



Titre: Title:	The Complexity of Bipartite Gaussian Boson Sampling	
Auteurs: Authors:	Daniel Grier, Daniel J. Brod, Juan Miguel Arrazola, Marcos Benicio de Andrade Alonso, & Nicolás Quesada	
Date:	2022	
Type:	Article de revue / Article	
Référence: Citation:	Grier, D., Brod, D. J., Arrazola, J. M., de Andrade Alonso, M. B., & Quesada, N. (2022). The Complexity of Bipartite Gaussian Boson Sampling. Quantum, 6, 44 pages. https://doi.org/10.22331/q-2022-11-28-863	

Document en libre accès dans PolyPublie Open Access document in PolyPublie

URL de PolyPublie: PolyPublie URL:	https://publications.polymtl.ca/52701/
Version:	Version finale avant publication / Accepted version Révisé par les pairs / Refereed
Conditions d'utilisation: Terms of Use:	Creative Commons Attribution 4.0 International (CC BY)

Document publié chez l'éditeur officiel Document issued by the official publisher

Titre de la revue: Journal Title:	Quantum (vol. 6)
	Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften
URL officiel: Official URL:	https://doi.org/10.22331/q-2022-11-28-863
Mention légale: Legal notice:	This Paper is published in Quantum under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Copyright remains with the original copyright holders such as the authors or their institutions.

The Complexity of Bipartite Gaussian Boson Sampling

Daniel Grier^{1,2}, Daniel J. Brod³, Juan Miguel Arrazola⁴, Marcos Benicio de Andrade Alonso³, and Nicolás Quesada⁵

Gaussian boson sampling is a model of photonic quantum computing that has attracted attention as a platform for building quantum devices capable of performing tasks that are out of reach for classical devices. There is therefore significant interest, from the perspective of computational complexity theory, in solidifying the mathematical foundation for the hardness of simulating these devices. We show that, under the standard Anti-Concentration and Permanent-of-Gaussians conjectures, there is no efficient classical algorithm to sample from ideal Gaussian boson sampling distributions (even approximately) unless the polynomial hierarchy collapses. The hardness proof holds in the regime where the number of modes scales quadratically with the number of photons, a setting in which hardness was widely believed to hold but that nevertheless had no definitive proof.

Crucial to the proof is a new method for programming a Gaussian boson sampling device so that the output probabilities are proportional to the permanents of submatrices of an arbitrary matrix. This technique is a generalization of Scattershot BosonSampling that we call BipartiteGBS. We also make progress towards the goal of proving hardness in the regime where there are fewer than quadratically more modes than photons (i.e., the high-collision regime) by showing that the ability to approximate permanents of matrices with no repetitions. The reduction suffices to prove that GBS is hard in the constant-collision regime.

1 Introduction

The quest for quantum computational advantage has given rise to a surprisingly fruitful relationship between computer science and physics: theorems provide a foundation for experiments, while practical considerations set challenges for new mathematics. Consider for instance the role of BosonSampling [1]. It gave strong complexity-theoretic evidence that even a weak photonic device could perform a task that is classically intractable. This work motivated future experimental demonstrations [2, 3, 4, 5] that inspired new theoretical models [6, 7], which in turn resulted in further experiments [8, 9, 10].

Gaussian boson sampling (GBS) is a paradigm where a Gaussian state of light is prepared, then measured in the photon-number basis [7, 11, 12]. This approach offers several benefits. First, Gaussian states can be prepared by a combination of squeezing, displacement, and linear interferometers, which can in principle be applied deterministically and implemented with nanophotonic integrated circuits [13]. This means they can potentially be mass-produced and scaled rapidly [14, 15]. Moreover, the inclusion of squeezing and displacements allows more versatility in programming GBS devices, a property that is leveraged in several GBS-based algorithms [16, 17, 18, 19, 20, 21, 22].

¹Institute for Quantum Computing, University of Waterloo, Canada

²Department of Computer Science and Engineering and Department of Mathematics, University of California, San Diego, US

³Instituto de Física, Universidade Federal Fluminense, Niterói, RJ, 24210-340, Brazil

⁴Xanadu, Toronto, ON, M5G 2C8, Canada

⁵Department of Engineering Physics, École Polytechnique de Montréal, Montréal, QC, H3T 1JK, Canada

GBS has already been used to claim quantum computational advantage [23, 24, 25], and there are several more proposals for hard-to-simulate GBS experiments [6, 7, 26]. That said, we believe these experiments reveal that significant progress is still required to bridge the gap between our theoretical hardness arguments and what is currently achievable in the lab. For example, state-of-art GBS experiments consisted of 216 modes with up to 125 photons [23] and 144 modes and at most 113 clicks [25], but all hardness arguments currently assume that the number of modes is at least quadratic in the number of photons, and sometimes worse.¹

Furthermore, the underlying physics of GBS is such that the probability of a given output state is described by the hafnian, a matrix function similar, but not identical to, the permanent. Because of this, new conjectures tailored to this paradigm are sometimes required—see, for example, the Hafnian-of-Gaussians conjecture in [7] which parallels the Permanent-of-Gaussians conjecture in [1]. While the goal of this paper is not to compare the merits of the individual conjectures, we do feel that having fewer standard conjectures is generally preferable.

This paper introduces Bipartite Gaussian Boson Sampling² (BipartiteGBS) as a new method for programming a GBS device which will begin to address some of these challenges. The key property of BipartiteGBS experiments is that the output probabilities are proportional to the permanents of submatrices of arbitrary matrices. Contrasted with traditional BosonSampling, where the output probabilities are dictated by permanents of submatrices of a unitary matrix, BipartiteGBS can be seen as a powerful new tool on which to build hardness-of-simulation arguments. In particular, this paper will focus on the hardness of approximately sampling from BipartiteGBS distributions with a classical device.

As it turns out, our construction is a strict generalization of Scattershot BosonSampling [6], a different GBS setup for which the output probabilities are given by permanents of unitary matrices. Because of this, the computational hardness of Scattershot BosonSampling can be rooted in the same conjectures on which the hardness of BosonSampling is based—namely, that Gaussian permanent estimation is #P-hard and that Gaussian permanents anti-concentrate. However, this also means that Scattershot BosonSampling inherits the same technical caveats of BosonSampling. In particular, to guarantee that the $n \times n$ submatrices of an $m \times m$ unitary appear Gaussian, Aaronson and Arkhipov require that $m = \omega(n^5)$. Therefore, all their hardness proofs are technically within this regime. To be clear, it is widely assumed that $m = \omega(n^2)$ is sufficient, but to the authors' knowledge, there is currently no definitive proof.³ Even more general Gaussian boson sampling protocols suffer from this problem, and recent approximate average-case hardness proofs for GBS must conjecture directly that quadratically-many modes suffice [26].

The main technical contribution of this paper is to show that BipartiteGBS can be used to close this loophole. That is, the hardness of GBS can be based on the exact same set of conjectures as BosonSampling, while also working in the regime where the number of modes m is quadratic in the expected number of photons $\langle n \rangle$.⁴ Formally, we prove the following theorem:

¹This mismatch between theory and experiment occurs because of photon loss in the interferometer [27, 28]. For experiments requiring Haar random interferometers, each mode must be able to exchange light with every other mode since with probability one every entry of a random Haar unitary matrix is nonzero. In particular, since random m-mode interformeters are built from $\Theta(m^2)$ 2-mode beamsplitters [29, 30, 31, 32], this implies that the depth of the circuit implementing the interferometer is proportional to the width. If there is a fixed transmission per beamsplitter layer $\eta \leq 1$ the total transmission of the interferometer will decay exponentially as η^m . It is then clear that it is more desirable to have m scaling with n and not n^2 to have a significant fraction of the photons arriving into the detectors.

² "Bipartite" refers to the Husimi covariance matrix of our output states, which is characterized by a bipartite adjacency matrix.

³Perhaps the closest result is that of Jiang [33] who shows that the $n \times n$ submatrices of $m \times m$ real orthogonal matrices converge in total variation to real Gaussian matrices whenever $m = \omega(n^2)$. Unfortunately, Jiang does not bound the *rate* of this convergence, which is required by the BosonSampling hardness arguments. Moreover (though perhaps less importantly), the BosonSampling arguments are based on the submatrices of complex unitary matrices rather than real orthogonal ones. Finally, Jiang proves that $m = O(n^2)$ does not suffice, but we note that the results of this paper will hold even when $m = \Theta(n^2)$.

 $^{^4}$ Unlike BosonSampling where the number of photons is fixed, the number of photons n in a GBS experiment is itself

Theorem 1. Suppose there is a classical oracle \mathcal{O} which approximately samples from a BipartiteGBS distribution with $m = \Theta(\langle n \rangle^2)$. Then, $\#P \subseteq \mathsf{FBPP}^{\mathsf{NP}^{\mathcal{O}}}$ assuming the Permanent-of-Gaussians Conjecture and the Permanent Anti-Concentration Conjecture.

In other words, assuming the BosonSampling conjectures, there is no efficient classical algorithm for BipartiteGBS in this regime unless there is a collapse in the polynomial hierarchy that complexity theorists consider to be extremely unlikely. Unsurprisingly, the proof of this theorem will leverage the fact that the output probabilities of BipartiteGBS experiments are based on permanents of arbitrary matrices. Since the hardness of BosonSampling is based on Gaussian permanents anyway, an obvious-in-retrospect idea is to simply *start* with those Gaussian matrices. Note that, as in all photonic experiments, the probabilities are actually governed by submatrices of the transition matrix. Clearly, however, the submatrices of a Gaussian matrix—i.e., a matrix for which each entry is an i.i.d. complex Gaussian number—are also Gaussian. This trivializes the "Hiding Lemma" often required in other hardness arguments.⁵

That said, the proof of Theorem 1 is not itself trivial. In particular, we must prove that the Bi-partiteGBS experiments that we ask the classical oracle to simulate have sufficient probability mass on those outcomes which correspond to #P-hard permanents. To do so, we bound important normalization factors in the output distribution using results from random matrix theory on ensembles of random Gaussian matrices. While our proof unsurprisingly borrows many ideas from the original BosonSampling hardness argument [1], it can be entirely understood without direct reference to it, and we suspect that many will find our rigorous hardness proofs of GBS to be preferable to some in the existing literature.

To complement our formal analysis of the normalization factors (which are sufficient to obtain Theorem 1), we present a more heuristic analysis of other aspects of our experiment that might be relevant to experimentalists hoping to claim quantum advantage (most likely with a more speculative set of conjectures). In particular, we give formulas for the expectation and variance of the number of "clicks" in the output distribution—that is, the number of modes that contain at least one photon when measured. Intuitively, this is an important quantity for hardness arguments since the probability of any particular outcome is proportional to a permanent of a matrix whose rank is equal to the number of clicks in that outcome. Since classical intractability is tied to the complexity of these permanents, and we know that permanents of low-rank matrices have efficient classical algorithms [34], we would like to avoid distributions with low click rates. Thankfully, we prove that this is generally not the case by showing that the expected number of clicks in a BipartiteGBS experiment with Gaussian transition matrices is the harmonic mean of the number of modes and expected number of photons:

$$\frac{2}{\frac{1}{\mathbb{E}[\langle n \rangle]} + \frac{1}{m}}.\tag{1}$$

So, for example, even when we expect that there are only linearly-many more modes than photons, we have that the expected number of clicks is linear. This formula is obtained by assuming that the singular values of a Gaussian matrix are drawn independently and exactly from the quarter-circle law.⁶

a random variable that is based on the squeezing parameters of the system. Therefore, we cannot guarantee that the number of modes m is actually quadratic in the number of photons n since there may be some small probability for which there are many more/fewer photons. Instead, what we can say is that the number of modes is quadratic in the expected number of photons $\langle n \rangle$. Coupled with bounds on the variance of n, we can conclude that the output distribution is dominated by output states which obey the $m = \Theta(n^2)$ relationship.

⁵A notable exception is the proposal of Kruse et al. [11] which also circumvents traditional hiding arguments. Their approach is similar—program a GBS device with outputs proportional to hafnians of arbitrary symmetric matrices, which trivializes hiding symmetric matrices. While they conjecture a hardness argument in the regime where the number of photons is linear in the number of modes, our paper shows that the types of challenges that will be required to carry out the full complexity-theoretic argument.

 $^{^6\}mathrm{It}$ is only known that the entire distribution of singular values approaches the quarter-circle law. See Section 2.3 for more discussion.

We present numerical evidence (see Fig. 3) showing that these formulas accurately predict the click distributions of random BipartiteGBS experiments.

Finally, we make preliminary progress towards a GBS hardness result in the regime where there are fewer than quadratically many more modes than photons.⁷ In this regime, we can no longer guarantee that there is at most one photon per mode. These photon collisions imply that the output probabilities are no longer described by permanents of simple submatrices, but rather by submatrices which have some rows and columns repeated. To this end, we define a new "Permanent-of-Repeated-Gaussians" problem for which the goal is to approximate such permanents. We provide numerics suggesting that a classical simulation of GBS in the high-collision regime can be leveraged to solve the Permanent-of-Repeated-Gaussians problem based on a Stockmeyer counting argument similar to that in Theorem 1. In other words, we could show approximate average-case hardness for GBS in the $m = o(\langle n \rangle^2)$ regime if we made the following assumptions: the #P-hardness of the Permanent-of-Repeated-Gaussians problem; a plausible conjecture in random matrix theory. We caution that these assumptions remain relatively unexplored.

As a first step towards understanding the Permanent-of-Repeated-Gaussians, we show that there is some sense in which we can reduce arbitrary matrix permanents to permanents of matrices with repeated rows and columns:

Theorem 2. Given an oracle \mathcal{O} that can approximate $\operatorname{Per}(B)$ for any matrix $B \in \mathbb{C}^{n \times n}$ that has k row/column repetitions to additive error ϵ , there is an $\operatorname{\mathsf{FBPP}}^{\mathcal{O}}$ algorithm that can approximate $\operatorname{Per}(A)$ for arbitrary matrices $A \in \mathbb{C}^{(n-k) \times (n-k)}$ to additive accuracy $O(\epsilon/1.498^k)$.

The reduction in this theorem has the additional nice property that if the matrix A is Gaussian, then the oracle is only queried on matrices B that are also Gaussian. This makes it an ideal candidate for use in a hardness reduction because we can only require a classical simulator to sample from distributions that we can sample from quantumly. Unfortunately, while we are getting an exponential improvement in the accuracy to Per(A), we show that it is still insufficient to conclude the #P-hardness of Permanent-of-Repeated-Gaussians from the usual Permanent-of-Gaussians conjecture. That said, if we assume that there are only constantly-many collisions we can show such a reduction. Furthermore, in this regime we can work directly with the magnitude of the permanent, avoiding the need for an additional anti-concentration conjecture.

Theorem 3. Given an oracle \mathcal{O} that can approximate $|\operatorname{Per}(B)|^2$ for any matrix $B \in \mathbb{C}^{n \times n}$ that has k row/column repetitions to additive error ϵ , there is an $\operatorname{FBPP}^{\mathcal{O}}$ algorithm that can approximate $|\operatorname{Per}(A)|^2$ for arbitrary matrices $A \in \mathbb{C}^{(n-k) \times (n-k)}$ to additive accuracy which is exponential in k, but polynomial in n and ϵ .

This theorem follows almost immediately by combining Theorem 2 with the polynomial interpolation techniques used to prove the classical hardness of BosonSampling with constantly-many lost photons [36].

The rest of this paper is organized as follows: Section 2 provides a brief introduction to GBS as well the BipartiteGBS protocol for programming GBS devices according to arbitrary transition matrices. It then states important properties of BipartiteGBS with Gaussian transition matrices: two lemmas concerning photon-number statistics and normalization constants (proofs in Appendix A), and analytical formulas for the distribution of click statistics backed up by numerics (proofs in Appendix C). Section 3 contains the proof that classical simulation of BipartiteGBS in the no-collision regime is hard (Theorem 1). Section 4 deals with BipartiteGBS in the collision regime and contains proofs of theorems 2 and 3 (though proofs of some important lemmas in Appendix B). Numerical evidence for extending our arguments beyond the dilute limit are given in Appendix D.

⁷It is worth noting that for the task of exact sampling, it was already known that BosonSampling is hard in the $m=o(n^2)$ regime, and in fact, even if m=n. Specifically, Grier and Schaeffer give a BosonSampling sampling experiment in the m=n regime for which a particular output probability is #P-hard to approximate [35]. Combined with the Stockmeyer counting arguments of [1], this shows classical intractability of the exact sampling task predicated on the non-collapse of the polynomial hierarchy.

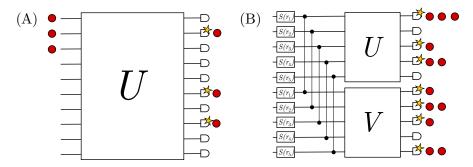


Figure 1: (A) Standard BosonSampling. The input to the interferometer are single photon states, which are detected at the output. (B) BipartiteGBS setup to encode arbitrary matrices. For a matrix C with singular value decomposition $C = U \operatorname{diag}(\sigma_i) V^T$, two-mode squeezed states with squeezing parameter $r_i = \tanh^{-1}(\sigma_i)$ are generated. These are sent to independent interferometers applying the unitaries U and V, respectively.

2 BipartiteGBS: Gaussian boson sampling with arbitrary transition matrices

2.1 Gaussian boson sampling introduction

Gaussian boson sampling is a model of photonic quantum computation where a multi-mode Gaussian state is prepared and then measured in the photon-number basis [7]. Gaussian states receive this name because their Wigner function—a quasi-probability representation of quantum states of light—is a Gaussian distribution [37]. We consider pure Gaussian states without displacements, which can be prepared from the vacuum by a sequence of single-mode squeezing gates followed by linear interferometry.

In contrast to BosonSampling, which uses single photons, GBS employs squeezed states as the input to the linear interferometer. In terms of the creation and annihilation operators a_i and a_i^{\dagger} on mode i, a squeezing gate is given by $S(r_i) = \exp[r_i(a_i^{\dagger 2} - a_i^2)/2]$, where r_i is a squeezing parameter. A squeezed state can be prepared by applying a squeezing gate to the vacuum. A linear interferometer transforms the operators as

$$a_i \to a_i' = \sum_j U_{ij} a_j,\tag{2}$$

where U is a unitary matrix.

Let $S = (s_1, s_2, \dots, s_{2m})$ encode a measurement outcome on 2m modes, where s_i is the number of photons in mode i. The probability of observing sample S when measuring a pure Gaussian state in the photon-number basis is given by [7]:

$$\Pr(S) = \frac{1}{\mathcal{Z}} \frac{\left| \operatorname{Haf}(A_S) \right|^2}{s_1! s_2! \cdots s_{2m}!}.$$
(3)

Here

$$\mathcal{Z} = \prod_{i=1}^{2m} \cosh(r_i),\tag{4}$$

$$A = U \operatorname{diag}(\sigma_1, \sigma_2, \dots, \sigma_{2m}) U^T,$$
(5)

$$\sigma_i = \tanh(r_i), \quad 0 \le \sigma_i < 1$$
 (6)

where r_i is the input squeezing parameter in the *i*th mode, U is the unitary describing the interferometer, and $\operatorname{Haf}(\cdot)$ is the hafnian.⁸ The matrix A_S is constructed from A by repeating the *i*th row

⁸For a $2n \times 2n$ symmetric matrix A, $\operatorname{Haf}(A) = \frac{1}{n!2^n} \sum_{\sigma \in \mathcal{S}_{2n}} \prod_{i=1}^n A_{\sigma(2i-1),\sigma(2i)}$. See references [38, 39, 40] for more detailed discussion of the hafnian and its complexity.

and column of A s_i -many times (e.g., if $s_i = 0$, both the corresponding row and column are removed entirely). Notice that Eq. (5) implies that A is symmetric.

The total mean photon number in the distribution is given by [19]

$$\sum_{i=1}^{2m} \frac{\sigma_i^2}{1 - \sigma_i^2}.\tag{7}$$

For future convenience we introduce n as the random variable describing the total number of photon pairs in a given sample. Its quantum-mechanical expectation is simply given by

$$\langle n \rangle = \frac{1}{2} \sum_{i=1}^{2m} \frac{\sigma_i^2}{1 - \sigma_i^2}.$$
 (8)

Notice that it is possible to choose a parameter $\lambda > 0$ and perform a rescaling $A_{\lambda} := \lambda A$ so that the singular values of A_{λ} become $\lambda \sigma_i$. One can check that the mean number of photon pairs $\langle n_{\lambda} \rangle$ is continuous for $\lambda \in [0, 1/\sigma_{\text{max}})$ and grows arbitrary large, and so it is possible to set λ such that $\langle n_{\lambda} \rangle$ is any desired non-negative number.

2.2 BipartiteGBS

We now introduce BipartiteGBS, a specific strategy for programming a GBS device such that the resulting distribution depends on an arbitrary transition matrix, not just a symmetric one. This scheme is illustrated in Fig. 1. First, we construct a device with 2m modes and generate photons by applying two-mode squeezing gates to modes i and i+m for $i=1,2,\ldots,m$. The two-mode squeezing gate is defined as $S_2(r_i) = \exp[r_i(a_i^{\dagger}a_{i+m}^{\dagger} - a_ia_{i+m})]$. It can be decomposed as two single-mode squeezing gates with identical parameters (i.e., $r_i = r_{i+m}$) followed by a 50:50 beamsplitter. The subsequent interferometer is configured by applying a unitary U to the first m modes and a separate unitary V to the second half of the modes. In this sense, this construction is a generalization of Scattershot BosonSampling [6] and Twofold Scattershot BosonSampling [41]. In the former the second interferometer is fixed to V = 1 and $r_i = r$ for all the two-mode squeezers, while in the latter only the squeezing parameters are fixed.

In this setup, the GBS distribution is also given by Eq. (3), but in this case the A matrix satisfies

$$A = \begin{pmatrix} 0 & C \\ C^T & 0 \end{pmatrix},\tag{9}$$

$$C = U \operatorname{diag}(\tanh r_i)V^T. \tag{10}$$

The expression for C is equivalent to the singular value decomposition $C = U\Sigma W^{\dagger}$ of an arbitrary complex matrix with singular values $\sigma_i = \tanh(r_i) \in [0,1)$, where we simply set $V^T = W^{\dagger}$. Thus, it is possible to choose C to be an arbitrary complex matrix, up to a rescaling $C \to \lambda C$ for some appropriate $\lambda > 0$ such that the singular values satisfy $\sigma_i \in [0,1)$.

