle^{\otimes n}|0\rangle^{\otimes N-n}$, where r is a parameter that measures the strength of the optical squeezing, which we shall assume to be equal for all sources. Again, we have two tasks: first, to construct the set of ξ_p and c_p, and apply the pairing rules and Eqs. (10) and (12) to determine the simulability criterion, and second to rewrite Eq. (6) to enable efficient sampling.

We will do these tasks in order. The strategy will be to build on the results from SBS. We will show that in the limit of weak squeezing, GBS reduces to SBS with a single additional constraint. This will enable us to reuse many of the results from SBS. We will also only consider the limit of weak squeezing (so that we can neglect multiphoton components). This is justified since this is the limit in which the hardness of GBS has been most rigorously demonstrated.

Without loss of generality, we exclusively consider two-mode squeezed vacuum states here. There also exist single-mode squeezed vacuum states, in which photon pairs are emitted into a single mode. In boson sampling, the two cases are

equivalent, as a layer of 50:50 beam splitters converts pairs of single-mode squeezed states into two-mode squeezed states, and vice versa. If we think of U as made up of a layer of 50:50 beam splitters and a remaining matrix U', by the definition of the Haar measure, the probability density of drawing U is as large as drawing U'. Hence, every boson sampling experiment with single-mode squeezers can be thought of as equivalent to an equiprobable boson sampling experiment with two-mode squeezers. This shows that the two are equivalent.

We begin by determining the set of ξ. We again consider the case of n modes containing photons and m detected photons. Compared to SBS, there is an additional constraint on ξ, which is due to the bipartite nature of the state: When choosing a photon from a particular mode, we must also chose a photon from the conjugate mode. We shall refer to such a pair of photons as a biphoton. Second, the multipair terms in $|\psi\rangle$ mean we also have an additional options in picking ξ, because we may now pick more than one photon pair from a given pair of modes. Therefore, the set of ξ_p contains $\binom{n/2-m/2-1}{m/2}$ elements, reflecting the fact that we are picking $m/2$ biphotons with replacement. To determine the c_p, we note that due to the exponential prefactors in $|\psi\rangle$, all ξ_p occur with the same probability amplitude, namely $\tanh(r)^{m/2}$; the phases may be absorbed into the action of the interferometer. However, their multiplicities are no longer equal: $c_p = \tanh(r)^{m/2}/\sqrt{\mu(\xi_p)}$.

To simplify our calculations, we will focus on the case of weak squeezing. We stress that this is also the limit in which the strongest arguments for hardness of GBS (via reduction to scattershot boson sampling) hold. If $(m/2)^2 \ll n/2$, then the fraction of ξ_p terms where we pick two biphotons from the same source can be made arbitrarily small. In that case, we can ignore the arbitrarily small fraction with nonunit multiplicity and renormalize the c_p to postselect on a particular photon number outcome (as we did for SBS), giving $c_p = \sqrt{\binom{n/2}{m/2}}$. Hence, in the limit of low squeezing, there are $\binom{n/2}{m/2}$ choices for ξ, all equiprobable.

Note that the above shows that in the limit of sufficiently weak squeezing, GBS strongly resembles SBS. The only difference is the additional constraint of picking biphotons. This fact will enable us to use the results from SBS with minor modifications.

We again seek to compute $\mathrm{var}(c_j)$ by looking for pairs of terms with matching permutations. For $j = 0$ and $j = m$, our results from SBS carry over in a straightforward fashion. For $j = 0$, we again have complete freedom in chosing ξ_p and χ_r, but the remaining variables are fixed by the pairing rules ($\sigma = \tau = I$, $\xi_p = \xi_q$, $\chi_s = \chi_r$). Therefore, we have $\binom{n/2}{m/2}^2$ ways of creating matching R, and $\mathrm{var}(c_j) = m!^2/N^{2m}$, as before. Similarly for $j = m$, all ξ_p, ξ_q, and σ are allowed, but their combination completely fixes τ, χ_r, and χ_s. Hence, we have $\mathrm{var}(c_m) < m!^2/N^{2m}$ as before.