The resulting output distribution can be more elegantly expressed directly in terms of the matrix C, which we refer to as the transition matrix. We introduce the notation $(S;T) = (s_1, \ldots, s_m; t_1, \ldots, t_m)$ to denote a sample. Here s_i is the number of photons in mode i and t_i is the number of photons in mode i + m, for $i = 1, 2, \ldots, m$. From Eq. (9), the outcomes S determine the rows of C and the columns of C^T that are kept or repeated when defining the submatrix A_S . Similarly, the outcomes T determine the columns of C and the rows of C^T . With this in mind, we employ the identity

$$\operatorname{Haf}\left[\begin{pmatrix} 0 & C \\ C^T & 0 \end{pmatrix}\right] = \operatorname{Per}(C),$$

to express the GBS distribution as:

$$\Pr(S;T) = \frac{1}{\mathcal{Z}} \frac{\left| \operatorname{Per}(C_{S,T}) \right|^2}{\prod_i s_i! \prod_i t_i!}.$$
(11)

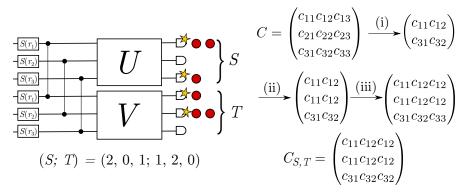


Figure 2: Constructing submatrices $C_{S,T}$ of the GBS distribution in Eq. (11). In this example, a device with 2m=6 modes is configured according to an m-dimensional matrix $C=U\operatorname{diag}(\tanh r_i)V^T$. The output is the photon pattern (S;T)=(2,0,1;1,2,0), where S determines the rows of the submatrix and T determines the columns. In step (i), since $s_2=0$ and $t_3=0$, we remove the second row and the third column. In step (ii), because $s_1=2$, we repeat the first row twice and finally in step (iii) since $t_2=2$, we repeat the second column twice to arrive at the submatrix $C_{S,T}$.

The notation $C_{S,T}$ corresponds to a submatrix obtained as follows: if $s_i = 0$, the *i*th row of C is removed. If $s_i > 0$, it is instead repeated s_i times. Similarly, if $t_i = 0$, the *i*th column of C is removed and if $t_i > 0$, it is repeated t_i times. See Fig. 2 for an example.

Since the permanent is only defined for square matrices, the number of photons detected in the first half of the modes should be equal to the total number of photons detected in the second set of modes, i.e., $\sum_{i=1}^m s_i = \sum_{i=1}^m t_i = n$. Physically, this corresponds to the action of the two-mode squeezing gate, which generates pairs of photons such that every photon in the first m modes has a twin photon in the remaining m modes. This observation, or the fact that $r_i = r_{i+m}$, allows us to write the expectation of the number of pairs as $\langle n \rangle = \sum_{i=1}^m \frac{\sigma_i^2}{1-\sigma_i^2}$, where the sum only extends to m.

Eq. (11) is the starting point for the computational problem we consider. In this formulation, our GBS construction is almost identical to a standard BosonSampling setup [1]. In both cases, probabilities are given in terms of the permanents of submatrices constructed in the same manner. The main difference is that, in our case, we employ 2m modes to encode an arbitrary $m \times m$ complex matrix, whereas the corresponding matrix in BosonSampling must be unitary, namely equal to the matrix that describes the interferometer. Another crucial difference is the normalization factor \mathcal{Z} . It is necessary to account for the fact that the space of outcomes includes events with different total photon numbers, and it will influence the behaviour of errors in our final result in Section 3.

2.3 BipartiteGBS with Gaussian matrices

As discussed above, it is possible to encode an arbitrary matrix in the GBS output distribution. In this section, we specialize this to the case of Gaussian random matrices. Let $\mathcal{N}(\mu, \Sigma^2)^{m \times m}_{\mathbb{C}}$ be the distribution over $m \times m$ matrices whose entries are independent complex Gaussians variables with mean μ and variance Σ^2 . We choose $C \sim \mathcal{N}(0, \Sigma^2)^{m \times m}_{\mathbb{C}}$, for Σ to be specified shortly.

By choosing C to be Gaussian, this mirrors the case of BosonSampling [1], where sufficiently small submatrices of uniformly Haar random unitaries are approximately also Gaussian. Therefore, this will allow us to support our hardness-of-simulation result on the same set of conjectures. However, in our construction, any $n \times n$ submatrix of C is also Gaussian, for any scaling between m and n. Contrast this with BosonSampling, where requiring submatrices to be approximately Gaussian formally constrains the number of modes to be much larger than the number of photons—rigorously, $m = \omega(n^5)$. We now prove a few important facts about GBS with Gaussian matrices.

First, several important quantities, such as the squeezing parameters, the normalization constant

 \mathcal{Z} , and the mean photon number $\langle n \rangle$, depend on the list of singular values of C, which we denote by $\{\sigma_i\}_{i=1...m}$. For m-dimensional random complex matrices with mean 0 and variance 1/m, in the asymptotic limit $m \to \infty$, the distribution $p(\sigma)$ of singular values converges to

$$p(\sigma) = \frac{1}{\pi} \sqrt{4 - \sigma^2}.\tag{12}$$

This is the quarter-circle law for random matrices [42]. Importantly, this result states that singular values are constrained within a finite interval, in this case $\sigma \in [0, 2]$, with high probability. The probability that the largest singular value is greater than $2 + \epsilon$ decays exponentially as $m \exp(-m\epsilon^2/8)$ [43]. See also Lemma 15 in Appendix A.2. For now, we simply assume that the singular values are in fact within this range, though we return to this issue in the next section. The above equation holds as the limit for the empirical distribution over singular values of C, and therefore it also corresponds to the limit marginal distribution satisfied by any single σ_i . However, we cannot assume in general that the σ_i are drawn independently from it.

Recall from Eq. (7) that for C to be amenable to encoding in the GBS device, its largest eigenvalue must lie in the range [0,1). Furthermore, it will be useful to tune the relation between m and $\langle n \rangle$. To address both of these issues, we rescale the matrix by a further factor of $1/\alpha$, for some $\alpha > 2$. Alternatively, we can choose C from $\mathcal{G} := \mathcal{N}(0,1/(\alpha^2 m))^{m \times m}_{\mathbb{C}}$. By doing this, the limiting distribution for the singular values in the interval $[0,2/\alpha]$ is

$$p_{\alpha}(\sigma) = \frac{\alpha}{\pi} \sqrt{4 - \alpha^2 \sigma^2}.$$
 (13)

We now consider the typical behaviour of two quantities that will be important for our main result. The first is the number of photon pairs n. The size of the matrix permanent associated with an output probability—and hence its complexity—is directly determined by the number of photons observed in a given experimental run. However, in GBS the total photon number is not fixed, and the fluctuations in photon number depend on the matrix C via its singular values. Therefore, it will be important to prove that fluctuations around the mean photon number are small enough so that our main argument is stable.

The other important quantity is the normalization constant \mathcal{Z} , which appears as a multiplicative factor between the matrix permanent and the output probabilities. For this reason, it will directly affect the error bounds of our main results, and we need to prove that it is typically not too large.

In the remainder of this section we give an intuitive analysis of these quantities, together with suitably formal bounds that are proven in Appendix A.

2.3.1 Fluctuations of the total photon number

From Eq. (13), we can compute the expected mean number of photon pairs as we vary 9 C over the Gaussian ensemble:

$$\mathbb{E}[\langle n \rangle] = \sum_{i=1}^{m} \mathbb{E}\left[\frac{\sigma_i^2}{1 - \sigma_i^2}\right] = \frac{m}{2} \left(\alpha^2 \left(1 - \sqrt{1 - 4/\alpha^2}\right) - 2\right). \tag{14}$$

Throughout our main argument, we usually consider a regime where m scales faster than $\Omega(n)$, i.e., as $\Theta(n^2)$. This corresponds to a regime where α is large and we can write $\alpha^2(1-\sqrt{1-4/\alpha^2})-2\approx \alpha^2(1-1+2/\alpha^2+2/\alpha^4)-2=2/\alpha^2$. Therefore in this regime we have

$$\mathbb{E}[\langle n \rangle] \approx m/\alpha^2. \tag{15}$$

⁹Because it can sometimes be confusing, let us reiterate here that we use $\langle \cdot \rangle$ to denote the "quantum-mechanical average", i.e. average over the photon number distribution for given transition matrix, and $\mathbb{E}[\cdot]$ to denote the expected value over the Gaussian ensemble. We use $\Delta^2[\cdot]$ and $\mathrm{Var}[\cdot]$ for the "quantum-mechanical variance" and variance due to the Gaussian ensemble, respectively.

For instance, if we want the GBS device to operate in a regime where $m = c\mathbb{E}[\langle n \rangle]^2$ for some c > 0, it suffices to choose $\alpha = (cm)^{1/4} = \Theta(m^{1/4})$. We will refer to this regime from now on as the dilute limit.

Even assuming the singular values follow the quarter-circle distribution exactly, computing the expectation of $\langle n \rangle$ is not enough. The complexity implied in Eq. (11) depends on the observed number of photons, not on $\langle n \rangle$. Therefore we must prove that, with high probability, n is not so far from its expectation as to invalidate our conclusions. We show the following formal bound, which follows from Theorem 19 and Theorem 21 in Appendix A:

Lemma 4. For any $\delta > 0$, we have

$$\Pr\left[\left|\langle n\rangle - \frac{m}{\alpha^2}\right| \ge \frac{512m}{\alpha^4} + \frac{1}{\alpha^2}\sqrt{\frac{2}{\delta}}\right] \le \delta$$

$$\Pr\left[\left|n - \frac{m}{\alpha^2}\right| \ge \frac{2\sqrt{m}}{\alpha\sqrt{\delta}} + \frac{3}{\alpha\delta^{3/4}} + \frac{84\sqrt{m}}{\alpha^2\sqrt{\delta}} + \frac{512m}{\alpha^4}\right] \le \delta$$

whenever $\alpha \geq 6$, and $m \geq \ln(1/\delta)$. The first probability is over the choice of Gaussian matrix C, whereas the second probability is over both the choice of C and over the photon number distribution.

Notice that in the dilute limit this statement implies that with high probability over the choices of Gaussian matrix and over the photon number distribution, the observed photon number is linear in m/α^2 to leading order.

2.3.2 The normalization factor

Let us now consider the typical behaviour of the normalization factor \mathcal{Z} . Recall that we can write

$$\mathcal{Z} = \prod_{i=1}^{m} \cosh(\operatorname{arctanh} \sigma_i)^2 = \prod_{i=1}^{m} \frac{1}{1 - \sigma_i^2}.$$
 (16)

Assuming $\sigma_i \in [0, 2/\alpha]$, it holds that $\mathcal{Z} \in [1, 1/(1-4/\alpha^2)^m]$. Asymptotically, the upper bound can be written as

$$\mathcal{Z}_{\text{max}} = 1/(1 - 4/\alpha^2)^m \approx e^{4m/\alpha^2}.$$
 (17)

We want \mathcal{Z} to be as small as possible. As will become clear in the next sections, the bound in Eq. (17) is not sufficiently tight for our purposes. On the other hand, if each singular value σ_i is drawn independently from the quarter-circle distribution (Eq. (13)), then the expectation of \mathcal{Z} would scale more favourably as e^{m/α^2} . Although the singular values are not independent, we prove that this heuristic argument in fact provides the right scaling (see Appendix A.2). More specifically, we give the following bound:

Lemma 5. For any $\delta > 0$

$$\Pr\left[\mathcal{Z} \ge \frac{2}{\delta} e^{m/\alpha^2} e^{272m/\alpha^4}\right] \le \delta,\tag{18}$$

whenever $\alpha \geq 6$, and $m \geq \ln(1/\delta)$.

Recall that $\alpha = \Theta(m^{1/4})$ in the dilute limit, so Lemma 5 implies that \mathcal{Z} is bounded asymptotically by e^{m/α^2} .

2.3.3 Collision statistics beyond the dilute limit

A BipartiteGBS sample is specified by M=2m non-negative integers giving the number of photons measured in each of the modes:

$$(S,T) = (s_1, \dots, s_m; t_1, \dots, t_m).$$
 (19)

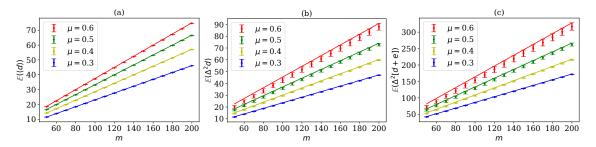


Figure 3: We compare the analytical results in Eqs. (24), (25) with numerics. For each value of m we generate 2500 random matrices where each of the m^2 elements is drawn i.i.d. from the standard complex-normal distribution and then each matrix is scaled so as to fix the total mean photon number to be $2\langle n\rangle = \mu \times 2m$. We then calculate (a) the average of the mean number of clicks in the first half of the modes $\mathbb{E}(\langle d\rangle)$, (b) its variance $\mathbb{E}\left[\Delta^2 d\right]$ and (c) the variance in the total number of clicks $\mathbb{E}\left[\Delta^2 (d+e)\right]$. The error bars in (a), (b) and (c) are given by $\sqrt{\mathrm{Var}(\langle d\rangle)}$, $\sqrt{\mathrm{Var}\left[\Delta^2 d\right]}$ and $\sqrt{\mathrm{Var}\left[\Delta^2 [d+e]\right]}$ respectively and are obtained by using the 2500 Monte Carlo samples for each value of μ and m. The lines are the predictions from theory. We do not show the mean $\mathbb{E}(\langle e\rangle)$ and variance $\mathbb{E}\left[\Delta^2 e\right]$ of the clicks in the second half of the modes since they are indistinguishable from the corresponding results of the first half. Note that for densities $\mu < 0.5$ the theory agrees very well with the numerics. For larger values of the density the first-order Taylor expansions used to derive Eqs. (25) in Appendix C.4 no longer hold, giving a significant deviation from the numerically obtained values as seen above for the variances in Figs. 3 (b) and 3 (c). The numerical calculation was performed using The Walrus [44].

As mentioned before, the total number of photons detected in the first m modes should be equal to the total number of detected in the second m modes—that is, $n = \sum_{i=1}^{m} s_i = \sum_{j=1}^{m} t_j$.

Another useful variable to consider is the number of clicks. A click sample is obtained from a photon number sample by "thresholding" the events, mapping any event with more than zero photons into outcome 1 while mapping vacuum events to 0. We write these thresholded samples as

$$(D, E) = (d_1, \dots, d_m; e_1, \dots, e_m)$$
 (20)

where $d_i, e_j = 1$ if $s_i, t_j \ge 1$ and 0 otherwise. It is also useful to write the total number of clicks in either half

$$d = \sum_{i=1}^{m} d_i \le n, \quad e = \sum_{j=1}^{m} e_j \le n, \tag{21}$$

where we also state the obvious fact that the total number of clicks is always smaller or equal to the total number of photons detected. Note that, unlike the photon number, the number of clicks in both halves of the modes need not be the same, thus in general, $d \neq e$. Whenever the number of clicks is less than the number of photons, there must be collisions (at least one mode with more than one photon).

To understand why the number of clicks is an important random variable consider the expression for the probabilities (recall Eq. (11)) depending on

$$Per(C_{S,T}). (22)$$

A priori, while the $n \times n$ matrix $C_{S,T}$ may be large, its rank $(\min(d, e))$ may be small. Because matrices of small rank have efficient algorithms [34], it is useful to understand the statistics of the clicks in each of the two halves of the modes. In the dilute limit no-collision events dominate and thus d = e = n.

Beyond the dilute limit, we can find simple expressions for the first and second moments of the total number of clicks in either half of the 2m modes. The detailed derivation of these results is provided in Appendix C. These expressions are written in terms of the photon number density

$$\mu = \mathbb{E}(\langle a_i^{\dagger} a_i \rangle) = \frac{\mathbb{E}(\langle n \rangle)}{m}.$$
 (23)

For the first order moments (means) we simply invoke the fact that the reduced states of two-mode squeezed states are thermal states, and scrambling from the random interferometers leads also to locally thermal states to find,

$$\mathbb{E}(\langle d \rangle) = \mathbb{E}(\langle e \rangle) = m \frac{\mu}{1 + \mu}.$$
 (24)

By rewriting the density μ in terms of the mean number of pairs and the number of modes one easily derives Eq. (1) for $\mathbb{E}(\langle d \rangle + \langle e \rangle)$. For the second order moments we need to invoke the quarter circle law and use a Taylor expansion to obtain

$$\mathbb{E}(\Delta^2 d) = \mathbb{E}(\Delta^2 e) = m \frac{\mu (1 - \mu^2 + \mu)}{(1 - \mu)(\mu + 1)^3} = (1 - \mu^2 + \mu) \mathbb{E}(\text{Cov}(d, e)), \tag{25a}$$

$$\mathbb{E}(\Delta^2(d+e)) = 2m \frac{(2-\mu)\mu}{(1-\mu)(1+\mu)^2}.$$
 (25b)

In the dilute limit $\mu \sim 1/\sqrt{m} \ll 1$ we have $\mathbb{E}(\Delta^2 d) = \mathbb{E}(\Delta^2 e) \approx (1+\mu)\mathbb{E}(\operatorname{Cov}(d,e))$ which tells us that d and e are very strongly correlated, as one would expect since in this limit clicks reduce to photon numbers, which should be equal in the two halves of the modes. Even beyond the dilute limit, the equations above predict strong correlations between the number of clicks in either half of the modes. We can for example consider the linear correlation ratio

$$Corr(e,d) = \frac{\mathbb{E}(Cov(d,e))}{\sqrt{\mathbb{E}(\Delta^2(d))\mathbb{E}(\Delta^2(e))}} = \frac{1}{1+\mu-\mu^2}$$
 (26)

which for a non-negligible density of $\mu = 0.3$ gives $Corr(e, d) \approx 0.83$. We find excellent agreement between the results in the equations above against exact numerical calculations for varying photon-number densities as the number of modes increased in Fig. 3.

We conclude this section with a discussion of the bosonic birthday paradox, a bound on how likely we are to observe collision events in a BosonSampling experiment, which will also turn out to govern BipartiteGBS [45]. Specifically, for a BosonSampling experiment with m modes and $n \sim c m^{(j-1)/j}$ photons, the bosonic birthday paradox says that the number of output modes where we expect to observe exactly j photons converges to a Poisson random variable with mean c^j . Collision outcomes with more than j photons are suppressed.

The key idea in the proof of the bosonic birthday paradox is that, upon applying a Haar-random matrix to an n-photon Fock state, the output state is the maximally mixed state over the n-photon, m-mode Hilbert space. While BipartiteGBS experiments do not in general have Fock input states (in fact, the total photon number is a random variable), one can show that the bosonic birthday paradox holds when we restrict to a fixed subspace of n photons since that restricted input state is a Fock state. Moreover, in the singular value decomposition of a Gaussian matrix, the two unitary matrices U and V corresponding to the interferometers in BipartiteGBS are Haar-random [46]. Therefore, it will still be the case that the output state, when averaged over the Gaussian ensemble, can be seen as a copy of the maximally mixed state on each set of modes, as per Fig. 1.

For example, because Lemma 4 shows that n is highly concentrated around m/α^2 in the dilute limit, we can apply the bosonic birthday paradox in this regime (at j=2). In particular, this implies that some constant fraction of the output distribution of a BipartiteGBS experiment has no collisions, a fact which is somewhat implicit¹⁰ in the proof of Theorem 9.

 $^{^{10}}$ Theorem 9 implies the following weaker statement: supposing the Permanent Anti-Concentration Conjecture, a polynomially large fraction of the BipartiteGBS output distribution has no collisions. *Proof.* Suppose that the no-collision subspace $\mathcal H$ does not comprise at least an inverse-polynomial fraction of the BipartiteGBS distribution—that is, $\sum_{S\in\mathcal H} p_S = o(1/\mathsf{poly}(m))$ with overwhelming probability over Gaussian matrices. Then, a randomly chosen probability p_S in the no-collision subspace must be $o(1/(|\mathcal H|\mathsf{poly}(m)))$ with high probability. Expanding out the expression for the

3 Approximate average-case hardness for GBS in the dilute limit

In this section we prove our main result, namely that GBS is hard to simulate efficiently on a classical computer in the dilute limit and up to the same complexity conjectures as in BosonSampling [1]. To that end, we start with the following computational problem:

Problem 6. ($|\text{GPE}|_{\pm}^2$ [1]) Given as input a Gaussian matrix $X \sim \mathcal{N}(0,1)_{\mathbb{C}}^{n \times n}$ together with error bounds $\epsilon, \delta > 0$, output a complex number \tilde{P} such that $\left|\tilde{P} - |\text{Per }X|^2\right| \leq \epsilon n!$ with probability at least $1 - \delta$.

Our goal is to prove that, if there exists an efficient classical algorithm to simulate the output of a GBS experiment to high precision in total variation distance, then $|\mathrm{GPE}|_{\pm}^2$ can be solved in the complexity class $\mathsf{FBPP}^\mathsf{NP}$. It is conjectured¹¹ that $|\mathrm{GPE}|_{\pm}^2$ is $\#\mathsf{P}$ -hard, and so we obtain that the polynomial hierarchy (PH) collapses to its third level by the usual chain of inclusions:

$$\mathsf{P}^{\mathsf{PH}} = \mathsf{P}^{\#\mathsf{P}} \subseteq \mathsf{BPP}^{|\mathsf{GPE}|_{\pm}^2} \subseteq \mathsf{BPP}^{\mathsf{FBPP}^{\mathsf{NP}}} \subseteq \Sigma_3^p. \tag{27}$$

Since the polynomial hierarchy is widely conjectured to be infinite, such an efficient classical algorithm is unlikely to exist [1, 47].

We will break the proof of our main result in two parts. In Section 3.1, we apply a Stockmeyer counting argument to leverage an efficient classical algorithm that samples from the GBS output distribution into a FBPP^{NP} algorithm that produces an estimate of any probability within this distribution. This follows the corresponding argument in [1], though we emphasize some aspects of the proof that are particular to our scenario. In Section 3.2, we then use Eq. (11) to connect the output probability of an event to the permanent of a Gaussian matrix, and discuss how this affects the corresponding error bounds.

3.1 From distributions to probabilities

Consider the probability distribution \mathcal{D}_C at the output of a GBS experiment which implements an arbitrary¹² transition matrix C, as described in Section 2.1. Let \mathcal{O} be a deterministic classical algorithm that takes as input a description of C, together with an error bound β and a uniformly-random number r, and outputs a sample drawn according to distribution \mathcal{D}'_C . We write this as

$$\mathcal{O}(C, 0^{1/\beta}, r) \sim \mathcal{D}'_C, \tag{28}$$

as $r \in \{0,1\}^{\mathsf{poly}(m)}$ is sampled from the uniform distribution. Suppose also that

$$\|\mathcal{D}_C - \mathcal{D}_C'\| \le \beta. \tag{29}$$

probability and rearranging, we get

$$\frac{|\mathcal{H}|}{\mathcal{Z}\alpha^{2n}m^n}|\operatorname{Per}(C_S)|^2 = o(1/\mathsf{poly}(m))$$

where $C \sim \mathcal{N}(0,1)_{\mathbb{C}}^{m \times m}$ and S is chosen uniformly at random from \mathcal{H} . We show in Theorem 9 (Eq. (44) and Eq. (45)) that $\mathcal{Z}\alpha^{2n}m^n/|\mathcal{H}|n!=\mathsf{poly}(m)$ with overwhelming probability. Plugging in this bound, we get $|\operatorname{Per}(C_S)|^2=o(n!/\mathsf{poly}(m))$. However, this contradicts the Gaussian Permanent Anti-Concentration Conjecture which says that $|\operatorname{Per}(C_S)|^2=\Omega(n!/\mathsf{poly}(m))$ with high probability.

 11 In [1], the conjecture is decomposed into two parts: the Permanent-of-Gaussians Conjecture asserts that the multiplicative version of the Gaussian permanent estimation problem is hard, while the Permanent Anti-Concentration Conjecture implies that this multiplicative version is poly-time equivalent to $|GPE|_{+}^{2}$.

¹²If the singular values of C are greater than 1, then we cannot implement the corresponding GBS experiment and \mathcal{D}_C is undefined. For such matrices, let us just assume that \mathcal{D}_C is the distribution that always outputs 0 photons in every mode.

Our goal is to use Stockmeyer's theorem [48] to show that a BPP^{NP} machine, with access to \mathcal{O} , can produce a good approximation to the probability of some particular outcome.

At this point, for generality, we do not yet impose the restrictions that we are in the dilute limit, nor that C is a Gaussian matrix. We only use the fact that the Gaussian ensemble is *invariant* under permutations of its rows and columns. At a high level, this allows us to randomly permute the rows and columns of C in order to hide the outcome S that we care about within \mathcal{D}_C . The classical algorithm \mathcal{O} is allowed to incur some total error β in its probabilities, but it is unlikely to concentrate too much of that error in the specific hidden outcome.

Now let S be a particular GBS outcome, 13 and let $\mathcal{H}_S \subseteq \mathbb{N}^{2m}$ be the subset of GBS outcomes that contains (without multiplicity) all outcomes $S_{\pi,\sigma} = (s_{\pi(1)}, \ldots, s_{\pi(m)}; t_{\sigma(1)}, \ldots, t_{\sigma(m)})$ for all permutations $\pi, \sigma \in \mathbb{S}_m$. This set is, by definition, invariant under permutations within the first m modes and within the last m modes, which correspond to permutations of rows or columns of C, respectively. To reiterate, to work in the dilute limit, S will be a 0/1-vector, but such a restriction is not yet needed. Indeed, we could choose S to be any possible GBS outcome, even those with many collisions. That said, in order to obtain an approximation to the underlying permanent question, the subspace \mathcal{H}_S must correspond to outcomes with a sufficiently large probability mass.

We then have the following result, which closely follows Theorem 3 of [1]:

Theorem 7. There exists an FBPP^{NP°} algorithm which, given an outcome S, produces an estimate of the corresponding probability $\mathcal{D}_C(S)$ to additive error $\varepsilon/|\mathcal{H}_S|$ with probability at least $1-\delta$ (over choices of C) in time $\mathsf{poly}(m, 1/\varepsilon, 1/\delta)$.

Proof. At a high level, we will construct an algorithm which counts all settings of the random bits that cause \mathcal{O} to output S. However, \mathcal{O} is allowed to err with some constant probability β . In order to prevent \mathcal{O} from concentrating more of that error on S, we will first permute the rows and columns of C uniformly at random. Equivalently, we can say that \mathcal{O} cannot tell which outcome in \mathcal{H}_S we desire to approximate.

Set $\beta := \varepsilon \, \delta/16$ and feed the input $\langle C, 0^{1/\beta}, r \rangle$ to \mathcal{O} . Recall that \mathcal{O} returns a sample from \mathcal{D}'_C such that $\|\mathcal{D}_C - \mathcal{D}'_C\| \le \beta$ if r is sampled uniformly. Furthermore, if we let $p_X := \mathcal{D}_C(X)$ and $q_X := \mathcal{D}'_C(X)$ for any outcome $X \in \mathbb{N}^{2m}$, we have

$$q_S = \Pr_{r \in \{0,1\}^{\text{poly}(m)}} \left[\mathcal{O}(C, 0^{1/\beta}, r) = S \right]. \tag{30}$$

We use Stockmeyer's approximate counting method [48] on this probability. If we can also show that q_S is close to p_S , this will imply that we can use Stockmeyer's method to estimate the desired quantity p_S as well.

Let $\Delta_X := |p_X - q_X|$. We get

$$\mathbb{E}_{X \in \mathcal{H}_S}[\Delta_X] \le \frac{1}{|\mathcal{H}_S|} \sum_{X \in \mathbb{N}^{2m}} \Delta_X = \frac{2}{|\mathcal{H}_S|} \|\mathcal{D}_C - \mathcal{D}_C'\| \le \frac{2\beta}{|\mathcal{H}_S|}.$$
 (31)

By Markov's inequality, we get

$$\Pr_{X \in \mathcal{H}_S} \left[\Delta_X > \frac{2\beta k}{|\mathcal{H}_S|} \right] < \frac{1}{k}. \tag{32}$$

Setting $k := 4/\delta$ we have $2\beta k = \varepsilon/2$, and so

$$\Pr_{X \in \mathcal{H}_S} \left[\Delta_X > \frac{\varepsilon}{2} \cdot \frac{1}{|\mathcal{H}_S|} \right] < \frac{\delta}{4}. \tag{33}$$

Our goal is to bound Δ_S , the error in probability for the specific outcome we are trying to compute. Recall that rows and columns of C were randomized before feeding it to \mathcal{O} . Since the distribution over

 $^{^{13}}$ For convenience, we are writing (S,T) as S here. That is, S is being used to represent both the row collisions and column collisions.

C is permutation invariant, we have that a random outcome of C has the same probability as a fixed outcome which has been permuted randomly, that is,

$$\Pr_{(\pi,\sigma),C} \left[\Delta_S > \frac{\varepsilon}{2} \cdot \frac{1}{|\mathcal{H}_S|} \right] < \frac{\delta}{4},\tag{34}$$

where π, σ are the permutations used to randomize the rows and columns of C^{14} .

Let us now formally describe the approximation we get for q_S using Stockmeyer counting [48]. For any $\theta > 0$, we can obtain an estimate \tilde{q}_S such that

$$\Pr\left[\left|\widetilde{q}_S - q_S\right| > \theta \, q_S\right] < \frac{\delta}{2} \tag{35}$$

with an $\mathsf{FBPP}^{\mathsf{NP}^{\mathcal{O}}}$ machine running in time polynomial in m, $1/\theta$, and $1/\delta$. Since q_X is a probability, observe that $\mathbb{E}_{X \in \mathcal{H}_S}[q_X] \leq 1/|\mathcal{H}_S|$ since $q_X \leq 1$. Once again we apply Markov's inequality to get

$$\Pr_{X \in \mathcal{H}_S} \left[q_X > \frac{k}{|\mathcal{H}_S|} \right] < \frac{1}{k}. \tag{36}$$

For similar reasons as before, it follows that

$$\Pr_{(\pi,\sigma),C} \left[q_S > \frac{k}{|\mathcal{H}_S|} \right] < \frac{1}{k}. \tag{37}$$

Finally, we set $\theta := \varepsilon \delta/8$ and $k := 4/\delta$. Combing all the above with the union bound, we get

$$\Pr\left[\left|\widetilde{q}_{S}-p_{S}\right| > \frac{\varepsilon}{|\mathcal{H}_{S}|}\right] \leq \Pr\left[\left|\widetilde{q}_{S}-q_{S}\right| > \frac{\varepsilon}{2} \cdot \frac{1}{|\mathcal{H}_{S}|}\right] + \Pr\left[\left|q_{S}-p_{S}\right| > \frac{\varepsilon}{2} \cdot \frac{1}{|\mathcal{H}_{S}|}\right]$$

$$\leq \Pr\left[q_{S} > \frac{4}{\delta|\mathcal{H}_{S}|}\right] + \Pr\left[\left|\widetilde{q}_{S}-q_{S}\right| > \theta q_{S}\right] + \Pr\left[\Delta_{S} > \frac{\varepsilon}{2} \cdot \frac{1}{|\mathcal{H}_{S}|}\right]$$

$$< \frac{\delta}{4} + \frac{\delta}{2} + \frac{\delta}{4}, \tag{38}$$

where the probability is over choices of C, (π, σ) and the randomness of the approximate counting procedure.

3.2 From probabilities to permanents

We just showed that, from the assumption that there exists a classical algorithm \mathcal{O} that simulates the output distribution of a GBS experiment to within total variation distance β , it follows that there exists an FBPP^{NP°} algorithm that produces an estimate to any outcome S to within error $\varepsilon/|\mathcal{H}_S|$. We now show how to leverage this result to obtain an estimate for the permanent of a Gaussian matrix, i.e., to solve $|\text{GPE}|_{\pm}^2$.

We begin by embedding the matrix we care about, X, within the GBS transition matrix C. Recall that $X \sim \mathcal{N}(0,1)^{n \times n}_{\mathbb{C}}$, and so there is an efficient algorithm which takes X as input and produces a matrix $C' \sim \mathcal{N}(0,1)^{m \times m}_{\mathbb{C}}$ which contains X as its submatrix, occurring in a random position.

Now recall that C' cannot be directly implemented in the GBS setup, since we require its singular values to be between [0,1). Furthermore, we wish to work in the dilute limit. Therefore, we rescale the matrix C' by a factor of $1/(\alpha\sqrt{m})$, as discussed in Section 2.3, resulting in the transition matrix $C \sim \mathcal{G}$.

Suppose now, without loss of generality, that $X/(\alpha\sqrt{m})$ appears as the top left submatrix of C. Consider now the 2n-photon no-collision outcome, S, which contains a single photon in each of the

 $^{^{14}}$ We note that this step takes place of the "Hiding Lemma" in [1].

modes $\{1, \ldots, n\}$ and of the modes $\{m+1, \ldots m+n\}$. By the discussion in Section 2.1, the probability of this outcome is

 $\Pr(S) = \frac{1}{\mathcal{Z}} \frac{1}{m^n \alpha^{2n}} \left| \operatorname{Per} X \right|^2.$ (39)

From this correspondence between the probability of outcome S and $|\operatorname{Per} X|^2$, together with Theorem 7, we obtain the following result:

Corollary 8. Let $X \sim \mathcal{N}(0,1)^{n \times n}_{\mathbb{C}}$. There exists an FBPP^{NP°} algorithm which estimates $|\text{Per }X|^2$ to additive error

$$\varepsilon \frac{\mathcal{Z}m^n \alpha^{2n}}{|\mathcal{H}_S|},\tag{40}$$

with probability at least $1 - \delta$ over choice of X in time $poly(m, 1/\epsilon, 1/\delta)$.

Proof. We begin by following the procedure described previously in order to embed X as a submatrix of the GBS transition matrix C. The entire procedure will fail if C has singular values greater than 1 since such a transition matrix does not correspond to a valid GBS experiment. However, by Lemma 15, the maximum singular value of C is greater than $2/\alpha + \epsilon$ with probability at most $me^{-m\alpha^4\epsilon/8}$. Let us now set ϵ to $1-2/\alpha$. If $me^{-m\alpha^4\epsilon/8}$ is less than, say $\delta/2$, we can assume the singular values are less than 1 by the union bound. However, one can check that the experiment might fail with probability greater than $\delta/2$ if $1/\delta = \Omega(\exp(m\alpha^4))$. In this case, however, our algorithm can run in time exponential in m since every polynomial in $\exp(m)$ is bounded by another polynomial in $1/\delta$ (and recall we allow that our algorithm runs in time $\exp(1/\delta)$). Since $|\operatorname{Per} X|^2$ can be computed exactly in time exponential in m by Ryser's formula, the theorem still holds even in this extreme regime of δ .

We now apply Theorem 7, which provides an estimate of $\Pr(S)$ to within additive error $\varepsilon/|\mathcal{H}_S|$. By Eq. (39), this immediately yields an additive estimate for $|\Pr(X)|^2$ to within the desired error. \square

Notice that we have yet to use the fact that we are working in the dilute limit—indeed, the previous theorems have been stated for a general scaling parameter α . We will now need to make this restriction. That is, we will show that any classical oracle \mathcal{O} which approximately samples from the output distribution of our GBS experiment in the dilute limit can be leveraged to compute Gaussian permanents.

The main outstanding step in this proof is to find a sufficiently large subspace \mathcal{H}_S in which to hide the outcome S, such that the error bound obtained in Corollary 8 matches the required error to solve $|\text{GPE}|_{\pm}^2$. For that purpose, we restrict ourselves to the no-collision subspace with n photon pairs, where n is the size of the matrix permanent we wish to estimate in $|\text{GPE}|_{\pm}^2$. More specifically, we set the scale parameter α such that the expected number of photon pairs is exactly n, and restrict our attention to states $S = (s_1, \ldots s_m; t_1, \ldots t_m)$ where each s_i or t_j is only equal to 0 or 1, and where $\sum s_i = \sum t_j = n$. To clarify, the size of the matrix is fixed and not a random number. Nevertheless, we call this parameter n suggestively since the total photon number of the experiment is unlikely to be too far from $\mathbb{E}[\langle n \rangle]$ by Lemma 4. Remarkably, however, our proof never explicitly invokes this fact.

Theorem 9. $|GPE|_{\pm}^2 \in FBPP^{NP^{\mathcal{O}}}$.

Proof. Recall that we are given a matrix $X \sim \mathcal{N}_{\mathbb{C}}^{n \times n}$, and we wish to construct an $\mathsf{FBPP}^{\mathsf{NP}^{\mathcal{O}}}$ algorithm which estimates $|\operatorname{Per}(X)|^2$ to additive error $\epsilon n!$ with probability at least $1 - \delta$ in time $\mathsf{poly}(n, 1/\epsilon, 1/\delta)$.

We employ Corollary 8, which gives an $\mathsf{FBPP}^\mathsf{NP^\circ}$ algorithm that estimates of $|\mathrm{Per}\,X|^2$ to within additive error

$$\varepsilon \frac{\mathcal{Z}m^n \alpha^{2n}}{|\mathcal{H}_S|} \tag{41}$$

with probability at least $1 - \Delta$ in time $\mathsf{poly}(m, 1/\varepsilon, 1/\Delta)$. We will set $\Delta = \delta/2$ and set ε such that Eq. (41) is bounded by $\epsilon n!$. To this end, let us define the ratio

$$\mathcal{I} := \epsilon/\varepsilon = \frac{\mathcal{Z}\alpha^{2n}m^n}{|\mathcal{H}_S|n!}.\tag{42}$$

Our goal now is to bound how large \mathcal{I} can be. If it is at most polynomially large, then our proof is concluded since we can set $\varepsilon = \epsilon/\mathsf{poly}(n, 1/\delta)$ and thus achieve the precision required by $|\mathsf{GPE}|_{\pm}^2$. First recall that we are working in the dilute limit, so let us set $m = cn^2$ and $\alpha^2 = \sqrt{cm}$ where c > 1 is some constant to be determined later. Furthermore, in each half of the output we have n photons in m modes without collisions, and so

$$|\mathcal{H}_S| = {m \choose n}^2 \ge \left(1 - \frac{1}{c}\right)^2 \left(\frac{m^n}{n!}\right)^2 := \gamma_c \left(\frac{m^n}{n!}\right)^2 \tag{43}$$

where we have used¹⁵ that $m = cn^2$ and implicitly defined a new constant $\gamma_c > 0$. From this, we have

$$\mathcal{I} = \frac{\mathcal{Z}\alpha^{2n}m^n}{\binom{m}{n}^2 n!} \le \frac{\mathcal{Z}\alpha^{2n}n!}{\gamma_c^2 m^n} \le \frac{2\mathcal{Z}\sqrt{2\pi n}}{\gamma_c^2} \left(\frac{\alpha^2 n}{me}\right)^n = \frac{2\mathcal{Z}\sqrt{2\pi n}}{\gamma_c^2} \frac{1}{e^n}$$
(44)

where the first inequality uses Eq. (43), and the second uses a Stirling approximation bound. By Lemma 5, we have that

$$\Pr\left[\mathcal{Z} \ge \frac{4e^{272/c}}{\delta}e^n\right] \le \frac{\delta}{2} \tag{45}$$

as long as $m \geq \ln(2/\delta)$. Once again, if $m < \ln(2/\delta)$, we get $1/\delta = \Omega(e^{cn^2})$ and so we could have computed the permanent of X explicitly using Ryser's formula. Therefore, let us assume Eq. (45) and combine it with Eq. (44) by the union bound to conclude that $\mathcal{I} = O(\sqrt{n}/\delta)$ with probability at least $1 - \delta$ over the Gaussian ensemble. This completes the proof.

4 Dealing with collision outcomes

In this section, we establish connections between permanents of matrices with repeated rows and columns (corresponding to collision outcomes) and permanents of matrices without repetitions. The end goal is to establish a reduction between the hardness of computing permanents of Gaussian random matrices with and without repetitions. These results apply generally and may also be useful for BosonSampling.

To start, notice that we need to define new variants of the Gaussian Permanent Estimation problem where the goal is to estimate permanents of matrices with repeated rows and columns. Much like how the error tolerance for GPE is based on the expectation of the permanent, the error tolerance for these new problems is based on the expectation of permanents with repetitions for which we need the following result (proof in Appendix B):

Lemma 10. Let $A \sim \mathcal{N}(0,1)^{c \times c}_{\mathbb{C}}$. Then

$$\mathbb{E}\left[\left|\operatorname{Per}(A_{S,T})\right|^{2}\right] = n! \prod_{i=1}^{c} s_{i}! \prod_{j=1}^{c} t_{j}! \tag{46}$$

for any repetition vectors $S, T \in \mathbb{N}^c$ such that $n = \sum_{i=1}^c s_i = \sum_{j=1}^c t_j$.

We can now define the following problems that generalize $|GPE|_{\pm}^2$ and GPE_{\pm} :

Definition 11. Given $A \sim \mathcal{N}(0,1)^{c \times c}_{\mathbb{C}}$, vectors $S, T \in \mathbb{N}^c$ with $n = \sum_{i=1}^c s_i = \sum_{j=1}^c t_j$, and accuracy parameters $\varepsilon, \delta > 0$, we define the following problems for which we are required to output a complex number \tilde{P} such that

$${m \choose n} = \frac{\prod_{i=0}^{n-1} (m-i)}{n!} \ge \frac{(m-n)^n}{n!} = \frac{(1 - \frac{1}{cn})^n m^n}{n!} \ge \left(1 - \frac{1}{c}\right) \frac{m^n}{n!}.$$

¹⁵Whenever $c, n \geq 1$, we have

1.
$$|RGPE|_{\pm}^2 : \left| \tilde{P} - |Per(A_{S,T})|^2 \right| \le \varepsilon n! \prod_{i,j=1}^c s_i! t_j!$$

2.
$$\operatorname{RGPE}_{\pm} : \left| \tilde{P} - \operatorname{Per}(A_{S,T}) \right| \le \varepsilon \sqrt{n! \prod_{i,j=1}^{c} s_i! t_j!},$$

with probability at least $1 - \delta$ over the randomness of A in time $poly(n, c, 1/\varepsilon, 1/\delta)$. 16

Given that we do not have a conclusive proof of hardness in the high-collision regime, some may wonder whether we have defined $|RGPE|_{\pm}^2$ correctly. Do we need to make the error scale with the expectation as is the case with $|GPE|_{\pm}^2$? In Appendix D, we give numerical evidence that this is the right choice—i.e., Stockmeyer counting on the BipartiteGBS distribution in the high-collision regime still solves the $|RGPE|_{\pm}^2$ problem. Moreover, if the error scaling were much smaller, then the estimate it produced would be insufficient.

In [1], the authors show that $|\text{GPE}|_{\pm}^2$ and GPE_{\pm} are, in fact, equivalent, up to the Permanent Anti-Concentration Conjecture (PACC), which states that the permanents of random Gaussian matrices are not too concentrated around 0.17 With some amount of work, it is possible to prove a similar equivalence between problems $|\text{RGPE}|_{\pm}^2$ and $|\text{RGPE}|_{\pm}^2$ assuming a version of the PACC which includes permanents of Gaussian matrices with repetitions. Unfortunately, even this generalization has its limits. It can be shown that, for a matrix composed of c copies of a single row of Gaussian elements, the permanent does not anti-concentrate. Therefore, even if the PACC is true, we expect its generalization to fail above a certain number of row/column repetitions. The hope is that the generalized PACC might nonetheless hold for those repetition patterns which occur naturally in the GBS distribution. We do not pursue this here as this equivalence is not required for the arguments that follow.