For the case of arbitrary j, the situation is a little more complex. The biphoton nature of the light field introduces additional constraints. In particular, if we have biphotons which are not wholly included in either the fixed points or in σ_p, then these reduce the number of possible covariance pairings, since the presence of exactly one photon of a pair in

σ_p implies the presence of its partner photon in the fixed points (since we must respect the biphoton pairing). Moreover, the presence of an "incomplete" biphoton in σ_p implies that the same biphoton must occur in τ_p as well, thereby partially determining the fixed points of τ as well.

This is best illustrated with an example. Consider the case of $j = 2$, for large m and n. There are two options for constructing σ_p: Either σ_p contains two complete biphotons, one in each row, for example, $\sigma_p = \binom{1\ 2}{3\ 4}$, or it contains two incomplete biphotons, for example, $\sigma_p = \binom{1\ 3}{3\ 1}$. Note that these are the only options allowed if we take into account the fact that the remaining elements of σ must be fixed points. For example, if we had $\sigma_p = \binom{1\ 3}{2\ 1}$, then we cannot complete the permutation using only fixed points, since the presence of a 1 in ξ_p implies that mode 2 must be selected as well, but since mode 2 already enters in $\sigma(\xi_q)$, it cannot also form a fixed point. Furthermore, for the case of $\sigma_p = \binom{1\ 3}{3\ 1}$, we know that both τ and σ must include the fixed points 2 and 4. This example illustrates that the constraint placed on σ and τ is that biphotons which are not completely in σ_p (or τ_p) must occur both in σ and τ.

Since biphotons which straddle σ_p and the fixed points determine a fixed point in both σ and τ, while biphotons contained in the fixed points can be independently chosen between σ and τ, the leading contribution to the covariance (expressed in powers of n and m) in Eq. (12) at large m and n will be the one with the fewest straddling fixed points. Therefore, for even j, it is the one where no fixed points are straddling σ_p and the fixed points, and for odd j, there will be only one. This removes the constraints between fixed points and σ_p introduced by the biphotons, meaning that we can follow the argument from SBS to count the number of variance pairs, substituting $n \rightarrow n/2$, $m \rightarrow m/2$, $j \rightarrow \lceil j/2 \rceil$ in all terms related to filling the fixed points and σ_p. This means that the number of terms with covariance is $\#R = j! \binom{n/2}{m/2}\binom{m/2}{\lceil j/2 \rceil}\binom{n/2-m/2+\lceil j/2 \rceil}{\lceil j/2 \rceil}\binom{n/2-\lceil j/2 \rceil}{m/2-\lceil j/2 \rceil}$, leading to a covariance of $\mathrm{cov}(c_j) = \#R\,\mathrm{var}(R) = \binom{n}{m}^{-1}\binom{m/2}{\lceil j/2 \rceil}\binom{n/2-m/2+\lceil j/2 \rceil}{\lceil j/2 \rceil}\binom{n/2-\lceil j/2 \rceil}{m/2-\lceil j/2 \rceil}\binom{m}{j}(m-j)!^2$

$j!^2/N^{2m} < m!^2/N^{2m}$ for all j, using the inequality $\binom{m}{j} > \binom{m/2}{j/2}^2$.

From this, we conclude that, like superposition boson sampling, the higher order interference terms of Gaussian boson sampling are at least as sensitive to photon distinguishability as in regular boson sampling. Note that unlike superposition boson sampling this derivation does not rely on numerical results. When we numerically evaluate Eq. (12) for Gaussian boson sampling, we observe a sawtooth pattern in the c_j, where odd j are strongly suppressed. This is a reflection of the biphoton nature of the light, and it suggests that further analysis might develop a tighter bound.

Finally, we turn to the issue of rewriting Eq. (6) into a form which can be efficiently sampled from. We will see that it is possible to regroup terms to gather identical quantum interference terms, as we did for the case of SBS. Furthermore, this results in polynomially many quantum interference terms of fixed size, which means they can be efficiently computed. Surprisingly, it is the classical interference terms which are problematic: We obtain an expression for the classical

interference terms for which no known classical algorithm exists to compute the output probability.