This rest of this section will show how solutions to RGPE can be used to solve GPE. In Section 4.1, we show how the permanent of a matrix with repetitions can be written as a polynomial that depends on the permanent of the matrix without repetitions. In Section 4.2, we use this result to prove an efficient reduction between $RGPE_{\pm}$ and GPE_{\pm} and an inefficient reduction between $|RGPE|_{\pm}^2$ and $|GPE|_{\pm}^2$. The latter result shows that BipartiteGBS is hard when there are O(1) collisions in the outcomes.

4.1 Permanents of matrices with repeated rows and columns

Given a matrix A and repetition patterns S, T, our goal for this section is to construct some matrix B such that $Per(B_{S,T})$ can be expressed as a polynomial whose constant coefficient is proportional to Per(A). Therefore, given an oracle to compute the permanents of matrices with this collision pattern, we can infer (through polynomial interpolation) the permanents of matrices that have no repeated rows or columns. The degree of the polynomial will depend on the total number of collisions $k := k_S + k_T$, i.e., the number of collisions in repetition patterns S and T, respectively. Note that here we are allowing $k_S \neq k_T$. If this is the case, the matrix B must be rectangular so that the matrix $B_{S,T}$ is square (for which the permanent is well-defined). Formally, we show the following:

Theorem 12. Let $A \in \mathbb{C}^{c \times c}$ and $S, T \in \mathbb{N}^c$ such that $s_i, t_j \geq 1$ be given. There is a $(c + k_T) \times (c + k_S)$ matrix B := B[z; x, y] obtained by adding k_S columns and k_T rows to the matrix A with entries that

¹⁶Note that we have chosen to define A as a $c \times c$ Gaussian matrix, rather than an $m \times m$ matrix. This choice was motivated by the fact that the complexity of the RGPE problem will depend on the number of total clicks in the repetition patterns, rather than the total number of modes. In fact, this idea will be important later in Theorem 12, where we require that there is at least one photon in every mode $(s_i, t_i \ge 1)$.

 $^{^{17}}$ Formally, the Permanent Anti-Concentration Conjectures states that there is some polynomial p such that for all c and $\delta>0, \, \Pr_{A\sim \mathcal{N}(0,1)_c^c\times^c}[|\operatorname{Per}(A)|^2 < c!/p(c,1/\delta)] < \delta.$

depend on complex variables z, $x = \{x_{i,j}^{(\ell)}\}$, and $y = \{y_{i,j}^{(\ell)}\}$ such that

$$Per(B_{S,T}) = \xi Per(A) + \gamma_1 z + \gamma_2 z^2 + \dots + \gamma_k z^k$$

$$(47)$$

$$\xi = \prod_{i=1}^{c} s_i! t_i! \left(\prod_{j=1}^{t_i-1} x_{i,j}^{(i)} \right) \left(\prod_{j=1}^{s_i-1} y_{i,j}^{(i)} \right).$$
 (48)

The coefficients $\gamma_1, \ldots, \gamma_k$ are functions of the entries of A, S, T, x and y. Furthermore, whenever $x_{i,j}^{(\ell)}$, $y_{i,j}^{(\ell)}$ and all entries of A are i.i.d. standard complex Gaussians, and |z| = 1, we have that $B \sim \mathcal{N}(0,1)_{\mathbb{C}}^{(c+k_T)\times(c+k_S)}$.

While the majority of this proof is given Appendix B, let us briefly describe the form of the construction, which results from two separate (but essentially identical) steps. We first consider the case where only the rows of A are repeated. That is, we construct a rectangular matrix B' such that when *only* the row repetition pattern S is applied (yielding a square matrix B'_{S}) the analogous version of Eq. (47) holds (i.e., the constant coefficient of the permanent polynomial is proportional to the permanent of A). In the second step, we simply treat B'_{S} as if it were the original matrix A and do the same construction except for the repeated columns specified by the repetition pattern T.

Let us now describe this first step, the construction of B'. For each ℓ such that $s_{\ell} > 1$, append to A the $c \times (s_{\ell} - 1)$ matrix $V^{(\ell)}$ with entries

$$V_{ij}^{(\ell)} = \begin{cases} y_{i,j}^{(\ell)} & \text{if } i = \ell \\ zy_{i,j}^{(\ell)} & \text{o.w.} \end{cases}$$
 (49)

For example, the 2×2 matrix $A = \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix}$ with repetition vector S = (3,2) has

$$V^{(1)} = \begin{pmatrix} y_{1,1}^{(1)} & y_{1,2}^{(1)} \\ z \, y_{2,1}^{(1)} & z \, y_{2,2}^{(1)} \end{pmatrix} \quad \text{and} \quad V^{(2)} = \begin{pmatrix} z \, y_{1,1}^{(2)} \\ y_{2,1}^{(2)} \end{pmatrix},$$

which leads to the matrix

$$B' = \begin{pmatrix} A & | & V^{(1)} & | & V^{(2)} \end{pmatrix} = \begin{pmatrix} a_{1,1} & a_{1,2} & y_{1,1}^{(1)} & y_{1,2}^{(1)} & z y_{1,1}^{(2)} \\ a_{2,1} & a_{1,2} & z y_{2,1}^{(1)} & z y_{2,2}^{(1)} & y_{2,1}^{(2)} \end{pmatrix}.$$

When we apply the row repetition pattern S, we get

$$B_S' = \begin{pmatrix} a_{1,1} & a_{1,2} & y_{1,1}^{(1)} & y_{1,2}^{(1)} & z y_{1,1}^{(2)} \\ a_{1,1} & a_{1,2} & y_{1,1}^{(1)} & y_{1,2}^{(1)} & z y_{1,1}^{(2)} \\ a_{1,1} & a_{1,2} & y_{1,1}^{(1)} & y_{1,2}^{(1)} & z y_{1,1}^{(2)} \\ a_{2,1} & a_{2,2} & z y_{2,1}^{(1)} & z y_{2,2}^{(1)} & y_{2,1}^{(2)} \\ a_{2,1} & a_{2,2} & z y_{2,1}^{(1)} & z y_{2,2}^{(1)} & y_{2,1}^{(2)} \end{pmatrix}.$$

To give some intuition for the proof of correctness, one can show that each monomial in the expansion of the permanent for B'_S cannot have two $a_{i,j}$ terms from the same row or column of A unless there is also z term. Therefore, the only monomials which do not have a factor of z are those that also appear in the expansion for the permanent of A (albeit with an extra factor of a product of y variables which is always the same). To complete the proof, it suffices to count the multiplicity of these monomials.

This result can now be used to show that the permanent of a matrix without repetitions can be estimated from estimates of larger matrices with repetitions. The first straightforward way of doing this is just setting z=0 and $x_{i,j}^i=y_{i,l}^i=1$ in our construction. This can be viewed as a "worst-case reduction" between these two problems, in the sense that we are assuming we have full control

over the choice of the larger matrix B. However, this is not sufficient to give a reduction between RGPE and GPE since the matrices there are required to have independent (up to the corresponding repetitions) Gaussian matrix elements, and the matrix B[0; x, y] is very far from Gaussian. The next section describes two (related) techniques for dealing with Eq. (47) when the matrices are random.

4.2 Polynomial interpolation techniques to estimate Gaussian permanents

The first polynomial interpolation technique to deal with Eq. (47) is to query the polynomial at the (k+1)th roots of unity. This leads to the error scaling described in Theorem 2 in the Introduction, which we rephrase below (proof in Appendix B):

Theorem 13. Given access to an oracle for RGPE_± with error $\varepsilon \sqrt{(c+k)! \prod s_i! t_j!}$, it is possible to use k+1 calls to the oracle to solve GPE_± with error $\varepsilon \sqrt{c!}$ where $\varepsilon := \varepsilon (\frac{1}{|\xi|} \sqrt{((c+k)!/c!) \prod s_i! t_j!})$ with ξ as in Eq. (48).

Notice that if we disregard the need to convert between the allowable $(\sqrt{(c+k)!\prod s_i!t_j!})$ -error fluctuations for RGPE $_\pm$ and the $(\sqrt{c!})$ -error fluctuations for GPE $_\pm$, then this theorem implies an error of ϵ for the repeated matrix leads to an error of $\epsilon/|\xi|$ for the unrepeated matrix. Moreover, we have that $\epsilon/|\xi| = O(\epsilon/1.498^k)$ with high probability (also proved in Appendix B), and so we actually obtain an exponential improvement in the error accuracy. Unfortunately, once the target error bounds for RGPE $_\pm$ and GPE $_\pm$ are incorporated we are in opposite situation— $\epsilon := \epsilon(\frac{1}{|\xi|}\sqrt{((c+k)!/c!)\prod s_i!t_j!})$ is exponentially large. Therefore, Theorem 13 cannot be used to reduce GPE $_\pm$ to RGPE $_\pm$ without an exponential error blowup.

That said, we believe that Theorem 13 provides weak evidence of a formal connection between RGPE_\pm and GPE_\pm (and recall the latter is believed to be $\#\mathrm{P-hard}$). Furthermore, it is worth noting that Theorem 13 leads to a somewhat surprising error scaling, as the error in the polynomial interpolation does not depend explicitly on the polynomial degree k. We hope the relatively benign error scaling of Theorem 13 might make it a useful building block in a hardness proof for GBS or BosonSampling in the regime where high-collision outcomes dominate, and we leave filling the above gaps for future research.

However, a crucial aspect of this reduction in Theorem 13 is that it is between amplitudes (whose phases are needed to arrange the cancellation of unwanted terms), whereas hardness arguments must be stated in terms of probabilities. We now describe an alternative polynomial interpolation technique which does provide an efficient reduction from $|GPE|_{\pm}^2$ to $|RGPE|_{\pm}^2$, but only in the regime where k = O(1). While the scaling in Theorem 13 will be much better and would also suffice in the constant-collision regime, we note that working directly with permanent magnitudes allows us to avoid creating a new anti-concentration conjecture for permanents with repeated rows and columns—i.e., to reduce from $RGPE_{\pm}$ to $|RGPE|_{\pm}^2$:

Theorem 14. Given access to an oracle for $|RGPE|_{\pm}^2$ with error $\varepsilon(c+k)! \prod s_i! t_j!$, it is possible to use O(k) calls to the oracle to solve $|GPE|_{\pm}^2$ with additive error $\epsilon c!$, for

$$\epsilon := O\left(\frac{(c+k)^{2k}k^{2k}}{\delta^{3k+1/2}2^k}\varepsilon\right) \tag{50}$$

Whenever k = O(1), we have $\epsilon = \mathsf{poly}(c, \varepsilon, \delta)$.

The main reason why the error above scales so poorly in k is that we take the absolute squared value of Eq. (47), and then apply a polynomial interpolation method. Contrary to Theorem 13, we cannot compute the polynomial for values of z in the unit circle, since many of its monomials depend on $|z|^2$ and so we cannot orchestrate the proper cancellations. As a result, we use a least squares estimator to approximate the value of the polynomial at z = 0 by sampling O(k) values of it in a small

range around z = 1. This causes a blowup that is exponential in the degree of the polynomial, leading to the scaling above.

In Appendix B we outline the proof of this result, though we omit some of the important technical lemmas that can be taken directly from [36] with regards to the hardness of lossy BosonSampling.

Theorem 14 can be plugged, in a relatively direct manner, into our main argument or that of [1] to prove that the complexity of both BosonSampling and GBS remains unchanged when we move from the no-collision subspace to a subspace with a (fixed) constant number of collisions. While this does not provide evidence of hardness in a regime where $m = O(n^d)$ for d < 2, we believe Theorem 14 is well-motivated in two senses: first, as a preliminary step towards a stronger result that does take into account more that constantly-many collisions, which might re-purpose some of the intermediate results; second, as a level of "experimental robustness", where allowing for a few collision outcomes to be included in the output distribution might improve the count-rate for finite-sized experiments (recall that even in the no-collision regime, some constant-fraction of the output distribution may still have collisions).

5 Discussion

We introduced BipartiteGBS as a method for programming a Gaussian boson sampling device so that the output probabilities are proportional to the permanents of arbitrary matrices. This allowed us to rigorously prove the hardness of approximately sampling from GBS distributions in the $m = \Theta(\langle n \rangle^2)$ regime under the same set of conjectures as those used for BosonSampling [1]. To recap the advantage of our approach, recall the two reasons why BosonSampling is required to operate in the $m = \Omega(n^2)$ (dilute) regime. The first is that the argument is predicated on the hardness of permanents with unrepeated rows and columns, and so we need that many more modes than photons for no-collision outcomes to dominate the output distribution. Perhaps more importantly, the second reason is that it is required that the $n \times n$ submatrices of $m \times m$ Haar-random matrices appear approximately Gaussian. In fact, while it is widely believed that such a statement holds whenever $m = \omega(n^2)$, it was only formally proved for $m = \omega(n^5)$. An important aspect of our work is that it has removed this obstacle in the case of GBS. Since we can implement arbitrary transition matrices, we can just choose them directly from the Gaussian ensemble, and thus their submatrices already look Gaussian even for $m = \Theta(n)$. Thus, there is the tantalizing prospect of improving on our work to show hardness for GBS in this regime, which would be much more experimentally friendly. We provide a blueprint for how this argument might go in Appendix D.

Let's also contrast our work with another recent paper of Deshpande et al. on the hardness of Gaussian boson sampling [26]. There, the authors also prove a worst-to-average case hardness result for approximate GBS in the dilute limit (c.f. Theorem 9). To do this, they must conjecture two plausible and yet new conjectures in complexity and random matrix theory. Furthermore, Deshpande et al. actually inherit the same problem that appears in BosonSampling—there are only rigorous proofs that the $n \times n$ submatrices of $m \times m$ are approximately Gaussian when $m = \omega(n^5)$. To circumvent this problem, the authors must conjecture directly that $m = \Omega(n^2)$ suffices (technically, they require that submatrices of the unitary product UU^T approximate random XX^T matrices for Gaussian X). Therefore, they will require fundamentally new ideas to operate in regimes beyond the dilute limit.

We reiterate that the main open problem left by our work is exactly that—prove hardness of GBS in a regime where the number of modes is subquadratic in the number of photons. We outline two possible approaches. First, one could attempt to improve the reduction in Section 4 so that an efficient algorithm for RGPE $_{\pm}$ implies an efficient algorithm for GPE $_{\pm}$. This would allow for the possibility that the hardness of GBS in the high-collisions could primarily be based on the same hardness conjectures as those used in BosonSampling. That said, several issues still need to be addressed. For example, consider our bound on $\mathcal Z$ in Lemma 5. There, we have a multiplicative $\exp(\Omega(m/\alpha^4))$ term which is constant in the dilute limit, but grows exponentially for smaller α . Some change in scaling is inevitable, but this bound will not lead to a tight enough estimation of the permanent to solve RGPE $_{\pm}$ via our

Stockmeyer counting argument. While we give strong evidence in Appendix D that such a bound exists, proving it rigorously may be challenging. Moreover, in order to use a reduction akin to that given in Theorem 13, one would need to prove an equivalence between the multiplicative and additive versions of RGPE. This would required a generalization of the Permanent Anti-Concentration Conjecture for matrices with repeated rows and columns which is not known to be equivalent to the original conjecture, and in fact is provably false in extreme settings. A second approach would be to try to strengthen the evidence for the #P-hardness of the Permanent-of-Repeated Gaussians problem without reference to the BosonSampling conjectures. We note that even in this case much of the same work outlined above would be required.

Another major open direction left in our work—and indeed in many of the proposed hardness arguments based on linear optics—is incorporating realistic experimental noise. To be clear, we show the classical hardness of approximately sampling (in total variation distance) from ideal GBS distributions. However, even in the high-collision regime, current experiments *also* do not sample approximately from their true distributions due to a variety of sources of noise (e.g., photon loss). Therefore, incorporating this noise is a critical to closing the gap between theory and experiment.

That said, it is likely that the flexibility of BipartiteGBS can help mitigate the effects of noise, particularly in near-term experiments. For instance, assume one wants to perform Scattershot Boson-Sampling with an unitary transition matrix W, which corresponds to BipartiteGBS where we set all squeezing parameters to be equal, U = W, and $V = \mathbb{I}$ (cf. Fig. 1). If W is Haar random, then it typically requires depth equal to m [29, 30]. But now note that we can shift some of the beam splitters from U to V, in the sense that any combination of these two matrices for which $W = UV^T$ leads to the same problem instance. Thus, we can program an m-depth decomposition of W, such as that of [29, 30], in an m/2-depth BipartiteGBS instance simply by mapping half of the beam splitter layers to U and half to V. Since losses scale exponentially with depth, this reduction of a factor of two for the depth corresponds to having square root of the losses.

Another less straightforward example is as follows. Suppose we measure the complexity of simulating the device by the computational cost of producing a single 2n-photon sample (from the exact distribution) using state-of-the-art algorithms. A 2m-mode implementation of the standard GBS model would typically require an arbitrary 2m-mode interferometer, and samples of its output probabilities can be generated in time $O(mn^32^n)$ [40, 49, 50]. A BipartiteGBS instance with comparable computational cost (up to polynomial factors) would also have 2m total modes, and its 2n-photon probabilities would be given by permanents of $n \times n$ matrices, with cost-per-sample of $O(n2^n)$ time [51, 52, 53, 54]. But an arbitrary BipartiteGBS interferometer only requires depth m, and thus half the depth (and square root of the losses) as an arbitrary implementation of standard GBS. Combining these two observations, it is possible to reduce the depth by a factor of 4 if one moves from arbitrary GBS to an implementation of BipartiteGBS based on an unitary transition matrix.

Finally, we ask what other uses BipartiteGBS might have in proving stronger forms of classical intractability—either from an experimental or theoretical viewpoint. For instance, because we are no longer constrained to have unitary transition matrices, one might envision a hardness result predicated on an entirely different distribution of matrices (e.g. Bernoulli instead of Gaussian). The flexibility of BipartiteGBS allows the possibility of more contrived distributions that nevertheless have a stronger complexity-theoretic foundation, and we hope that BipartiteGBS motivates others to explore these possibilities.

Acknowledgments

We would like to thank Torben Krüger for useful discussions. N.Q. thanks Z. Vernon for valuable discussions and the Ministère de l'Économie et de l'Innovation du Québec and the Natural Sciences and Engineering Research Council of Canada for financial support. D. J. Brod acknowledges support from Instituto Nacional de Ciência e Tecnologia de Informação Quântica (INCT-IQ, CNPq) and FAPERJ. M. B. A. Alonso acknowledges support from CNPq.

References

- [1] Scott Aaronson and Alex Arkhipov. "The computational complexity of linear optics". Theory of Computing 9, 143–252 (2013).
- [2] Max Tillmann, Borivoje Dakić, René Heilmann, Stefan Nolte, Alexander Szameit, and Philip Walther. "Experimental boson sampling". Nature Photonics 7, 540–544 (2013).
- [3] Justin B. Spring, Benjamin J. Metcalf, Peter C. Humphreys, W. Steven Kolthammer, Xian-Min Jin, Marco Barbieri, Animesh Datta, Nicholas Thomas-Peter, Nathan K. Langford, Dmytro Kundys, James C. Gates, Brian J. Smith, Peter G. R. Smith, and Ian A. Walmsley. "Boson sampling on a photonic chip". Science 339, 798–801 (2013).
- [4] Andrea Crespi, Roberto Osellame, Roberta Ramponi, Daniel J Brod, Ernesto F Galvao, Nicolo Spagnolo, Chiara Vitelli, Enrico Maiorino, Paolo Mataloni, and Fabio Sciarrino. "Integrated multimode interferometers with arbitrary designs for photonic boson sampling". Nature photonics 7, 545–549 (2013).
- [5] Matthew A. Broome, Alessandro Fedrizzi, Saleh Rahimi-Keshari, Justin Dove, Scott Aaronson, Timothy C. Ralph, and Andrew G. White. "Photonic boson sampling in a tunable circuit". Science 339, 794–798 (2013).
- [6] Austin P Lund, Anthony Laing, Saleh Rahimi-Keshari, Terry Rudolph, Jeremy L O'Brien, and Timothy C Ralph. "Boson sampling from a Gaussian state". Phys. Rev. Lett. 113, 100502 (2014).
- [7] Craig S. Hamilton, Regina Kruse, Linda Sansoni, Sonja Barkhofen, Christine Silberhorn, and Igor Jex. "Gaussian boson sampling". Phys. Rev. Lett. 119, 170501 (2017).
- [8] Marco Bentivegna, Nicolò Spagnolo, Chiara Vitelli, Fulvio Flamini, Niko Viggianiello, Ludovico Latmiral, Paolo Mataloni, Daniel J Brod, Ernesto F Galvão, Andrea Crespi, Roberta Ramponi, Roberto Osellame, and Fabio Sciarrino. "Experimental scattershot boson sampling". Science Advances 1, e1400255 (2015).
- [9] Hui Wang, Yu He, Yu-Huai Li, Zu-En Su, Bo Li, He-Liang Huang, Xing Ding, Ming-Cheng Chen, Chang Liu, Jian Qin, Jin-Peng Li, Yu-Ming He, Christian Schneider, Martin Kamp, Cheng-Zhi Peng, Sven Höfling, Chao-Yang Lu, and Jian-Wei Pan. "High-efficiency multiphoton boson sampling". Nature Photonics 11, 361 (2017).
- [10] Han-Sen Zhong, Li-Chao Peng, Yuan Li, Yi Hu, Wei Li, Jian Qin, Dian Wu, Weijun Zhang, Hao Li, Lu Zhang, Zhen Wang, Lixing You, Xiao Jiang, Li Li, Nai-Le Liu, Jonathan P. Dowling, Chao-Yang Lu, and Jian-Wei Pan. "Experimental Gaussian boson sampling". Science Bulletin 64, 511–515 (2019).
- [11] Regina Kruse, Craig S. Hamilton, Linda Sansoni, Sonja Barkhofen, Christine Silberhorn, and Igor Jex. "Detailed study of Gaussian boson sampling". Phys. Rev. A 100, 032326 (2019).
- [12] Thomas R Bromley, Juan Miguel Arrazola, Soran Jahangiri, Josh Izaac, Nicolás Quesada, Alain Delgado Gran, Maria Schuld, Jeremy Swinarton, Zeid Zabaneh, and Nathan Killoran. "Applications of near-term photonic quantum computers: software and algorithms". Quantum Science and Technology 5, 034010 (2020).
- [13] J. M. Arrazola, V. Bergholm, K. Brádler, T. R. Bromley, M. J. Collins, I. Dhand, A. Fumagalli, T. Gerrits, A. Goussev, L. G. Helt, J. Hundal, T. Isacsson, R. B. Israel, J. Izaac, S. Jahangiri, R. Janik, N. Killoran, S. P. Kumar, J. Lavoie, A. E. Lita, D. H. Mahler, M. Menotti, B. Morrison, S. W. Nam, L. Neuhaus, H. Y. Qi, N. Quesada, A. Repingon, K. K. Sabapathy, M. Schuld, D. Su, J. Swinarton, A. Száva, K. Tan, P. Tan, V. D. Vaidya, Z. Vernon, Z. Zabaneh, and Y. Zhang. "Quantum circuits with many photons on a programmable nanophotonic chip". Nature 591, 54–60 (2021).
- [14] Jianwei Wang, Fabio Sciarrino, Anthony Laing, and Mark G. Thompson. "Integrated photonic quantum technologies". Nature Photonics 14, 273–284 (2020).

- [15] Z. Vernon, N. Quesada, M. Liscidini, B. Morrison, M. Menotti, K. Tan, and J.E. Sipe. "Scal-able squeezed-light source for continuous-variable quantum sampling". Phys. Rev. Applied 12, 064024 (2019).
- [16] Joonsuk Huh, Gian Giacomo Guerreschi, Borja Peropadre, Jarrod R. McClean, and Alán Aspuru-Guzik. "Boson sampling for molecular vibronic spectra". Nature Photonics 9, 615–620 (2015).
- [17] Juan Miguel Arrazola and Thomas R. Bromley. "Using Gaussian boson sampling to find dense subgraphs". Phys. Rev. Lett. 121, 030503 (2018).
- [18] Leonardo Banchi, Mark Fingerhuth, Tomas Babej, Christopher Ing, and Juan Miguel Arrazola. "Molecular docking with Gaussian boson sampling". Science Advances 6, eaax1950 (2020).
- [19] Soran Jahangiri, Juan Miguel Arrazola, Nicolás Quesada, and Nathan Killoran. "Point processes with Gaussian boson sampling". Phys. Rev. E 101, 022134 (2020).
- [20] Maria Schuld, Kamil Brádler, Robert Israel, Daiqin Su, and Brajesh Gupt. "Measuring the similarity of graphs with a Gaussian boson sampler". Phys. Rev. A 101, 032314 (2020).
- [21] Soran Jahangiri, Juan Miguel Arrazola, Nicolás Quesada, and Alain Delgado. "Quantum algorithm for simulating molecular vibrational excitations". Physical Chemistry Chemical Physics **22**, 25528–25537 (2020).
- [22] Leonardo Banchi, Nicolás Quesada, and Juan Miguel Arrazola. "Training Gaussian boson sampling distributions". Phys. Rev. A 102, 012417 (2020).
- [23] Lars S. Madsen, Fabian Laudenbach, Mohsen Falamarzi. Askarani, Fabien Rortais, Trevor Vincent, Jacob F. F. Bulmer, Filippo M. Miatto, Leonhard Neuhaus, Lukas G. Helt, Matthew J. Collins, Adriana E. Lita, Thomas Gerrits, Sae Woo Nam, Varun D. Vaidya, Matteo Menotti, Ish Dhand, Zachary Vernon, Nicolás Quesada, and Jonathan Lavoie. "Quantum computational advantage with a programmable photonic processor". Nature 606, 75–81 (2022).
- [24] Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian Qin, Dian Wu, Xing Ding, Yi Hu, Peng Hu, Xiao-Yan Yang, Wei-Jun Zhang, Hao Li, Yuxuan Li, Xiao Jiang, Lin Gan, Guangwen Yang, Lixing You, Zhen Wang, Li Li, Nai-Le Liu, Chao-Yang Lu, and Jian-Wei Pan. "Quantum computational advantage using photons". Science 370, 1460–1463 (2020).
- [25] Han-Sen Zhong, Yu-Hao Deng, Jian Qin, Hui Wang, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Dian Wu, Si-Qiu Gong, Hao Su, et al. "Phase-programmable Gaussian boson sampling using stimulated squeezed light". Phys. Rev. Lett. 127, 180502 (2021).
- [26] Abhinav Deshpande, Arthur Mehta, Trevor Vincent, Nicolás Quesada, Marcel Hinsche, Marios Ioannou, Lars Madsen, Jonathan Lavoie, Haoyu Qi, Jens Eisert, Dominik Hangleiter, Bill Fefferman, and Ish Dhand. "Quantum computational advantage via high-dimensional Gaussian boson sampling". Science Advances 8, eabi7894 (2022).
- [27] Raúl García-Patrón, Jelmer J Renema, and Valery Shchesnovich. "Simulating boson sampling in lossy architectures". Quantum 3, 169 (2019).
- [28] Haoyu Qi, Daniel J. Brod, Nicolás Quesada, and Raúl García-Patrón. "Regimes of classical simulability for noisy Gaussian boson sampling". Phys. Rev. Lett. 124, 100502 (2020).
- [29] Michael Reck, Anton Zeilinger, Herbert J. Bernstein, and Philip Bertani. "Experimental realization of any discrete unitary operator". Phys. Rev. Lett. **73**, 58–61 (1994).
- [30] William R Clements, Peter C Humphreys, Benjamin J Metcalf, W Steven Kolthammer, and Ian A Walsmley. "Optimal design for universal multiport interferometers". Optica 3, 1460–1465 (2016).
- [31] Hubert de Guise, Olivia Di Matteo, and Luis L. Sánchez-Soto. "Simple factorization of unitary transformations". Phys. Rev. A 97, 022328 (2018).
- [32] Bryn A Bell and Ian A Walmsley. "Further compactifying linear optical unitaries". APL Photonics 6, 070804 (2021).

- [33] Tiefeng Jiang. "How many entries of a typical orthogonal matrix can be approximated by independent normals?". The Annals of Probability 34, 1497–1529 (2006).
- [34] Alexander I Barvinok. "Two algorithmic results for the traveling salesman problem". Mathematics of Operations Research 21, 65–84 (1996).
- [35] Daniel Grier and Luke Schaeffer. "New hardness results for the permanent using linear optics". In 33rd Computational Complexity Conference (CCC 2018). Volume 102 of Leibniz International Proceedings in Informatics (LIPIcs), pages 19:1–19:29. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2018).
- [36] Scott Aaronson and Daniel J. Brod. "BosonSampling with lost photons". Phys. Rev. A 93, 012335 (2016).
- [37] Christian Weedbrook, Stefano Pirandola, Raúl García-Patrón, Nicolas J. Cerf, Timothy C. Ralph, Jeffrey H. Shapiro, and Seth Lloyd. "Gaussian quantum information". Rev. Mod. Phys. 84, 621– 669 (2012).
- [38] Eduardo R Caianiello. "On quantum field theory—I: explicit solution of Dyson's equation in electrodynamics without use of Feynman graphs". Il Nuovo Cimento (1943-1954) **10**, 1634–1652 (1953).
- [39] Alexander Barvinok. "Combinatorics and complexity of partition functions". Volume 276. Springer. (2016).
- [40] Andreas Björklund, Brajesh Gupt, and Nicolás Quesada. "A faster hafnian formula for complex matrices and its benchmarking on a supercomputer". Journal of Experimental Algorithmics (JEA) 24, 11 (2019).
- [41] L. Chakhmakhchyan and N. J. Cerf. "Boson sampling with Gaussian measurements". Phys. Rev. A 96, 032326 (2017).
- [42] Jianhong Shen. "On the singular values of Gaussian random matrices". Linear Algebra and its Applications 326, 1–14 (2001).
- [43] Uffe Haagerup and Steen Thorbjørnsen. "Random matrices with complex Gaussian entries". Expositiones Mathematicae 21, 293–337 (2003).
- [44] Brajesh Gupt, Josh Izaac, and Nicolás Quesada. "The Walrus: a library for the calculation of hafnians, Hermite polynomials and Gaussian boson sampling". Journal of Open Source Software 4, 1705 (2019).
- [45] Alex Arkhipov and Greg Kuperberg. "The bosonic birthday paradox". Geometry & Topology Monographs 18, 1–7 (2012).
- [46] Antonia M Tulino and Sergio Verdú. "Random matrix theory and wireless communications". Now Publishers Inc. (2004).
- [47] Michael J. Bremner, Richard Jozsa, and Dan J. Shepherd. "Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy". Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences (2010).
- [48] Larry Stockmeyer. "The complexity of approximate counting". In Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing. Page 118–126. STOC '83. Association for Computing Machinery (1983).
- [49] Nicolás Quesada, Rachel S. Chadwick, Bryn A. Bell, Juan Miguel Arrazola, Trevor Vincent, Haoyu Qi, and Raúl García-Patrón. "Quadratic speed-up for simulating Gaussian boson sampling". PRX Quantum 3, 010306 (2022).
- [50] Jacob FF Bulmer, Bryn A Bell, Rachel S Chadwick, Alex E Jones, Diana Moise, Alessandro Rigazzi, Jan Thorbecke, Utz-Uwe Haus, Thomas Van Vaerenbergh, Raj B Patel, et al. "The boundary for quantum advantage in Gaussian boson sampling". Science Advances 8, eabl9236 (2022).

- [51] Herbert John Ryser. "Combinatorial mathematics". Volume 14. American Mathematical Soc. (1963).
- [52] Alex Neville, Chris Sparrow, Raphaël Clifford, Eric Johnston, Patrick M Birchall, Ashley Montanaro, and Anthony Laing. "Classical boson sampling algorithms with superior performance to near-term experiments". Nature Physics 13, 1153–1157 (2017).
- [53] Peter Clifford and Raphaël Clifford. "The classical complexity of boson sampling". Pages 146–155. Society for Industrial and Applied Mathematics. (2018).
- [54] Peter Clifford and Raphaël Clifford. "Faster classical boson sampling" (2020). arXiv:2005.04214.
- [55] Philip J Hanlon, Richard P Stanley, and John R Stembridge. "Some combinatorial aspects of the spectra of normally distributed random matrices". Contemporary Math 138, 151–174 (1992).
- [56] D Maiwald and D Kraus. "Calculation of moments of complex Wishart and complex inverse Wishart distributed matrices". IEE Proceedings - Radar, Sonar and Navigation 147, 162– 168 (2000).
- [57] S.M. Barnett and P.M. Radmore. "Methods in theoretical quantum optics". Clarendon Press. (2002).
- [58] Nathaniel R Goodman. "Statistical analysis based on a certain multivariate complex Gaussian distribution (an introduction)". The Annals of Mathematical Statistics **34**, 152–177 (1963).
- [59] Irina Shevtsova. "On the absolute constants in the Berry-Esseen-type inequalities". Doklady Mathematics 89, 378–381 (2014).
- [60] Alessio Serafini. "Quantum continuous variables: A primer of theoretical methods". CRC Press. (2017).
- [61] Nicolás Quesada, Juan Miguel Arrazola, and Nathan Killoran. "Gaussian boson sampling using threshold detectors". Phys. Rev. A 98, 062322 (2018).
- [62] Nicolás Quesada and Juan Miguel Arrazola. "Exact simulation of Gaussian boson sampling in polynomial space and exponential time". Phys. Rev. Research 2, 023005 (2020).
- [63] Peter D. Drummond, Bogdan Opanchuk, A. Dellios, and M. D. Reid. "Simulating complex networks in phase space: Gaussian boson sampling". Phys. Rev. A 105, 012427 (2022).
- [64] Alan Edelman. "Eigenvalues and condition numbers of random matrices". SIAM journal on matrix analysis and applications 9, 543–560 (1988).

A Bounds on n and $\mathcal Z$ from random matrix theory

In this appendix, we prove Lemma 4 and Lemma 5, corresponding to bounds on n and \mathcal{Z} , respectively. As in the main text, let $\mathcal{G} = \mathcal{N}(0, \frac{1}{\alpha^2 m})_{\mathbb{C}}^{m \times m}$ be the distribution over $m \times m$ matrices whose entries are independent complex Gaussians with mean 0 and variance $1/\alpha^2 m$.

The quantities we will bound depend crucially on the set $\{\lambda_i\}_{i=1,\dots,m}$, where $\lambda_i = \sigma_i^2$ and σ_i is the *i*th singular value of Gaussian matrix $C \sim \mathcal{G}$. Alternatively, λ_i is the *i*th eigenvalue of the matrix $A := CC^{\dagger}$. The matrix A is called a *complex Wishart matrix*. The complex Wishart distribution is usually denoted by $\mathcal{W}(\mu, \Sigma)$ when the columns of C are sampled from the multivariate complex normal distribution with mean vector μ and covariance matrix Σ . In our case we can write $A \sim \mathcal{W}(0, \frac{1}{\alpha^2 m} \mathbb{I}_m)$, and the latter ensemble will be represented by \mathcal{W} , for short.

Much is known about the spectrum of complex Wishart matrices from random matrix theory, and we will invoke several known results regarding W throughout this appendix. For example, it is known that the maximum eigenvalue of $A \sim W$ (denoted by $\lambda_{\max}(A)$) can be bounded as follows:

Lemma 15 (Haagerup and Thorbjørnsen [43]).

$$\Pr[\lambda_{\max}(A) \ge 4/\alpha^2 + \epsilon] \le me^{-m\alpha^4 \epsilon^2/8}.$$
 (51)

Let us recall the notation we will use for the source of randomness in these calculations. We write $\langle \cdot \rangle$ for the average of a quantity (e.g., the total photon number) over the distribution of photon numbers that arises in a GBS setup of Section 2.1, for a particular transition matrix C. We write $\mathbb{E}[\cdot]$ as the expectation of a particular quantity over the randomness of $C \sim \mathcal{G}$.

A.1 Bounds on observed photon number

We wish to upper bound the probability that the observed number of photon pairs, n, is far from its expectation. Let us begin by reviewing the computation of the expectation value itself. We can write

$$\langle n \rangle = \sum_{i=1}^{m} \frac{\lambda_i}{1 - \lambda_i} = \sum_{i=1}^{m} \sum_{k=1}^{\infty} \lambda_i^k = \sum_{k=1}^{\infty} \text{Tr}(A^k), \tag{52}$$

where λ_i is the *i*th eigenvalue of matrix $A \sim \mathcal{W}$. We can compute the exact expected value of $\langle n \rangle$ over random A using the following lemma:

Lemma 16 (Hanlon, Stanley, Stembridge [55]). For $A \sim \mathcal{W}(0, \sigma^2 \mathbb{I}_m)$,

$$\mathbb{E}[\text{Tr}(A^k)] = \frac{\sigma^{2k}}{k} \sum_{i=1}^k (-1)^{i-1} \frac{(m+1-i)_k^2}{(k-i)!(i-1)!},\tag{53}$$

where $(a)_k := a(a+1)\cdots(a+k-1)$ is the rising Pochhammer symbol.

Therefore, for $A \sim \mathcal{W}(0, \frac{1}{\alpha^2 m} \mathbb{I}_m) = \mathcal{W}$, the first several instances (starting at k = 1) of $\mathbb{E}[\text{Tr}(A^k)]$ are

$$\frac{m}{\alpha^2}, \frac{2m}{\alpha^4}, \frac{5m^2+1}{\alpha^6 m}, \frac{14m^2+10}{\alpha^8 m}, \frac{42m^4+70m^2+8}{\alpha^{10}m^3}, \dots$$
 (54)

This makes the expression for $\mathbb{E}[\langle n \rangle]$ challenging to write down and manipulate. Notice, however, that when α is large, the higher order terms are negligible. Formally, we state the following proposition:

Proposition 17. $\langle n \rangle \leq \text{Tr}(A) + 2m\lambda_{\max}(A)^2$ whenever $\lambda_{\max}(A) \leq 1/2$.

Proof. Start by writing $\langle n \rangle = \sum_{i=1}^m \sum_{k=1}^\infty \lambda_i^k$. We can bound all the higher-order terms (k > 1) as

$$\sum_{i=1}^{m} \sum_{k=2}^{\infty} \lambda_i^k \le m \sum_{k=2}^{\infty} \lambda_{\max}(A)^k \le 2m \lambda_{\max}(A)^2, \tag{55}$$

whenever $\lambda_{\max}(A) \leq 1/2$. What remains (the k=1 term) is simply the trace of A.

Importantly, we know that the maximum eigenvalue of $A \sim W$ is small (for large α) with high probability, due to Lemma 15. We also know the expectation of the trace:

Lemma 18 (Maiwald and Kraus [56]). $\mathbb{E}[\operatorname{Tr}(A)] = \frac{m}{\alpha^2}$ and $\mathbb{E}[\operatorname{Tr}(A)^2] = \frac{m^2+1}{\alpha^4}$.

We arrive at the following conclusion about the distribution of $\langle n \rangle$ over the choice of A:

Theorem 19.

$$\Pr\left[\left|\langle n\rangle - \frac{m}{\alpha^2}\right| \ge \frac{32\beta^2 m}{\alpha^4} + \frac{1}{\alpha^2} \sqrt{\frac{2}{\delta}}\right] \le \delta,\tag{56}$$

whenever $\beta > 4$, $\alpha^2 > 8\beta$, and $m > \beta^{-2} \ln(1/\delta)$.

Proof. By Lemma 15, for any $\beta \geq 4$, we have

$$\Pr[\lambda_{\max}(A) \ge 4\beta/\alpha^2] \le me^{-2m(\beta-1)^2} \le \frac{\delta}{2},\tag{57}$$

whenever $m \ge \beta^{-2} \ln(1/\delta)$, so let us assume $\lambda_{\max}(A) \le 4\beta/\alpha^2 \le 1/2$. Let us write $\langle n \rangle = \operatorname{Tr}(A) + \xi(A)$ where $\xi(A) := \sum_{k=2}^{\infty} \operatorname{Tr}(A^k)$. Using Proposition 17 and the fact that A is positive semidefinite, we have that $0 \le \xi(A) \le 32\beta^2 m/\alpha^4$.

Therefore, to understand the deviation of $\langle n \rangle$, it suffices to consider the variance of Tr(A). By Lemma 18, we have $\text{Var}[\text{Tr}(A)] = \mathbb{E}[\text{Tr}(A)^2] - \mathbb{E}[\text{Tr}(A)]^2 = 1/\alpha^4$, and so, by Chebyshev's inequality, we get

$$\Pr\left[|\operatorname{Tr}(A) - \mathbb{E}[\operatorname{Tr}(A)]| \ge \frac{1}{\alpha^2} \sqrt{\frac{2}{\delta}}\right] \le \frac{\delta}{2}.$$
 (58)

Combining our two bounds with the union bound, we get

$$|\langle n \rangle - \mathbb{E}[\text{Tr}(A)]| \le |\langle n \rangle - \text{Tr}(A)| + |\text{Tr}(A) - \mathbb{E}[\text{Tr}(A)]|$$

$$\le |\xi(A)| + \frac{1}{\alpha^2} \sqrt{\frac{2}{\delta}}$$

$$\le \frac{32\beta^2 m}{\alpha^4} + \frac{1}{\alpha^2} \sqrt{\frac{2}{\delta}},$$
(59)

which (using $\mathbb{E}[\text{Tr}(A)] = m/\alpha^2$) gives the theorem.

Therefore, assuming n is roughly equal to its expected value (over both the choice of A and the inherent randomness of the photon number), we can approximate n as m/α^2 whenever $\alpha = \Omega(m^{1/4})$ from Proposition 17. What we would like to know is how much n deviates from this value, for which we will need to compute the variance

$$\Delta^2[n] = \langle n^2 \rangle - \langle n \rangle^2 \tag{60}$$

Let us now write $n = \sum_{i=1}^{m} n_i$ where n_i is the number of photons generated at source in mode *i*. Expanding out the $\langle n^2 \rangle$ term of the variance, we get

$$\langle n^2 \rangle = \sum_{i=1}^m \langle n_i^2 \rangle + \sum_{i \neq j} \langle n_i \rangle \langle n_j \rangle = \langle n \rangle^2 + \langle n \rangle + \sum_{i=1}^m \langle n_i \rangle^2$$
 (61)

where the first equality uses $\langle n_i n_j \rangle = \langle n_i \rangle \langle n_j \rangle$ (i.e., the photons are generated independently at each mode), and the second equality uses the following identity for one half of a two-mode squeezed state:¹⁸

$$\langle n_i^2 \rangle = 2 \langle n_i \rangle^2 + \langle n_i \rangle. \tag{62}$$

Therefore, we can write the variance as

$$\Delta^{2}[n] = \langle n \rangle + \sum_{i=1}^{m} \left(\frac{\lambda_{i}}{1 - \lambda_{i}} \right)^{2} = \langle n \rangle + \sum_{i=1}^{m} \sum_{k=2}^{\infty} (k - 1) \lambda_{i}^{k}.$$
 (63)

Once again, we make the observation that when the maximum eigenvalue of A is small, the higher order terms in the Taylor expansion converge.

Proposition 20. $\Delta^2[n] \leq \langle n \rangle + 4m\lambda_{\max}(A)^2$, whenever $\lambda_{\max}(A) \leq 1/2$.

¹⁸See, for example, Eq. (3.5.15) of Barnett and Radmore [57]. For completeness, we give a self-contained proof of this fact in Appendix A.3.

Proof. We have $\Delta^2[n] = \langle n \rangle + \sum_{i=1}^m \sum_{k=2}^\infty (k-1) \lambda_i^k$. We bound the terms of the sum as

$$\sum_{i=1}^{m} \sum_{k=2}^{\infty} (k-1)\lambda_i^k \le m \sum_{k=2}^{\infty} (k-1)\lambda_{\max}(A)^k \le 4m\lambda_{\max}(A)^2$$
 (64)

whenever $\lambda_{\max}(A) \leq 1/2$.

We are finally ready to state and prove the main result for this section: a bound on how much n deviates from m/α^2 (roughly, its expected value) over both the randomness of $A \sim \mathcal{W}$ and the randomness of the photon number.