We begin by following the same strategy as with SBS, by grouping terms which correspond to the same quantum interference process (i.e., the same permuted part of the permutation σ_p), but which have different fixed points. This can be done straightforwardly:

$$
P = \sum_{j=0}^{m} \sum_{\sigma_p^j} \sum_{\rho} \mathrm{Perm}\big(M_{I_{\sigma_p},\rho} \circ M_{\sigma_p,\rho}^{\dagger}\big) \sum_{\bar{\sigma}_p^j} \mathrm{Perm}\big(|M_{\sigma_p^j}|^2\big),
$$

(17)

where $\bar{\sigma}_p^j$ is the complement of σ_p^j, in the sense that it contains a way of choosing fixed points such that the combination of σ_p^j and $\bar{\sigma}_p^j$ forms a "legal" permutation, i.e., one that corresponds to a particular choice of ξ_p, ξ_q, and σ. Note that there are multiple ways of completing a permutation from its unfixed points, and hence there is a sum over the various possible $\bar{\sigma}_p^j$ in Eq. (17).

To assess the complexity of sampling from Eq. (17), many of the arguments from Eq. (16) carry over: There are $\binom{m}{2j}$ ways of choosing σ_p^j, and hence there are polynomially many quantum interference terms of size at most j. However, the classical interference term $\sum_{\bar{\sigma}_p^j} \mathrm{Perm}(|M_{\sigma_p^j}|^2)$ cannot be computed efficiently. This can be seen already when considering $j = 0$: In this case, there are $\binom{n-m/2+1}{m/2}$ ways of picking sets of m fixed points, which grows exponentially with m.

As we have done with Eq. (16), we will try to recast each classical interference term in Eq. (17) as a permanent of a positive matrix. We will not succeed because of the biphoton nature of the light field, which places additional constraints on $\bar{\sigma}_p^j$. This is in contrast to SBS, where every choice of fixed points was allowed.

For biphotons which belong half to a fixed point, half to σ_p, a straightforward construction similar to that described for SBS can be followed. For the fixed points where both parts of the biphoton belong to the fixed points, a more elaborate construction is necessary, which we will show for the case $j = 0$. Define an $m + n$ by $m + n$ matrix M_{gfp} :

$$
M_{\mathrm{gfp}} = \begin{pmatrix} \bar{M} & 0 \\ P & S \end{pmatrix}
$$

(18)

with $\bar{M}$ an m by n matrix containing the mod-squared values of the elements of M. P is an n by n matrix of the following form:

$$
P = \begin{pmatrix}
1 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\
1 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 1 & \cdots & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 1 & \cdots & 0 & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\
0 & 0 & 0 & 0 & \cdots & 1 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & \cdots & 1 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & \cdots & 0 & 0 & 1 & 1 \\
0 & 0 & 0 & 0 & \cdots & 0 & 0 & 1 & 1
\end{pmatrix},
$$

(19)

and S and n by m matrix of the following form:

$$
S = \begin{pmatrix}
1 & 1 & \cdots & 1 & 1 \\
-1 & -1 & \cdots & -1 & -1 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
1 & 1 & \cdots & 1 & 1 \\
-1 & -1 & \cdots & -1 & -1
\end{pmatrix}.
$$

(20)

The remaining quadrant of M_{gfp} is composed of zeros. We can show that $\mathrm{Perm}(M_{\mathrm{gfp}})/2^m m! = \sum_{\bar{\sigma}_p^j} \mathrm{Perm}(|M_{\sigma_p^j}|^2)$ by recursively Laplace expansion along the ith and $(i+1)$-st rows of P and S simultaneously, for odd i. If we do this, the only 2 minors which have a nonzero prefactor are either those where we pick columns i and $i+1$, which corresponds to deleting a pair of columns from $\bar{M}$, or those where we pick any pair of columns from S, which only deletes zeros in the upper m rows. Either of these picks up a factor of 2, since $\mathrm{Perm}\big(\begin{smallmatrix} 1 & 1 \\ 1 & 1 \end{smallmatrix}\big) = \mathrm{Perm}\big(\begin{smallmatrix} 1 & 1 \\ -1 & -1 \end{smallmatrix}\big) = 2$. Any choice which mixes P and S has a zero prefactor, since $\mathrm{Perm}\big(\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\big) = 0$. Therefore, we can see that this construction respects the pairing of modes which arises from the biphoton nature of the light field: Either it deletes a pair of photons from consideration, or it "skips a turn" and deletes no photons. This ensures that when the auxilliary matrices P and S are exhausted via repeated pairwise Laplace expansion, only pairs of terms in $\bar{M}$ which correspond to the same biphoton remain, which completes the proof.