Theorem 21.

$$\Pr\left[\left|n - \frac{m}{\alpha^2}\right| \ge \frac{2\sqrt{m}}{\alpha\sqrt{\delta}} + \frac{3}{\alpha\delta^{3/4}} + \frac{21\beta\sqrt{m}}{\alpha^2\sqrt{\delta}} + \frac{32\beta^2m}{\alpha^4}\right] \le \delta,\tag{65}$$

whenever $\beta \geq 4$, $\alpha^2 \geq 8\beta$, and $m \geq 1/\beta^2 \ln(1/\delta)$.

Proof. Once again, we have $\lambda_{\max}(A) \leq 4\beta/\alpha^2 \leq 1/2$ with probability $\delta/4$ by Lemma 15, and so we have

$$\Delta^{2}[n] \le \langle n \rangle + \frac{64\beta^{2}m}{\alpha^{4}} \le \frac{m}{\alpha^{2}} + \frac{2}{\alpha^{2}\sqrt{\delta}} + \frac{96\beta^{2}m}{\alpha^{4}}$$
 (66)

where the first inequality comes from Proposition 20 and the second comes from Theorem 19 (which we assume holds with probability $\delta/2$). By Chebyshev's inequality, we get

$$\Pr\left[|n - \langle n \rangle| \ge \sqrt{\frac{4\Delta^2[n]}{\delta}}\right] \le \frac{\delta}{4},\tag{67}$$

where the probability is over the randomness in the photon number. And so, combining the previous inequalities with the union bound, we get

$$|n - m/\alpha^{2}| \leq |n - \langle n \rangle| + |\langle n \rangle - m/\alpha^{2}|$$

$$\leq \sqrt{\frac{4\Delta^{2}[n]}{\delta}} + \frac{2}{\alpha^{2}\sqrt{\delta}} + \frac{32\beta^{2}m}{\alpha^{4}}$$

$$\leq \frac{2}{\sqrt{\delta}}\sqrt{\frac{m}{\alpha^{2}}} + \frac{2}{\alpha^{2}\sqrt{\delta}} + \frac{96\beta^{2}m}{\alpha^{4}} + \frac{2}{\alpha^{2}\sqrt{\delta}} + \frac{32\beta^{2}m}{\alpha^{4}}$$

$$\leq \frac{2}{\sqrt{\delta}}\left(\sqrt{\frac{m}{\alpha^{2}}} + \sqrt{\frac{2}{\alpha^{2}\sqrt{\delta}}} + \sqrt{\frac{96\beta^{2}m}{\alpha^{4}}}\right) + \frac{2}{\alpha^{2}\sqrt{\delta}} + \frac{32\beta^{2}m}{\alpha^{4}}$$

$$\leq \frac{2\sqrt{m}}{\alpha\sqrt{\delta}} + \frac{3}{\alpha\delta^{3/4}} + \frac{21\beta\sqrt{m}}{\alpha^{2}\sqrt{\delta}} + \frac{32\beta^{2}m}{\alpha^{4}},$$
(68)

with probability δ , which completes the proof.

A.2 Bounds on normalization factor

We are now interested in the quantity

$$\mathcal{Z} = \frac{1}{\det(\mathbb{I}_m - CC^{\dagger})} = \prod_{i=1}^m \frac{1}{1 - \lambda_i}$$
 (69)

where λ_i is the *i*th eigenvalue of a complex Wishart matrix $A := CC^{\dagger} \sim \mathcal{W}$. We begin this section with the following fact about the moment generating function of $A \sim \mathcal{W}$:

Lemma 22 (Goodman [58]).

$$\mathbb{E}[\exp(\text{Tr}(A))] = \left(1 - \frac{1}{\alpha^2 m}\right)^{-m^2}.$$
 (70)

Our goal for this section is to prove the following:

Lemma 23.

$$\Pr\left[\mathcal{Z} \ge \frac{2}{\delta} e^{m/\alpha^2} e^{17\beta^2 m/\alpha^4}\right] \le \delta,\tag{71}$$

whenever $\beta \ge 4$, $\alpha^2 \ge 8\beta$, and $m \ge 1/\beta^2 \ln(1/\delta)$.

Proof. Recall that $\mathcal{Z} = \prod_{i=1}^{m} (1 - \lambda_i)^{-1}$. We have

$$\prod_{i=1}^{m} (1 - \lambda_i)^{-1} = \exp\left(-\sum_{i=1}^{m} \ln(1 - \lambda_i)\right) = \exp\left(\sum_{i=1}^{m} \sum_{k=1}^{\infty} \frac{\lambda_i^k}{k}\right),\tag{72}$$

where the last equality comes from the Taylor expansion of each term $\ln(1 - \lambda_i)$. First, we will show that we can bound the higher powers of the eigenvalues. By Lemma 15, for any $\beta \geq 4$, we have

$$\Pr[\lambda_{\max}(A) \ge 4\beta/\alpha^2] \le me^{-2m(\beta-1)^2} \le \frac{\delta}{2},\tag{73}$$

whenever $m \ge 1/\beta^2 \ln(1/\delta)$. Therefore, let's assume that $\lambda_{\max}(A) \le 4\beta/\alpha^2$. We obtain

$$\sum_{i=1}^{m} \sum_{k=2}^{\infty} \frac{\lambda_i^k}{k} \le m \sum_{k=2}^{\infty} \frac{(4\beta/\alpha^2)^k}{k} \le \frac{16\beta^2 m}{\alpha^4},\tag{74}$$

where we have assumed that $8\beta \leq \alpha^2$ for the last inequality. Combining (72) and (74), we get

$$\prod_{i=1}^{m} (1 - \lambda_i)^{-1} \le \exp\left(\sum_{i=1}^{m} \lambda_i\right) \exp\left(\frac{16\beta^2 m}{\alpha^4}\right) = \exp(\operatorname{Tr}(A)) \exp\left(\frac{16\beta^2 m}{\alpha^4}\right). \tag{75}$$

By Lemma 22, we have

$$\mathbb{E}[\exp \operatorname{Tr}(A)] = \left(1 - \frac{1}{\alpha^2 m}\right)^{-m^2} \le \exp(m/\alpha^2 + 1/\alpha^4) \tag{76}$$

where we've used a Taylor expansion and $\alpha \geq 2$ for the last inequality. By Markov's inequality we get

$$\exp \operatorname{Tr}(A) \le 2 \exp(m/\alpha^2 + 1/\alpha^4)/\delta,\tag{77}$$

with probability at least $1 - \delta/2$. The theorem follows from (75), (77), and the union bound.

A.3 Proof of thermal state identity

The purpose of this section is to give a self-contained proof of the identity

$$\langle n_i^2 \rangle = 2 \langle n_i \rangle^2 + \langle n_i \rangle, \tag{78}$$

which is used in Appendix A.1 to show that n is close to m/α^2 with high probability. To start, recall that in the main text we introduced the two-mode squeezing operator between modes i and i+m with squeezing parameter r as: $S_2(r) = \exp\left[r\left(a_i^{\dagger}a_{i+m}^{\dagger} - a_ia_{i+m}\right)\right] = \exp\left[rG\right]$, where $G = \frac{1}{2} \left[r\left(a_i^{\dagger}a_{i+m}^{\dagger} - a_ia_{i+m}\right)\right]$

 $a_i^{\dagger}a_{i+m}^{\dagger}-a_ia_{i+m}=-G^{\dagger}$ is an anti-Hermitian generator. We would like to investigate expectation values of operators O on the two-mode squeezed vacuum state

$$\langle O \rangle = \langle \psi | O | \psi \rangle = \langle 0_i 0_{i+m} | \underbrace{S_2^{\dagger}(r) O S_2(r)}_{\equiv O(r)} | 0_i 0_{i+m} \rangle \text{ with } | \psi \rangle = S_2(r) | 0_i 0_{i+m} \rangle.$$
 (79)

To make progress, we investigate the following quantity

$$\frac{\partial}{\partial r}O(r) = \left[\frac{\partial}{\partial r}e^{rG^{\dagger}}\right]Oe^{rG} + e^{rG^{\dagger}}O\left[\frac{\partial}{\partial r}e^{rG^{\dagger}}\right] = \left[-Ge^{rG^{\dagger}}\right]Oe^{rG} + e^{rG^{\dagger}}O\left[Ge^{rG}\right] = S_2(r)^{\dagger}[O, G]S_2(r). \tag{80}$$

In the last equation we used the fact that $G^{\dagger} = -G$ and introduced the commutator $[A, B] \equiv AB - BA$. Recalling the canonical bosonic commutation relations $[a_i, a_j] = [a_i^{\dagger}, a_j^{\dagger}] = 0$ and $[a_i, a_j^{\dagger}] = \delta_{i,j}$, we now set $O = a_i$ to obtain

$$\frac{\frac{\partial}{\partial r}a_i(r)}{\frac{\partial}{\partial r}a_{i+m}^{\dagger}(r)} = a_{i+m}^{\dagger}(r), \implies a_i(r) = a_i\cosh r + a_{i+m}^{\dagger}\sinh r, \\ \frac{\partial}{\partial r}a_{i+m}^{\dagger}(r) = a_i(r), \implies a_{i+m}^{\dagger}(r) = a_{i+m}^{\dagger}\cosh r + a_i\sinh r.$$
 (81)

Having derived these transformations we can now calculate

$$\langle \psi | a_i^{\dagger k} a_i^k | \psi \rangle = \langle 0_i 0_{i+m} | S_2^{\dagger}(r) a_i^{\dagger k} S_2(r) S_2^{\dagger}(r) a_i^k S_2(r) | 0_i 0_{i+m} \rangle$$

$$= \langle 0_i 0_{i+m} | \left[a_i^{\dagger} \cosh r + a_{i+m} \sinh r \right]^k \left[a_i \cosh r + a_{i+m}^{\dagger} \sinh r \right]^k | 0_i 0_{i+m} \rangle$$

$$= \langle 0_i 0_{i+m} | \left[a_{i+m} \sinh r \right]^k \left[a_{i+m}^{\dagger} \sinh r \right]^k | 0_i 0_{i+m} \rangle$$

$$= k! \sinh^{2k} r. \tag{82}$$

In the last equation we used the canonical commutation relations and the fact that the vacuum is annihilated by the annihilation operator. If we set k=1, we easily find the mean photon number to be $\langle n_i \rangle = \langle a_i^{\dagger} a_i \rangle = \sinh^2 r$. For the second moment, we write $\langle n_i^2 \rangle = \langle a_i^{\dagger} a_i a_i^{\dagger} a_i \rangle = \langle a_i^{\dagger} a_i^2 \rangle + \langle a_i^{\dagger} a_i \rangle$. The first term in the RHS of the last equality can be obtained by setting k=2 and we finally obtain $\langle n_i^2 \rangle = 2 \langle n_i \rangle^2 + \langle n_i \rangle$, as claimed.

B Technical lemmas for proving hardness of permanents with repetitions

In this appendix, we give proofs for the technical lemmas in Section 4. To start, let us show the proof of the expectation of a permanent of a matrix with repeated rows and columns which is used to define the allowable error fluctuations in the RGPE problem (see Definition 11).

Lemma 10. Let $A \sim \mathcal{N}(0,1)^{c \times c}_{\mathbb{C}}$. Then

$$\mathbb{E}\left[\left|\text{Per}(A_{S,T})\right|^{2}\right] = n! \prod_{i=1}^{c} s_{i}! \prod_{j=1}^{c} t_{j}!$$
(46)

for any repetition vectors $S, T \in \mathbb{N}^c$ such that $n = \sum_{i=1}^c s_i = \sum_{j=1}^c t_j$.

Proof. For simplicity, let us define the $n \times n$ matrix $X := A_{S,T}$ so that we have $\operatorname{Per}(X) = \sum_{\sigma \in \mathcal{S}_n} \prod_{i=1}^n X_{i,\sigma(i)}$. Because X has repeated rows and columns, notice that many terms $\prod_{i=1}^n X_{i,\sigma(i)}$ of the permanent are the same for different permutations σ . Let us rewrite the permanent so that such terms are grouped together. Notice that $A_{S,T}$ can be thought of as consisting of c^2 blocks, where the (i,j)th block is the constant matrix $A_{i,j}^{s_i \times t_j}$.

We block the permutations similarly. Concretely, given a permutation matrix over n elements, let $f: \mathcal{S}_n \to \mathbb{N}^{c \times c}$ be the function that sums all elements within the same block. For example, suppose we have S = T = (2, 2, 1). Then

$$\begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ \hline 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ \hline 1 & 0 & 0 & 0 & 0 \end{pmatrix} \xrightarrow{f} \begin{pmatrix} 1 & 0 & 1 \\ 0 & 2 & 0 \\ 1 & 0 & 0 \end{pmatrix}.$$

Let $\mathcal{M}_{S,T} := \{f(\sigma) : \sigma \in S_n\}$ be the set of matrices that arise from applying f to all permutations. We now ask, given a particular matrix $M \in \mathcal{M}_{S,T}$, how many times is M the image of some permutation σ ? Notice that if we start from a permutation matrix and permute say the first s_1 rows or t_1 columns, then its image under f is invariant. Generalizing, we get

$$|\{\sigma \in S_n : f(\sigma) = M\}| = \prod_{i=1}^{c} \prod_{j=1}^{c} \frac{s_i!t_j!}{M_{i,j}!}$$

where the denominator compensates for overcounting. Therefore, we can write

$$\operatorname{Per}(X) = \sum_{\sigma \in S_n} \prod_{i=1}^n X_{i,\sigma(i)}$$

$$= \sum_{\sigma \in S_n} \prod_{i=1}^c \prod_{j=1}^c A_{i,j}^{f(\sigma)_{i,j}}$$

$$= \sum_{M \in \mathcal{M}_{S,T}} \prod_{i=1}^c \prod_{j=1}^c \frac{s_i!t_j!}{M_{i,j}!} A_{i,j}^{M_{i,j}}$$

Recall that for a Gaussian random variable a, we have $\mathbb{E}[a^m(a^*)^{m'}]$ is m! when m=m', and 0 otherwise. Therefore,

$$\mathbb{E}[|\operatorname{Per}(X)|^{2}] = \mathbb{E}[\operatorname{Per}(X)\operatorname{Per}(X)^{*}]$$

$$= \sum_{M \in \mathcal{M}_{S,T}} \prod_{i=1}^{c} \prod_{j=1}^{c} \left(\frac{s_{i}!t_{j}!}{M_{i,j}!}\right)^{2} M_{i,j}!$$

$$= n! \prod_{i=1}^{c} s_{i}! \prod_{j=1}^{c} t_{j}!$$

where the last line comes from the fact

$$n! = |\{\sigma \in S_n\}| = \sum_{M \in \mathcal{M}_{S,T}} |\{\sigma \in S_n : f(\sigma) = M\}| = \sum_{M \in \mathcal{M}_{S,T}} \prod_{i=1}^{c} \prod_{j=1}^{c} \frac{s_i! t_j!}{M_{i,j}!}.$$

B.1 Expressing permanent of matrix with repetitions as polynomial

This section contains the proof of Theorem 12, which expresses the permanent of a matrix with repeated rows and columns as a polynomial whose constant coefficient is proportional to a permanent with no repetitions. We start with the proof of correctness for the construction of the matrix B' given in Section 4.1:

Lemma 24. Let $A \in \mathbb{C}^{c \times c}$ and $S \in \mathbb{N}^c$ such that $s_i \geq 1$ is given. Defining $n' := c + k_S$, let B' be a $c \times n'$ matrix obtained by adding k_S columns to the matrix A with entries that depend on complex variables z and $y = \{y_{i,j}^{(\ell)}\}$. Let B'_S be the $n' \times n'$ matrix which has s_i copies of row i of B'. We construct B' such that

$$Per(B'_S) = \xi_1 Per(A) + \gamma_1 z + \gamma_2 z^2 + \ldots + \gamma_{k_S} z^{k_S},$$
 (83)

$$\xi_1 = \prod_{i=1}^c \left(s_i! \prod_{j=1}^{s_i - 1} y_{i,j}^{(i)} \right) \tag{84}$$

where the coefficients $\gamma_1, \ldots, \gamma_{k_S}$ are functions of the entries of A, S, and $y_{i,j}^{(\ell)}$. Furthermore, whenever $A \sim \mathcal{N}(0,1)_{\mathbb{C}}^{c \times c}$, $y_{i,j}^{(\ell)} \sim \mathcal{N}(0,1)_{\mathbb{C}}^{c \times n'}$.

Proof. Recall that B' is constructed as follows: for each ℓ such that $s_{\ell} > 1$, append to A the $c \times (s_{\ell} - 1)$ matrix $V^{(\ell)}$. We have that $V_{ij}^{(\ell)} = y_{i,j}^{(\ell)}$ if $i = \ell$, and $zy_{i,j}^{(\ell)}$ otherwise. We refer the reader to Section 4.1 for an explicit example of this construction.

To prove the correctness of this procedure, recall that the permanent of any matrix $C = \{c_{i,j}\}_{i,j=1}^n$ can be written as the sum of monomials $\prod_{i=1}^n c_{i,\sigma(i)}$ over all permutations σ . Let us consider the monomial terms in the expression for $\operatorname{Per}(B_S')$ that are constant in z. We claim that for each such monomial, there must be at least one element from each row of the original matrix A. Suppose, by contradiction (and without loss of generality), that there is some monomial with no contribution from the elements in the first row of A. By construction, it follows that the monomial must have s_1 -many contributions from the first s_1 rows of the added columns (i.e., from $V^{(1)}, \ldots, V^{(c)}$). Notice that only $V^{(1)}$ has entries in those first s_1 rows that are not multiplied by z. However, $V^{(1)}$ only has $(s_1 - 1)$ -many columns, and hence the monomial in question must have some contribution from the the other added columns, which are all multiplied by the variable z.

Therefore, one element from each of the c rows of A appears in any monomial that is a constant in z. Observe that these c terms form a monomial of the permanent of A. The remaining k_S terms must come from the matrices $V^{(i)}$, for which the only entries that do not depend on z are of the form $y_{i,j}^{(i)}$. Furthermore, since we cannot chose two terms from the same row, this immediately implies the only choice for the remaining terms is $\prod_{j=1}^{s_i-1} y_{i,j}^{(i)}$. To complete the proof, we simply need to count how many such terms occur. For a particular monomial of the A, there are s_i ways to choose the element $a_{i,\sigma(i)}$, and there are $(s_i-1)!$ ways to choose the $\prod_{j=1}^{s_i-1} y_{i,j}^{(i)}$ term (since every permutation in the permanent of that submatrix suffices). Eq. (84) follows. The remaining conclusions of the lemma are direct consequences of the construction, which we leave to the reader.

Let us now describe in detail how Lemma 24 and its column version are combined to prove Theorem 12. Given $A \in \mathbb{C}^{c \times c}$, we will construct a $(c + k_T) \times (c + k_S)$ matrix B := B[z; x, y] to which we apply the repetition patterns $S = (s_1, \ldots, s_c, 1, \ldots, 1)$ and $T = (t_1, t_2, \ldots, t_c, 1, \ldots, 1)$. There are k_T trailing 1's in S, and k_S trailing 1's in T. The matrix B is obtained from A by appending the rows/columns which arise in the following two invocations of the Lemma 24:

- 1. Row version with pattern $S' = (s_1, \ldots, s_c)$ to A to obtain the $c \times n'$ matrix B' in variables z and $\{y_{i,j}^{(\ell)}\}$.
- 2. Column version with pattern T to the $n' \times n'$ matrix $B'_{S'}$ to obtain k_T new rows in variables z and $\{x_{i,j}^{(\ell)}\}$.

Correctness of this procedure follows from the lemmas:

$$Per(B_{S,T}) = \xi_2 Per(B'_{S'}) + zp_2(z) = \xi_2(\xi_1 Per(A) + zp_1(z)) + zp_2(z) = \xi Per(A) + zp(z),$$
(85)

where p(z) is a polynomial of degree k-1, and we have implicitly defined

$$\xi := \xi_1 \xi_2 = \prod_{i=1}^c s_i! t_i! \left(\prod_{j=1}^{t_i - 1} x_{i,j}^{(i)} \right) \left(\prod_{j=1}^{s_i - 1} y_{i,j}^{(i)} \right). \tag{86}$$

B.2 Polynomial interpolation techniques

Starting from Eq. (47), which we restate,

$$Per(B_{S,T}) = \xi Per(A) + \gamma_1 z + \gamma_2 z^2 + \dots + \gamma_k z^k, \tag{87}$$

we give the proofs for the polynomial interpolation theorems given in Section 4.2 to extract the constant coefficient. We begin with a general technique for approximating the constant term in a complex polynomial:

Lemma 25. Let $P(z) = \gamma_0 + \sum_{n=1}^k \gamma_n z^n$ be a complex polynomial. Assume that there is a procedure to obtain an estimate $\tilde{P}(z) = P(z) + e(z)$ for any z, where e(z) is the error in the estimate, satisfying $|e(z)| \le \epsilon$ for all z. Then

$$\tilde{\gamma}_0 = \frac{1}{k+1} \sum_{j=0}^k \tilde{P}\left(e^{2\pi i j/(k+1)}\right)$$
(88)

is an estimate of γ_0 such that $|\tilde{\gamma}_0 - \gamma_0| \leq \epsilon$.

Proof. The proof is by construction. For each $j=0,1,\ldots,k$, define $z_j=e^{2\pi i j/(k+1)}$. Then it holds that

$$\sum_{j=0}^{k} P(z_j) = \sum_{j=0}^{k} \left(\gamma_0 + \sum_{n=1}^{k} \gamma_n z_j^n \right)$$
$$= (k+1)\gamma_0 + \sum_{n=1}^{k} \gamma_n \sum_{j=0}^{k} z_j^n$$
$$= (k+1)\gamma_0,$$

where we used the fact that $\sum_{j=0}^{k} z_{j}^{n} = 0$ for all $n \leq k$. It follows that

$$\gamma_0 = \frac{1}{k+1} \sum_{j=0}^k P(z_j). \tag{89}$$

We then have

$$|\tilde{\gamma}_0 - \gamma_0| = \left| \frac{1}{k+1} \sum_{j=0}^k \tilde{P}(z_j) - P(z_j) \right|$$
$$= \frac{1}{k+1} \left| \sum_{j=0}^k e(z_j) \right| \le \epsilon,$$

as desired. \Box

Combining lemmas 24 and 25, we have shown that given an oracle to estimate $Per(B_{S,T})$, it is possible to also estimate Per(A) by calling the oracle k+1 times for values of z equally distributed along the unit circle. Indeed, defining $\tilde{P} = \tilde{\gamma}_0/\xi$, with ξ as in Eq. (48), we have $|\tilde{P} - Per(A)|$ is at most the error of the oracle times $1/|\xi|$. We now use this to provide a reduction from GPE_{\pm} to $RGPE_{\pm}$, as per Definition 11.

Theorem 13. Given access to an oracle for RGPE_± with error $\varepsilon \sqrt{(c+k)! \prod s_i! t_j!}$, it is possible to use k+1 calls to the oracle to solve GPE_± with error $\varepsilon \sqrt{c!}$ where $\varepsilon := \varepsilon (\frac{1}{|\xi|} \sqrt{((c+k)!/c!) \prod s_i! t_j!})$ with ξ as in Eq. (48).

Proof. Let $A \in \mathbb{C}^{c \times c}$ be the input matrix to GPE_{\pm} . Following Theorem 12, we build a matrix B such that $\text{Per}(B_{S,T}) = \xi \, \text{Per}(A) + zp(z)$. Recall that if we chose |z| = 1 and the variables $x_{i,j}^{(\ell)}, y_{i,j}^{(\ell)}$ to be from the same Gaussian distribution as the entries of A, this ensures that $B_{S,T}$ is a Gaussian random matrix (with repetitions) that can be fed into the oracle for RGPE_{\pm} .

Now repeat this construction for $z_j = e^{2\pi i j/(k+1)}$ to obtain the matrices $B[z_j; x, y]_{S,T}$ for each $j = 0, 1, \ldots, k+1$. For each such matrix, use the oracle to compute an estimate $\tilde{P}(z_j)$ of its permanent. Then we output an estimate \tilde{P} of $\xi \operatorname{Per}(A)$ given by

$$\tilde{P} = \frac{1}{(k+1)} \sum_{j=0}^{k} \tilde{P}(z_j). \tag{90}$$

Using Lemma 25, we have $|\tilde{P} - \xi \operatorname{Per}(A)| \leq \varepsilon \sqrt{(c+k)! \prod s_i! t_j!}$, and so the error in the estimate \tilde{P}/ξ for $\operatorname{Per}(A)$ is

$$|\tilde{P}/\xi - \operatorname{Per}(A)| = \frac{1}{|\xi|} |\tilde{P} - \xi \operatorname{Per}(A)| \le \left(\varepsilon \frac{1}{|\xi|} \sqrt{((c+k)!/c!) \prod s_i! t_j!}\right) \sqrt{c!} = \epsilon \sqrt{c!}$$

as claimed. \Box

To understand ϵ in Theorem 13, we must understand the size of $|\xi|$. To this end, in Appendix B.3 we will prove that $|\xi| = \Omega(1.498^k)$ with constant probability for sufficiently large k. To succeed with probability $1-\delta$ for any $\delta>0$, we note that the random Gaussian variables appearing in the definition of $|\xi|$ come from the additional Gaussian variables (the $x_{i,j}^{(\ell)}$ and $y_{i,j}^{(\ell)}$ terms) employed in the reduction. Therefore, if those variables are not sufficiently large, we can simply resample them $O(\log(1/\delta))$ -many times until they are.

Let us now turn to our second polynomial interpolation technique, which works directly with the magnitude of the permanent:

Theorem 14. Given access to an oracle for $|RGPE|_{\pm}^2$ with error $\varepsilon(c+k)! \prod s_i! t_j!$, it is possible to use O(k) calls to the oracle to solve $|GPE|_{\pm}^2$ with additive error $\epsilon c!$, for

$$\epsilon := O\left(\frac{(c+k)^{2k}k^{2k}}{\delta^{3k+1/2}2^k}\varepsilon\right) \tag{50}$$

Whenever k = O(1), we have $\epsilon = \mathsf{poly}(c, \varepsilon, \delta)$.

Proof. We begin by taking the absolute squared value of Eq. (47). This leads to an equation of the form

$$|\operatorname{Per}(B_{S,T})|^2 = |\xi|^2 |\operatorname{Per}(A)|^2 + q(z)$$
 (91)

where q(z) is polynomial in z and \bar{z} with q(0) = 0. Contrary to Theorem 13, we cannot compute the left-hand side for z in the unit circle and leverage the corresponding cancellations, since q(z) has monomials that depend on $|z|^2$. However, notice that when z is real, Eq. (91) is of the same form as Eq. (8) in [36]. Unsurprisingly, the remainder of this proof closely follows that in [36], and so, we will just walk through their proof outline and emphasize the differences that arise in our case but omit repetition of technical details.

The first step is to choose $x_{i,j}^{(i)}$ and $y_{i,j}^{(i)}$ in $B_{S,T}$ to be i.i.d. Gaussian values from $\mathcal{N}(0,1)_{\mathbb{C}}$ (i.e., the same distribution as the entries of A), and then choose different (real and positive) values of z in the range $[1-\gamma,1+\gamma]$, for a suitably small γ . Note that, if $\gamma=0$, then matrix B has only i.i.d. elements

from $\mathcal{N}(0,1)_{\mathbb{C}}$, as per Theorem 12. When we choose a different value of γ we induce a deformation in the distribution of B, since now some of its elements are drawn from $\mathcal{N}(0,z^2)_{\mathbb{C}}$ instead.

If γ is too large, then the oracle for $|\mathrm{RGPE}|_\pm^2$ may fail since it only succeeds with high probability over the Gaussian distribution. To bound how large γ can be before this happen, we compute the total variation distance between the distributions over B when $\gamma=0$ and when $\gamma>0$ (its deformed version). If that distance is $O(\delta)$, then the oracle for $|\mathrm{RGPE}|_\pm^2$ will still succeed with probability $O(\delta)$ and can be invoked to compute $|\mathrm{Per}(B_{S,T})|^2$ to the desired precision. Luckily, there is a closed form for the Kullback-Leibler divergence between two sets of i.i.d. Gaussian variables which, combined with Pinsker's inequality, gives a bound on the desired total variation distance (cf. Lemma 1 of [36]). The only thing remaining is to count how many matrix elements of B have actually been deformed via multiplication by z. Following the proof of Theorem 12, we can see that $(c^2 + k_S + k_T)$ elements of B are standard complex Gaussians, and the remaining $(c-1)k + k_S k_T$ have been deformed. This implies that we must set

$$\gamma \le \frac{\delta}{\sqrt{(c-1)k + k_S k_T}},\tag{92}$$

for the two distributions to be sufficiently close.

The second step is to apply the polynomial interpolation. This is done by invoking the oracle to estimate the left-hand side of Eq. (91), to error $\varepsilon(c+k)! \prod s_i!t_j!$, for 2k+1 values of c evenly spaced in the interval $[1-\gamma, 1+\gamma]$. This leads to a linear system which can be solved by a standard least squares method. If we denote by $\hat{\beta}_0$ the resulting estimator for the constant term in the polynomial of Eq. (91), the variance of this estimator is given by

$$\operatorname{Var}[\hat{\beta}_0] = \frac{\left(\varepsilon(c+k)! \prod s_i! t_j!\right)^2}{\gamma^{4k}},\tag{93}$$

which is the same as Eq. (20) in [36] when translated to our variables. By Chebyshev's inequality, we get

$$\Pr\left[\left||\operatorname{Per}(A)|^{2} - \hat{\beta}_{0}|\xi|^{-2}\right| \ge |\xi|^{-2}\sqrt{\operatorname{Var}[\hat{\beta}_{0}]/\delta}\right] \le \delta,\tag{94}$$

and so we obtain a bound for $|\text{Per}(A)|^2$ that holds with probability $1 - \delta$. We identify this bound with the error in our estimate for $|\text{Per}(A)|^2$, which needs to be at most $\epsilon c!$ by the definition of $|\text{GPE}|_{\pm}^2$. Combining all of the above, we obtain the following relation

$$\epsilon = \varepsilon \frac{(c+k)! \left((c-1)k + k_S k_T \right)^k}{c! X \delta^{2k+1/2} \prod s_i! t_i!} \le \varepsilon \frac{(c+k)^{2k} k^k}{X \delta^{2k+1/2} 2^k},\tag{95}$$

where $X = \prod_{j=1}^{t_i-1} |x_{i,j}^{(i)}|^2 \prod_{j=1}^{s_i-1} |y_{i,j}^{(i)}|^2$ is the product of k independent absolute-squared standard i.i.d. complex Gaussian variables (coming from the definition of ξ in Eq. (86)). Each of the k terms, which we will generically call x, follows an exponential distribution with CDF

$$\Pr[x \le g] = 1 - e^{-g}. (96)$$

If we set $g = \delta/k$, then $\Pr[x \leq g] \leq 1 - (1 - \delta)^{1/k}$, or in other words $\Pr[x > g] > (1 - \delta)^{1/k}$. By the independence of the k terms comprising X, we get $\Pr[1/X < 1/g^k] > 1 - \delta$. Therefore we can assume that, with probability at least $1 - \delta$, 1/X is at most k^k/δ^k . Plugging this into Eq. (95), we get that

$$\epsilon = O\left(\frac{(c+k)^{2k}k^{2k}}{\delta^{3k+1/2}2^k}\varepsilon\right),\tag{97}$$

which is polynomial in c, δ and ε as long as k = O(1), as desired.

B.3 Lower bound for $|\xi|$

The purpose of this section is to give a lower bound on the $|\xi|$ term coming from Eq. (48) in Theorem 13:

$$|\xi| = \prod_{i=1}^{c} s_i! t_i! \left(\prod_{j=1}^{t_i-1} |x_{i,j}^{(i)}| \right) \left(\prod_{j=1}^{s_i-1} |y_{i,j}^{(i)}| \right).$$

$$(98)$$

where $\boldsymbol{x}_{i,j}^{(i)}, \boldsymbol{y}_{i,j}^{(i)}$ are i.i.d. standard complex Gaussians. We show

Theorem 26. $\Pr[|\xi| \ge 1.498^k] \ge 1/2 - 4/\sqrt{k}$.

To prove this theorem, let us first focus on bounding the product of the k Gaussian magnitudes that appear in Eq. (98). To simplify the notation, let us simply rewrite this product as $X = \prod_{i=1}^k X_i$ where each X_i is the random variable for the magnitude of a standard complex Gaussian. One can check that $X_i \sim \text{Rayleigh}(1/2)$ such that $\Pr[X_i = x] = 2xe^{-x^2}$ for $x \in [0, \infty)$. To bound X, we will need the following standard central limit theorem:

Theorem 27 (Berry-Esseen Theorem). Let Z_1, \ldots, Z_k be real i.i.d. random variables satisfying

$$\mathbb{E}[Z_i] = \mu,$$

$$\mathbb{E}[(Z_i - \mu)^2] = \sigma^2 > 0,$$

$$\mathbb{E}[|Z_i - \mu|^3] = \rho < \infty$$

Let $Z = \sum_{i=1}^k Z_i$ and let $W \sim \mathcal{N}(\mu k, \sigma^2 k)_{\mathbb{R}}$ be a real Gaussian. Then, for all $x \in \mathbb{R}$,

$$|\Pr[Z > x] - \Pr[W > x]| \le \frac{C\rho}{\sigma^3 \sqrt{k}},\tag{99}$$

where C is some universal constant.

Lemma 28. $\Pr[X \ge .7493^k] \ge 1/2 - 4/\sqrt{k}$.

Proof. Letting $Z_i = \log(X_i)$ and $Z = \sum_{i=1}^k Z_i$, we begin by writing

$$X = e^{\sum_{i=1}^{k} \log(X_i)} = e^Z.$$
 (100)

One can check that $\Pr[Z_i = z] = 2e^{2z - e^{2z}}$ leading to the following moments:

$$\mathbb{E}[Z_i] = -\gamma/2,\tag{101}$$

$$\mathbb{E}[(Z_i + \gamma/2)^2] = \pi^2/24,\tag{102}$$

$$\mathbb{E}[|Z_i + \gamma/2|^3] < 0.5136 \tag{103}$$

where $\gamma \approx 0.5772$ is the Euler-Macheroni constant. Plugging Eqs. (101), (102), and (103) into the Berry-Esseen Theorem (using C < .4748 [59]), we get that for any $x \in \mathbb{R}$

$$|\Pr[Z > x] - \Pr[W > x]| < 4/\sqrt{k} \tag{104}$$

where W is real Gaussian with mean $-k\gamma/2$ and variance $k\pi/24$. This immediately gives the lower bound

$$\begin{split} \Pr[Z > -k\gamma/2] &= \Pr[Z > -k\gamma/2] - \Pr[W > -k\gamma/2] + \Pr[W > -k\gamma/2] \\ &\geq \Pr[W > -k\gamma/2] - |\Pr[Z > -k\gamma/2] - \Pr[W > -k\gamma/2]| \\ &\geq 1/2 - 4/\sqrt{k} \end{split}$$

where the last line uses Eq. (104) and the simple fact that a Gaussian is greater than its mean with probability 1/2. Of course, if $Z > -k\gamma/2$, then $X > (e^{-\gamma/2})^k$, which yields the lemma.

Using that $\prod_{i=1}^{c} s_i!t_i! \geq 2^k$ immediately gives Theorem 26.

C Moments of the click distribution

In this appendix we derive the results provided in Section 2.3.3 on click moments. In Appendix C.1, we present a brief review of the Gaussian formalism in terms of covariance matrices an apply it to derive the Husimi covariance matrix of a Gaussian state in which a transition matrix has been encoded. In Appendix C.2, we describe how to reduce the calculation of the mean number of clicks and its variance for an arbitrary multimode state to the calculation of one- and two-mode vacuum probabilities. In Appendix C.3, we show how to fix the scaling factor α in terms of the photon number density using the quarter circle law. Finally, in Appendix C.4, we put together all the results derived previously and derive Eqs. (24) and (25) giving the first and second order moments of the click distribution for the two halves of the modes respectively.

C.1 Gaussian states and covariance matrices

In this appendix we provide a self-contained review of the Gaussian formalism as relevant to the discussion in the main text and the subsequent appendices. We consider a Gaussian state ρ with M modes and zero mean. The latter condition means that $\langle a_i \rangle = \text{Tr}(a_i \rho) = 0$ for all $i \in \{1, ..., M\}$. The state ρ is uniquely characterized by its (complex-)Husimi covariance matrix Q with entries

$$Q_{i,j} = \frac{1}{2} \langle z_i z_j^{\dagger} + z_j^{\dagger} z_i \rangle + \frac{1}{2} \delta_{i,j} = \frac{1}{2} \text{Tr} \left(\rho \left[z_i z_j^{\dagger} + z_j^{\dagger} z_i \right] \right) + \frac{1}{2} \delta_{i,j}, \tag{105}$$

where

$$z = \left(a_1, \dots, a_M, a_1^{\dagger}, \dots, a_M^{\dagger}\right). \tag{106}$$

For Q to represent a valid quantum mechanical state it needs to satisfy the uncertainty relation [60]

$$Q + \frac{1}{2}Z - \frac{1}{2}\mathbb{I}_{2M} \succeq 0, \tag{107}$$

where $Z=\left(\begin{smallmatrix}\mathbb{I}_M&0\\0&-\mathbb{I}_M\end{smallmatrix}\right)$, with \mathbb{I}_M the $M\times M$ identity matrix.

The marginal state of a subset of the modes of a Gaussian state is another Gaussian state. In particular the covariance matrix of a subset of the modes i_1, \ldots, i_K is obtained by keeping rows and columns i_1, \ldots, i_K and $i_1 + M, \ldots, i_K + M$ from the original covariance matrix Q. Finally, note that the diagonal elements of the Husimi covariance matrix are related to the mean photon number in each of the modes as

$$Q_{i,i} = Q_{i+M,i+M} = \langle a_i^{\dagger} a_i \rangle + 1. \tag{108}$$

For a given zero-mean Gaussian state the probability of obtaining vacuum in all its modes when measuring its photon number statistics is given by

$$\Pr(\text{vac}) = \frac{1}{\sqrt{\det(Q)}}.$$
(109)

A simple expression exists linking a pure Gaussian state parametrized by an adjacency matrix $A = A^T \in \mathbb{C}^{M \times M}$ and its Husimi covariance matrix [7, 19]

$$Q = (\mathbb{I}_{2M} - \mathcal{X}(A \oplus A^*))^{-1}, \text{ where } \mathcal{X} = \begin{pmatrix} 0 & \mathbb{I}_M \\ \mathbb{I}_M & 0 \end{pmatrix}.$$
 (110)

We now consider the case where M=2m and we use A to encode a transition matrix $C\in\mathbb{C}^{m\times m}$ as follows

$$A = \begin{pmatrix} 0 & C \\ C^T & 0 \end{pmatrix}. \tag{111}$$

We write the following decomposition $C = U\sigma V^T$ with U and V $m \times m$ unitary matrices and σ a diagonal matrix with $0 \le \sigma_{i,i} < 1$. We can then write

$$\mathcal{X}(A \oplus A^*) = \begin{pmatrix} 0 & 0 & 0 & C \\ 0 & 0 & C^T & 0 \\ 0 & C^* & 0 & 0 \\ C^{\dagger} & 0 & 0 & 0 \end{pmatrix} = F \begin{pmatrix} \sigma & 0 & 0 & 0 \\ 0 & \sigma & 0 & 0 \\ 0 & 0 & -\sigma & 0 \\ 0 & 0 & 0 & -\sigma \end{pmatrix} F^{\dagger}, \tag{112}$$

where

$$F = \frac{1}{\sqrt{2}} \begin{pmatrix} U & 0 & -U & 0\\ 0 & V & 0 & -V\\ 0 & U^* & 0 & U^*\\ V^* & 0 & V^* & 0 \end{pmatrix}, \tag{113}$$

is a unitary matrix of size $4m \times 4m$. With this diagonalization it is direct to write

$$Q = \begin{pmatrix} X & 0 & 0 & W \\ 0 & Y & W^T & 0 \\ 0 & W^* & X^* & 0 \\ W^{\dagger} & 0 & 0 & Y^* \end{pmatrix}, \tag{114}$$

with

$$X = U \frac{\mathbb{I}_m}{\mathbb{I}_m - \sigma^2} U^{\dagger} = X^{\dagger}, \quad Y = V \frac{\mathbb{I}_m}{\mathbb{I}_m - \sigma^2} V^{\dagger} = Y^{\dagger}, \quad W = U \frac{\sigma}{\mathbb{I}_m - \sigma^2} V^T. \tag{115}$$

This form of the covariance matrix will be useful below.

C.2 Detector click statistics

In this section we derive simple expressions for the mean and variance of the number of clicks when a quantum state is probed using threshold detectors. We will obtain expressions for $\langle c \rangle$, $\langle c^2 \rangle$ and $\Delta^2 c$ that hold for an arbitrary M-mode quantum state. In the next subsections we will specialize these expressions to zero-mean Gaussian states encoding a random complex-Gaussian transition matrix.

Let ν_c be a subset of $\{1, 2, ..., M\}$ where ν_c contains exactly c elements. Furthermore, let $p(\nu_c)$ denote the probability of observing a click in all detectors with labels in ν_c . This probability is given by [61]

$$p(\nu_c) = \text{Tr}\left[\rho\left(\prod_{j\in\nu_c} P_1^{(j)} \prod_{j\notin\nu_c} P_0^{(j)}\right)\right],\tag{116}$$

where $P_0^{(j)} = |0_j\rangle\langle 0_j|$ is a projector onto the vacuum in mode j and $P_1^{(j)} = \mathbbm{1}^{(j)} - P_0^{(j)}$. The probability of observing c clicks is then given by the sum $p(c) = \sum_{\nu_c} p(\nu_c)$. Now consider its characteristic function:

$$\chi_{c}(x) = \langle e^{icx} \rangle = \sum_{c=0}^{M} e^{ixc} p(c)$$

$$= \sum_{c=0}^{M} e^{ixc} \sum_{\nu_{c}} \operatorname{Tr} \left[\rho \left(\prod_{k \in \nu_{c}} P_{1}^{(k)} \prod_{k \notin \nu_{c}} P_{0}^{(k)} \right) \right]$$

$$= \operatorname{Tr} \left[\rho \sum_{c=0}^{M} \sum_{\nu_{c}} \left(\prod_{k \in \nu_{c}} e^{ix} P_{1}^{(k)} \prod_{k \notin \nu_{c}} P_{0}^{(k)} \right) \right]$$

$$= \operatorname{Tr} \left[\rho \prod_{\ell=1}^{M} \left(e^{ix} P_{1}^{(\ell)} + P_{0}^{(\ell)} \right) \right] , \tag{117}$$

where in the last line we used the binomial expansion. The expected number of detectors that click $\langle c \rangle$ can then be computed from the characteristic function [18]

$$\langle c \rangle = -i \frac{d\chi_c(x)}{dx} \Big|_{x=0}$$

$$= -i \text{Tr} \left[\rho \frac{d}{dx} \prod_{\ell=1}^{M} \left(e^{ix} P_1^{(\ell)} + P_0^{(\ell)} \right) \right] \Big|_{x=0}$$

$$= \text{Tr} \left[\rho \sum_{j=1}^{M} P_1^{(j)} \prod_{\ell \neq j} \left(P_1^{(\ell)} + P_0^{(\ell)} \right) \right]$$

$$= \sum_{j=1}^{M} \text{Tr} \left[P_1^{(j)} \rho \right], \qquad (118)$$

where we have used the product rule for derivatives. Using the fact that $P_1^{(j)} + P_0^{(j)} = \mathbb{1}^{(j)}$, we conclude that

$$\langle c \rangle = M - \sum_{i=1}^{M} \langle P_0^{(j)} \rangle,$$
 (119)

where we have defined $\langle P_0^{(j)} \rangle := \text{Tr} \left[P_0^{(j)} \rho \right]$. Similarly, we can obtain the second moment of the number of clicks as

$$\langle c^2 \rangle = (-i)^2 \frac{d^2 \chi_c(x)}{dx^2} \bigg|_{x=0}$$
 (120)

We now write

$$\frac{d^2 \chi_c(x)}{dx^2} = \sum_{j=1}^M \text{Tr} \left(\rho \frac{d}{dx} \prod_{\ell \neq j} \left[e^{ix} P_1^{(\ell)} + P_0^{(\ell)} \right] i P_1^{(j)} e^{ix} \right)
= \sum_{j=1}^M \text{Tr} \left(\rho \left\{ i^2 P_1^{(j)} e^{ix} \prod_{\ell \neq j} \left[e^{ix} P_1^{(j)} + P_0^{(\ell)} \right] + i^2 P_1^{(j)} e^{ix} \sum_{k \neq j}^M P_1^{(k)} e^{ix} \prod_{\ell \neq j, \ell \neq k} \left[e^{ix} P_1^{(\ell)} + P_0^{(\ell)} \right] \right\} \right)$$

By setting x = 0 in the last equation we obtain

$$\langle c^2 \rangle = \sum_{j=1}^M \text{Tr}[P_1^{(j)} \rho_j] + \sum_{j=1}^M \sum_{j \neq k}^M \text{Tr}[P_1^{(j)} P_1^{(k)} \rho]. \tag{121}$$

Equivalently, we can use $P_1^{(j)} = \mathbb{1}^{(j)} - P_0^{(j)}$, which gives

$$\langle c^2 \rangle = M^2 - (2M - 1) \sum_{j=1}^{M} \langle P_0^{(j)} \rangle + 2 \sum_{j < k}^{M} \langle P_0^{(j)} P_0^{(k)} \rangle,$$
 (122)

where we have defined $\langle P_0^{(j)} P_0^{(k)} \rangle := \text{Tr} \left[P_0^{(j)} P_0^{(k)} \rho \right]$. Putting these results together, the variance of the number of clicks is

$$\Delta^{2}c = \sum_{j=1}^{M} \langle P_{0}^{(j)} \rangle \left(1 - \langle P_{0}^{(j)} \rangle \right) + 2\sum_{j < k} \left(\langle P_{0}^{(j)} P_{0}^{(k)} \rangle - \langle P_{0}^{(j)} \rangle \langle P_{0}^{(k)} \rangle \right) = \sum_{i,j=1}^{M} \left(\langle P_{0}^{(j)} P_{0}^{(k)} \rangle - \langle P_{0}^{(j)} \rangle \langle P_{0}^{(k)} \rangle \right). \tag{123}$$

Since reduced states of Gaussian states can be computed efficiently [62], it is possible to use these formulas to calculate the average and variance of the number of clicks for these states. Under the assumption of zero-mean Gaussian states, the two types of terms in the equations above can be written as

$$\langle P_0^{(j)} P_0^{(k)} \rangle = \frac{1}{\sqrt{\det(Q^{i,j})}}, \quad \langle P_0^{(j)} \rangle = \frac{1}{\sqrt{\det(Q^j)}}, \tag{124}$$

where $Q^{i,j}$ and Q^j are the reduced Husimi covariance matrices of modes i, j and j respectively. We note that these quantities can also be estimated efficiently for Gaussian states using Monte Carlo methods as recently shown in Ref. [63].