The reason why this expression for M_{gfp} does not result in an efficient algorithm is that in order to cast the classical interference as the permanent of a single matrix M_{gfp}, we needed to introduce negative numbers in that matrix to make this construction work. This means that we cannot apply the approximation algorithm of Jerrum et al. [48] to efficiently approximate the permanent of M_{gfp}, since that algorithm requires a matrix with only positive elements.

We leave the question of whether such an algorithm to efficiently approximate M_{gfp} exists as an open problem for the theory of matrix permanents. We can, however, show that the worst-case complexity of approximating M_{gfp} is at least as that of approximately counting perfect matchings in an arbitrary graph, which is an open problem [51]. This reduction was found by M. Jerrum [52]. The construction is as follows: Consider a graph with m vertices and $n/2$ edges, where m and n define the size of $\bar{M}$ as above. Construct $\bar{M}$ as follows: Index the rows by the vertices and the columns (pairwise) by the edges. Passing along each pair of columns, place a 1 whenever the incident vertex corresponding to that column is identical to the vertex in the row indexation, and a 0 otherwise. If we now embed $\bar{M}$ in M_{gfp} as above and evaluate the permanent, each nonzero term in the permanent either contains both vertices of a given edge or neither. Hence, each nonzero term contains a different set of edges of size $m/2$, which do not share any vertices, which is a perfect matching in m. We note that this is not the first time Gaussian boson sampling has been associated with counting perfect matchings in arbitrary graphs: A scheme to embed this quantity in the output of a Gaussian boson sampler has been proposed previously [53].

Unfortunately, this worst-case result is not a full answer: We are interested in the average case complexity of computing

M_{gfp} when $\bar{M}$ is a matrix of independent and identically distributed Gaussian variables, and the above argument requires $\bar{M}$ to be of a very specific form. It would be fine for our purposes if the approximation algorithm failed for some values of $\bar{M}$, provided that set had measure zero under the probability distribution of $\bar{M}$. On physical grounds, however, it is entirely possible that such an approximate algorithm exists, since the probability which we are interested in corresponds to classical transmission, and this also holds for all cases which arise when $j \neq 0$, which can be accounted for with a construction analogous to the one presented above. We note that nothing in the hardness proof of the original Fock state boson sampling proposal hinges on the existence of the JSV algorithm. However, as we noted above, at finite levels of imperfection, all quantum interference effects are accounted for by the quantum interference terms, which are of size k or smaller. It would be strange if the hardness of simulating boson sampling, which we think of as an intrinsically quantum phenomenon, depends on the hardness of computing a classical transmission probability.

We also briefly consider the scenario where such an efficient algorithm for approximating M_{gfp} could be proven not to exist. In that case, this would of course not constitute a hardness proof for noisy GBS, but it would show that the state of the art approach to simulating boson sampling does not work for this specific case. This would substantially enhance the interest in Gaussian boson sampling as a quantum advantage demonstration.

Another way to solve the difficulty of the classical interference terms would be to construct an algorithm that directly samples from an approximation of the state taking into account the effects of noise, instead of approximating an output probability and feeding that to an Markov Chain Monte Carlo sampler. This was recently achieved for Fock state boson sampling [30], although in the resulting algorithm, the truncation point k is a rising function of m, meaning that the algorithm is not efficient. Since $\mathrm{Perm}(M_{\mathrm{gfp}})$ represents a convex sum over positive terms, such a sampling algorithm would remove the need to compute the entire sum. Instead, a sample from the classical interference term would suffice, which can be produced readily by selecting one of the terms and simulating sending photons through the interferometer one by one [5].

Finally, we note that our noise model for Gaussian boson sampling does not capture all sources of imperfections. There are two effects which we did not consider and which we would like to highlight. First, we considered a limit in which double pair emission is negligible, i.e., in which all multiplicities are equal to one. Since double pairs reduce the complexity of the sampling problem by introducing repeated rows in the permanent, they constitute an imperfection. The second imperfection which we did not consider is an effect which we name *layer mixing*. This occurs when sampling from any input state with an indeterminate photon number, in the presence of loss. At the beginning of our derivation, we assumed that we knew the number of photons generated by the sources. In the presence of loss, this knowledge is imperfect; there is mixedness between the process where m photons are generated and none are lost, $m + 1$ and 1 is lost, and so on. This effect, which is not present for Fock state sampling, should not be confused with the quantum interference between photons originating

from different sources. The presence of layer mixing suggests that perhaps both GBS and SBS are more susceptible to loss than regular Fock state boson sampling. We leave these issues to future study.

However, since all previous examples of noise analysis in boson sampling have shown that noise sources compound, we speculate that the addition of these sources of noise will not reduce, but rather increase, the simulability of Gaussian boson sampling under imperfections. This means that until further results on this topic arise, it is not unreasonable to assume that our bound will also hold outside of the weak-squeezing regime for which it was derived.

VI. CONCLUSIONS

In this section, we briefly reprise the results of this work. First, we have demonstrated how to extend the analysis of the effect of imperfections on boson sampling to the case of arbitrary input states. We used this method to show that for two particular choices of input states, namely Gaussian states and superposition states, the output probability without noise cannot be efficiently approximated by our methods. This is a necessary but not sufficient condition for computational hardness of sampling problems using these states. For superposition sampling, this is an indication that this problem is of interest for sampling.

Then, we introduced the effects of noise. Using a combination of analytic results and numerical simulations, we showed that the effect of noise on both these sampling protocols is identical to that of regular boson sampling, at least for the particular type of noise (distinguishability) under consideration here; they are no more or less resilient to this kind of noise than Fock state boson sampling is. For superposition states, this leads immediately to a classical algorithm which can efficiently simulate imperfect superposition boson sampling. For Gaussian states, the existence of such an algorithm depends on a not-implausible conjecture in the theory of matrix permanents.

As a final warning not to take our results as a hardness proof for either Gaussian boson sampling or superposition boson sampling, we give an explicit example [54] where there is no computational complexity but where our simulation strategy does not work. This is the case where weak coherent states are incident on the interferometer. In this case, phases and amplitudes can be propagated efficiently through the interferometer, no entanglement builds up, and we can sample efficiently from the output distribution. However, if we follow the analysis outlined above, we would find that perfect coherent state sampling is not susceptible to our algorithm. In this case, the very first step in our approach, namely the projection onto Fock states, destroys the structure (namely the coherent state amplitudes and phases) which enables efficient classical simulation. This emphasizes the point that failure of any one simulation strategy is no guarantee that all simulation strategies will fail.

Finally, we address some open problems. There are two major open problems raised by this work. The first is to find a better way to compute $E(|c_j|)$, which does not depend on the inequality $\sqrt{\mathrm{var}(c_j)} > E(|c_j|)$, which is insufficiently tight for our purposes. The second open problem is how to compute the permanent of M_{gfp} efficiently.

[1] R. de Wolf, Ethics Inf. Technol. **19**, 271 (2017).

[2] J. Preskill, Quantum **2**, 79 (2018).

[3] A. Harrow and A. Montanaro, Nature (London) **549**, 203 (2017).

[4] 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 *et al.*, Nature (London) **574**, 505 (2019).

[5] S. Aaronson and A. Arkhipov, Theory Comput. **9**, 143 (2013).

[6] P. Clifford and R. Clifford, arXiv:1706.01260.

[7] A. Neville *et al.*, Nat. Phys. **13**, 1153 (2017).

[8] J. Wu *et al.*, Nat. Sci. Rev. **5**, 715 (2018).

[9] M. A. Broome *et al.*, Science **339**, 794 (2012).

[10] J. B. Spring *et al.*, Science **339**, 798 (2012).

[11] M. Tillmann *et al.*, Nat. Photon. **7**, 540 (2013).

[12] A. Crespi *et al.*, Nat. Photon. **7**, 545 (2013).

[13] J. Carolan *et al.*, Science **349**, 711 (2015).

[14] M. Bentivegna *et al.*, Sci. Adv. **1**, e1400255 (2015).

[15] H. Wang *et al.*, Nat. Photon. **11**, 361 (2017).