In the following, we use these results to estimate the average number of clicks for GBS devices that are configured to encode Gaussian random matrices.

C.3 Fixing the scaling factor in terms of the photon number density

In Eq. (14) we used the quarter circle law to find a relation between the scaling factor α , the number of modes m and the total mean photon number. From this equation we can solve

$$\alpha(\mu) = \frac{1+\mu}{\sqrt{\mu}}.\tag{125}$$

For future convenience we can also calculate the following matrix norms of the blocks of the Husimi covariance matrix (recall Eq. (115)) when the transition matrix has singular values satisfying the quarter circle law

$$\mathbb{E}\left(||X||^{2}\right) = \mathbb{E}\left(||Y||^{2}\right) = \mathbb{E}\left(\sum_{i=1}^{m} \left(\frac{1}{1-\sigma_{i}^{2}}\right)^{2}\right) = m \int_{0}^{2/\alpha(\mu)} p_{\alpha(\mu)}(\sigma) \left(\frac{1}{1-\sigma^{2}}\right)^{2} = m \frac{(1+\mu)}{1-\mu}, \quad (126)$$

$$\mathbb{E}\left(||W||^2\right) = \mathbb{E}\left(\sum_{i=1}^m \left(\frac{\sigma_i}{1-\sigma_i^2}\right)^2\right) = m \int_0^{2/\alpha(\mu)} p_{\alpha(\mu)}(\sigma) \left(\frac{\sigma}{1-\sigma^2}\right)^2 = m \frac{\mu(1+\mu)}{1-\mu}.$$
 (127)

In the equation above we used the fact that the Frobenius norm of a matrix, $||A||^2 = \sum_{i,j=1}^m |A_{i,j}|^2$ can also be expressed as the sum of the squares of the singular values of the same matrix.

C.4 Means and variances of the click distribution

With the preliminary calculations in the last subsections we are ready to tackle the calculation of the means and covariances of the click distribution between the two sets of modes when averaged over the set of Gaussian random matrices.

For this we start with the means, for which we find the covariance of mode i in the first half

$$Q^{i} = \begin{pmatrix} X_{i,i} & 0 \\ 0 & X_{i,i} \end{pmatrix} \longrightarrow \langle P_{0}^{(i)} \rangle = \frac{1}{X_{i,i}} = \frac{1}{1 + \langle a_{i}^{\dagger} a_{i} \rangle}, \tag{128}$$

and similarly if i = j + m is in the second half

$$Q^{i} = \begin{pmatrix} Y_{j,j} & 0 \\ 0 & Y_{j,j} \end{pmatrix} \longrightarrow \langle P_{0}^{(i)} \rangle = \frac{1}{Y_{j,j}} = \frac{1}{1 + \langle a_{i}^{\dagger} a_{i} \rangle}. \tag{129}$$

From the last two equations we find $\langle d \rangle = m - \sum_{i=1}^{m} \frac{1}{X_{j,j}}$ and similarly $\langle e \rangle = m - \sum_{i=1}^{m} \frac{1}{Y_{j,j}}$. Now we can let $X_{j,j} \to 1 + \mu$ and $Y_{j,j} \to 1 + \mu$ and rearrange to find precisely Eq. (24).

Now we consider the covariances between the clicks in the two halves. For this we need the twobody covariance matrices of the modes. Consider first the case where i and j refer to modes in the first half of the set, i.e., $i, j \leq m$. In that case their reduced covariance matrix is given by

$$Q^{i,j} = \begin{pmatrix} X_{i,i} & X_{i,j} & 0 & 0 \\ X_{i,j}^* & X_{j,j} & 0 & 0 \\ 0 & 0 & X_{i,i} & X_{i,j}^* \\ 0 & 0 & X_{i,j} & X_{j,j} \end{pmatrix} \longrightarrow \langle P_0^{(j)} P_0^{(i)} \rangle - \langle P_0^{(j)} \rangle \langle P_0^{(j)} \rangle = \frac{1}{X_{i,i} X_{j,j} - |X_{i,j}|^2} - \frac{1}{X_{i,i} X_{j,j}}$$

$$= \frac{1}{X_{i,i} X_{j,j}} \sum_{k=1}^{\infty} \left[\frac{|X_{i,j}|^2}{X_{i,i} X_{j,j}} \right]^k. \quad (131)$$

Note that the Taylor expansion is guaranteed to converge since the uncertainty relation for the Q covariance matrix (cf. Appendix C.1) guarantees that $X_{i,i}X_{j,j} > |X_{i,j}|^2$. If the two modes i, j are in the second half one can obtain a similar expression to the one above by letting $X \to Y$

Finally, if i and j are in different halves one has

$$Q^{i,j} = \begin{pmatrix} X_{i,i} & 0 & 0 & W_{i,j} \\ 0 & Y_{j,j} & W_{i,j} & 0 \\ 0 & W_{i,j}^* & X_{i,i} & 0 \\ W_{i,j}^* & 0 & 0 & Y_{j,j} \end{pmatrix} \longrightarrow \langle P_0^{(j)} P_0^{(i)} \rangle - \langle P_0^{(j)} \rangle \langle P_0^{(j)} \rangle = \frac{1}{X_{i,i} Y_{j,j} - |W_{i,j}|^2} - \frac{1}{X_{i,i} Y_{j,j}}$$

$$= \frac{1}{X_{i,i} Y_{j,j}} \sum_{k=1}^{\infty} \left[\frac{|W_{i,j}|^2}{X_{i,i} Y_{j,j}} \right]^k. \quad (133)$$

As before, the Taylor expansion converges by virtue of the uncertainty relation.

We can now look at the variance of the total number of clicks (recall Eq. (123)) over the two halves to write it as

$$\Delta^{2}[d+e] = \sum_{i=1}^{m} \left(\frac{1}{X_{i,i}} + \frac{1}{Y_{i,i}} - \frac{1}{X_{i,i}^{2}} - \frac{1}{Y_{i,i}^{2}} \right)$$

$$+ \sum_{i \neq j}^{m} \left(\frac{1}{X_{i,i}X_{j,j} - |X_{i,j}|^{2}} - \frac{1}{X_{i,i}X_{j,j}} + \frac{1}{Y_{i,i}Y_{j,j} - |Y_{i,j}|^{2}} - \frac{1}{Y_{i,i}Y_{j,j}} \right)$$

$$+ 2 \sum_{i,j=1}^{m} \left(\frac{1}{X_{i,i}Y_{j,j} - |W_{i,j}|^{2}} - \frac{1}{X_{i,i}Y_{j,j}} \right).$$

$$(134)$$

We will now subtract the quantity $\sum_{i=1}^{m} \frac{1}{X_{i,i}^2} + \frac{1}{Y_{i,i}^2}$ in the first term of the equation above and add it in the second. Then we will also Taylor expand the terms in the second and third line to write

$$\Delta^{2}[d+e] = \sum_{i=1}^{m} \left(\frac{1}{X_{i,i}} + \frac{1}{Y_{i,i}} - \frac{2}{X_{i,i}^{2}} - \frac{2}{Y_{i,i}^{2}} \right)$$

$$+ \sum_{i,j}^{m} \frac{|X_{i,j}|^{2}}{X_{i,i}^{2}X_{j,j}^{2}} + \frac{|Y_{i,j}|^{2}}{Y_{i,i}^{2}Y_{j,j}^{2}} + \sum_{i \neq j}^{m} \frac{1}{X_{i,i}X_{j,j}} \sum_{k=2}^{\infty} \left(\frac{|X_{i,j}|^{2}}{X_{i,i}X_{j,j}} \right)^{k} + \sum_{i \neq j}^{m} \frac{1}{Y_{i,i}Y_{j,j}} \sum_{k=2}^{\infty} \left(\frac{|Y_{i,j}|^{2}}{Y_{i,i}Y_{j,j}} \right)^{k}$$

$$+ 2 \sum_{i,j=1}^{m} \frac{|W_{i,j}|^{2}}{X_{i,i}^{2}Y_{j,j}^{2}} + 2 \sum_{i,j=1}^{m} \frac{1}{X_{i,i}Y_{j,j}} \sum_{k=2}^{\infty} \left(\frac{|W_{i,j}|^{2}}{X_{i,i}Y_{j,j}} \right)^{k} .$$

$$= \epsilon_{2}$$

In the last equation we introduced the errors $\epsilon_1, \epsilon_2 \geq 0$. We can then write the variance as

$$\Delta^{2}[d+e] = \sum_{i=1}^{m} \left(\frac{1}{X_{i,i}} + \frac{1}{Y_{i,i}} - \frac{2}{X_{i,i}^{2}} - \frac{2}{Y_{i,i}^{2}} \right) + \sum_{i,j=1}^{m} \left(\frac{|X_{i,j}|^{2}}{X_{i,i}^{2}X_{j,j}^{2}} + \frac{|Y_{i,j}|^{2}}{Y_{i,i}^{2}Y_{j,j}^{2}} + 2\frac{|W_{i,j}|^{2}}{X_{i,i}^{2}Y_{j,j}^{2}} \right) + \epsilon, \quad (136)$$

where $\epsilon = \epsilon_1 + \epsilon_2$. At this point we replace the diagonal elements $X_{i,i}, Y_{j,j}$ by their expectation $1 + \mu$ and use the expressions for the Frobenius norms of X, Y and W to find

$$\mathbb{E}\left(\Delta^{2}[d+e]\right) \approx 2m \frac{(2-\mu)\mu}{(1-\mu)(1+\mu)^{2}},\tag{137}$$

where we ignore the small error $\mathbb{E}(\epsilon)$.

For the variance in the two halves we can easily write

$$\Delta^2 d = \sum_{i=1}^m \left(\frac{1}{X_{i,i}} - \frac{1}{X_{i,i}^2} \right) + \sum_{i \neq j}^m \left(\frac{1}{X_{i,i} X_{j,j} - |X_{i,j}|^2} - \frac{1}{X_{i,i} X_{j,j}} \right), \tag{138}$$

$$\Delta^{2}e = \sum_{i=1}^{m} \left(\frac{1}{Y_{i,i}} - \frac{1}{Y_{i,i}^{2}} \right) + \sum_{i \neq j}^{m} \left(\frac{1}{Y_{i,i}Y_{j,j} - |Y_{i,j}|^{2}} - \frac{1}{Y_{i,i}Y_{j,j}} \right), \tag{139}$$

and use the same approximations used previously to derive the results in Eq. (25a). Finally, we can obtain the covariance between d and e and the variance of their difference by recalling that $\Delta^2[d\pm e] = \Delta^2 d + \Delta^2 e \pm \text{Cov}(d, e)$.

D Going beyond the dilute limit

The purpose of this section is to give evidence that Theorem 9 can be extended beyond the dilute limit to show: $|RGPE|_{\pm}^2 \in FBPP^{NP^{\mathcal{O}}}$ for some set of collision patterns that have high probability in the BipartiteGBS distribution. To do this, we first claim (without proof)¹⁹ that a generalization of Theorem 7 exists where the subspace \mathcal{H}_S can be replaced by any subspace \mathcal{H} for which $S \in \mathcal{H}$ implies $S_{\pi,\sigma} \in \mathcal{H}$ for all $\pi, \sigma \in S_m$. Specifically, let us choose \mathcal{H} to be the space of repetition patterns with $\mathbb{E}[\langle n \rangle]$ photon pairs and $\mathbb{E}[\langle c \rangle]/2$ clicks in each half of the modes (see Eq. (14) and Eq. (24), respectively). For this section, we will simply write these quantities as n and c, respectively. Therefore, we have

$$|\mathcal{H}| = {m \choose c}^2 {n-1 \choose n-c}^2. \tag{140}$$

Then, the generalization of Corollary 8 allows us to obtain a

$$\varepsilon \frac{\mathcal{Z}\alpha^{2n}m^n \prod_{i=1}^m s_i! \prod_{j=1}^m t_j!}{|\mathcal{H}|}$$
(141)

approximation to $|\operatorname{Per}(A_{S,T})|^2$ where the pattern (S,T) was chosen uniformly from \mathcal{H} . On the other hand, to solve $|\operatorname{RGPE}|_{\pm}^2$, we need an $(\epsilon n! \prod_{i=1}^m s_i! \prod_{j=1}^m t_j!)$ -additive approximation. Following the same logic from Theorem 9, to obtain an FBPP^{NP°} algorithm, we want

$$\mathcal{I} := \epsilon/\varepsilon = \frac{\mathcal{Z}\alpha^{2n}m^n \prod_{i=1}^m s_i! \prod_{j=1}^m t_j!}{|\mathcal{H}|n! \prod_{i=1}^m s_i! \prod_{j=1}^m t_j!} = \frac{\mathcal{Z}\alpha^{2n}m^n}{|\mathcal{H}|n!}.$$
 (142)

¹⁹This may at first seem counterintuitive since a classical sampler may choose to favor certain collision patterns over others. However, the critical point is that the classical sampler does not know which collision pattern to favor since it will be sampled randomly.

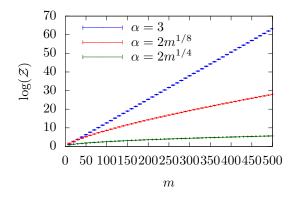


Figure 4: Comparison of $1/\mathbb{E}[\mathcal{Z}^{-1}]$ and \mathcal{Z} as a function of m and various α (shown on a log scale). Each point was calculated from Lemma 29, and the error bars show the sample standard deviation of $\log(\mathcal{Z})$ calculated from 50 complex Wishart matrices drawn from $\mathcal{W}(0, \frac{1}{\alpha^2 m} \mathbb{I}_m)$.

to be polynomially bounded. Unfortunately, we can no longer use Lemma 5 to bound this quantity. Because $\alpha = o(m^{1/4})$, the subdominant term of that expression (e^{272m/α^4}) would *still* contribute exponentially to Eq. (142). Since that expression was obtained through several approximations, it is no longer tight enough. Nevertheless, we claim that the empirical scaling of Eq. (142) still appears to be polynomial.

Let us now describe how those numerics were obtained. Given that we do not have an explicit bound on \mathcal{Z} , we must choose to estimate it somehow. One approach would be to simply draw a random Gaussian matrix and compute \mathcal{Z} from the eigenvalues of that matrix. Unfortunately, this approach is quite slow in practice and we would only be able to calculate \mathcal{Z} for relatively modest numbers of modes. Instead, we will choose to estimate \mathcal{Z} by an explicit formula. First, consider the quantity

$$Z^{-1} = \det(\mathbb{I}_m - A) = \prod_{i=1}^m (1 - \lambda_i)$$
(143)

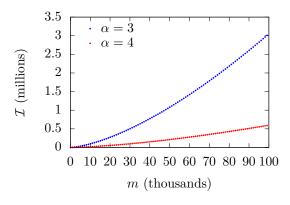
for complex Wishart matrix $A \sim W$ (see Appendix A for explanation of this ensemble). In other words, the *inverse* of \mathcal{Z} is the characteristic polynomial of A. The expectation of the characteristic polynomial for a Wishart matrix is well-understood:

Lemma 29 (Edelman [64]).
$$\mathbb{E}[\mathcal{Z}^{-1}] = m! \sum_{i=0}^{m} {m \choose i} \frac{(-1)^i}{(m-i)!(\alpha^2 m)^i}$$
.

In Fig. 4, we show that the quantity $1/\mathbb{E}[\mathcal{Z}^{-1}]$ serves as a good estimate for \mathcal{Z} by comparing it against its true value (computed directly with the eigenvalues). In Fig. 5, we compute \mathcal{I} from Eq. (142) using the estimate $1/\mathbb{E}[\mathcal{Z}^{-1}]$ for \mathcal{Z} and find a relatively benign scaling. We conclude by reiterating that if we could *prove* the accuracy of our estimator and that this scaling is only polynomial, then we arrive at our desired conclusion: $|\mathrm{RGPE}|_{\pm}^2 \in \mathrm{FBPP}^{\mathsf{NP}^{\mathcal{O}}}$. Therefore, if one were to also assume that $|\mathrm{RGPE}|_{\pm}^2$ is $\#\mathsf{P}$ -hard, we obtain approximate average-case hardness for GBS in the high-collision $m = \Theta(\mathbb{E}[\langle n \rangle])$ regime.

D.1 Justification for collision subspace

The purpose of this section is to justify that the subspace \mathcal{H} only contains collision patterns that suffice for the reduction given by Lemma 24. While it is known that Lemma 24 is insufficient to prove an efficient reduction between GPE_{\pm} and $RGPE_{\pm}$, it is still weak evidence that the permanents that appear in the probabilities are hard to estimate. To recap that procedure, we embed the permanent of matrix A into the permanent of a matrix $B_{S,T}$ with repeated rows/columns provided the repetition



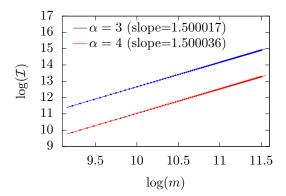


Figure 5: Left: $\mathcal{I}:=\frac{\mathcal{Z}\alpha^{2n}m^n}{|\mathcal{H}|n!}$ as a function of m and various α (domain and range have been scaled by 10^3 and 10^6 , respectively). Large values of α , such as those corresponding to the dilute limit are dwarfed by the scaling of \mathcal{I} at these small values of α , and so are not shown. Right: The same data as the left graph, except shown on a log-log plot and only values of m greater than 10^4 . Linear regression lines suggest that \mathcal{I} scales like $m^{3/2}$ for small constant α .

patterns S and T have a least k_T and k_S modes (respectively) with a single photon.²⁰ Notice that \mathcal{H} contains repetition patterns with $k_T = k_S = n - c$. So, to match the reduction, n - c of the c modes that click must have exactly one photon. We will show that under a relatively modest condition on α , all repetition patterns in \mathcal{H} will satisfy this property.

While this may seem like a rather strong condition, the fact that the number of modes that click is always a significant fraction of the total number of photons makes it possible. First notice that there are at most n-c modes (of our total budget of c clicks) that can have more than two photons. For the reduction, we then need a different set of n-c modes to click. That is, we want

$$2(n-c) \le c. \tag{144}$$

Using Eq. (14) and Eq. (24), this will hold whenever $\alpha \geq 3/\sqrt{2}$.

Finally, we claim that A will be a $(2c-n) \times (2c-n)$ matrix because the total number of clicks must be the dimension of A plus the n-c extra clicks required for the reduction (i.e., c=(2c-n)+(n-c)). In order to obtain a hardness proof, we would like A to be relatively large so that computing its permanent is #P-hard. We have that the size of A is at least half the total number of clicks in the distribution whenever $2c-n \ge c/2$. Notice that this is the same condition as that in Eq. (144), and so we only require $\alpha \ge 3/\sqrt{2}$. This is more than sufficient to guarantee hardness for all regimes from the dilute limit all the way to when the number of photons is linear in the number of modes. We stress again, however, that our reduction is only efficient given an oracle for $\text{Per}(B_{S,T})$, and the numerics of the previous section only suggest an oracle for $|\text{Per}(B_{S,T})|^2$.

 $^{^{20}}$ Recall that while k_T and k_S are the number of collisions, they are also the number of modes that we need to have without collisions. This is due to the fact that every repeated row/column requires an unrepeated column/row in the reduction.