[16] H. Wang, W. Li, X. Jiang, Y.-M. He, Y.-H. Li, X. Ding, M.-C. Chen, J. Qin, C.-Z. Peng, C. Schneider *et al.*, Phys. Rev. Lett. **120**, 230502 (2018).

[17] H. Wang *et al.*, Phys. Rev. Lett. **123**, 250503 (2019)..

[18] V. S. Shchesnovich, Phys. Rev. A **89**, 022333 (2014).

[19] V. S. Shchesnovich, Phys. Rev. A **91**, 063842 (2014).

[20] V. S. Shchesnovich, Phys. Rev. A **91**, 013844 (2015).

[21] V. Shchesnovich, J. Phys. A: Math. Theor. **50**, 505301 (2017).

[22] V. Shchesnovich, Sci. Rep. **7**, 31 (2017).

[23] V. Tamma and S. Laibacher, Quant. Inform. Proc. **15**, 1241 (2015).

[24] V. Tamma and S. Laibacher, Phys. Rev. Lett. **114**, 243601 (2015).

[25] M. C. Tichy, Phys. Rev. A **91**, 022316 (2015).

[26] S. Laibacher and V. Tamma, Phys. Rev. A **98**, 053829 (2018).

[27] G. Kalai and G. Kindler, arXiv:1409.3093.

[28] J. J. Renema, A. Menssen, W. R. Clements, G. Triginer, W. S. Kolthammer, and I. A. Walmsley, Phys. Rev. Lett. **120**, 220502 (2018).

[29] J. Renema, V. Shchesnovich, and R. Garcia-Patron, arXiv:1809.01953.

[30] A. E. Moylett, R. Garcia-Patron, J. J. Renema, and P. S. Turner, Quantum Sci. Technol. **5**, 015001 (2020).

[31] V. S. Shchesnovich, Phys. Rev. A **100**, 012340 (2019).

[32] V. Shchesnovich, arXiv:1904.02013.

[33] Y.-M. He *et al.*, Nat. Nanotechnol. **8**, 213 (2013).

[34] X. Ding *et al.*, Phys. Rev. Lett. **116**, 020401 (2016).

[35] N. Somaschi *et al.*, Nat. Photon. **10**, 340 (2016).

[36] R. Boyd, *Nonlinear Optics* (Academic Press, New York, 2008).

[37] H. Wang *et al.*, Nat. Photon. **13**, 770 (2019).

[38] J. F. Clauser, Phys. Rev. D **9**, 853 (1974).

[39] S. Aaronson [https://www.scottaaronson.com/blog/?p=1579].

[40] A. P. Lund, A. Laing, S. Rahimi-Keshari, T. Rudolph, J. L. OBrien, and T. C. Ralph, Phys. Rev. Lett. **113**, 100502(R) (2014).

[41] C. S. Hamilton *et al.*, Phys. Rev. Lett. **119**, 170501 (2017).

[42] L. Chakhmakhchyan and N. J. Cerf, Phys. Rev. A **96**, 032326 (2017).

[43] H. Qi, D. J. Brod, N. Quesada, and R. Garcia-Patron, Phys. Rev. Lett. **124**, 100502 (2020)..

[44] J. C. Loredo *et al.*, Nat. Photonics **13**, 803 (2019).

[45] V. Shchesnovich, arXiv:1712.03191.

[46] S. Scheel and S. Buhmann, Acta Physica Slovaca. **58**, 675 (2008).

[47] J. Carolan *et al.*, Nat. Photon. **8**, 621 (2014).

[48] M. Jerrum, A. Sinclair, and E. Vigoda, *Journal of the ACM* 51, 671 (2004).

[49] D. Glynn, Eur. J. Comb. **31**, 1887 (2010).

[50] H. Ryser, *Combinatorial Mathematics*, Carus Mathematical Monographs, Vol. 14 (Mathematical Association of America, New York, 1963).

[51] M. Rudelson, A. Samorodnitsky, and O. Zeitouni, Ann. Prob. **44**, 2858 (2016)..

[52] M. Jerrum (private communication).

[53] K. Brádler, P.-L. Dallaire-Demers, P. Rebentrost, D. Su, and C. Weedbrook, Phys. Rev. A **98**, 032310 (2018).

[54] This example is due to R. Garcia-Patron (private communication).