Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Signal and interference constraints on cryptographic uses of coarse-grained boson sampling

Hikaru Wakaura

hikaruwakaura@gmail.com

QuantScape Inc

Research Article

Keywords: boson sampling, coarse-graining, quantum computational advantage, classical simulation, quantum cryptography, proof-of-work

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11867655/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Signal and interference constraints on cryptographic uses of coarse-grained boson sampling

Hikaru Wakaura[1,*]

[1]QIRI (Quantum Integrated Research Institute Inc.), 1-16-3 Akasaka, Minato-ku, Tokyo 107–0061, Japan
[*]Author to whom any correspondence should be addressed.

E-mail: h.wakaura@qiri.co.jp

Keywords: boson sampling, coarse-graining, quantum computational advantage, classical simulation, quantum cryptography, proof-of-work

Abstract

Boson sampling is a leading route to quantum computational advantage with photons. Because individual outcome probabilities are exponentially small, applications coarse-grain the output into polynomially many bins, and cryptographic proposals (one-way functions, hash functions, proof-of-work) assume that the bin probabilities are accessible to a boson sampler but not to a classical computer. Classical algorithms now estimate linear coarse-grainings efficiently, whereas permuted and nonlinear binnings remain open. It is unclear which coarse-grainings retain a signal that is both readable from polynomially many samples and genuinely due to many-photon interference. Here we show that each proposal we examine fails at least one of these requirements. For a random permutation of the outcome space we prove that every bin probability equals the uniform value up to a deviation bounded by the collision probability of the output distribution; at six photons a uniform random sampler passes the reward test of a proposed proof-of-work protocol in all trials. A hash function built on three-mode photon-number correlations of Gaussian states is evaluated exactly from the covariance matrix, in under a second for 256 modes. For lexicographic binning the signal decays only polynomially but is carried mostly by single-particle transition probabilities: at 16 photons interference contributes about a quarter of it even with sixteen times more bins than modes, and a product formula for distinguishable particles ranks the most probable bin within its top 1% of candidates in over 90% of instances from ten photons on. A coarse-graining suitable for cryptography must therefore carry a signal that survives the removal of low-order correlations. These criteria give a quantitative test for future proposals and delimit where verifiable or cryptographic uses of photonic samplers can still be sought.

1 Introduction

Photonic devices have reached the regime in which sampling from their output distribution is beyond direct classical simulation [1, 2]. The underlying tasks, boson sampling with single photons [3] and its Gaussian variant [4], were designed to make this separation provable under complexity-theoretic assumptions rather than to compute anything of independent interest [5]. Whether such a sampler can do something useful has therefore become a central question for the platform.

Every proposed use has to deal with the same obstacle. The probability of an individual outcome is exponentially small in the number of photons, so no experiment can estimate it. What an experiment can estimate, from polynomially many samples, is the probability of a group of outcomes. Applications therefore coarse-grain: they map each outcome to one of polynomially many bins and work with the bin probabilities. Molecular vibronic spectra [6] and dense-subgraph search [7] are of this form, and so are certification with coarse-grained measurements [8] and the validation tests that group detectors or modes [9, 10, 11]. For the first two, classical algorithms that match the sampler are now known [12, 13].

Cryptography is the application in which coarse-graining carries the most weight. Nikolopoulos and Brougham proposed to bin the outcome space and to use the label of the most probable bin as the output of a function that a boson sampler evaluates and a classical computer supposedly cannot [14]. This idea underlies a one-way function [15], its experimental demonstration [16], an indistinguishability-based scheme [17], hash functions [18, 19] and a proof-of-work consensus protocol [20]. Recently Seron *et al.* showed that any coarse-graining through a linear function of the output occupation numbers, which includes the binning of Refs. [14, 15], can be estimated

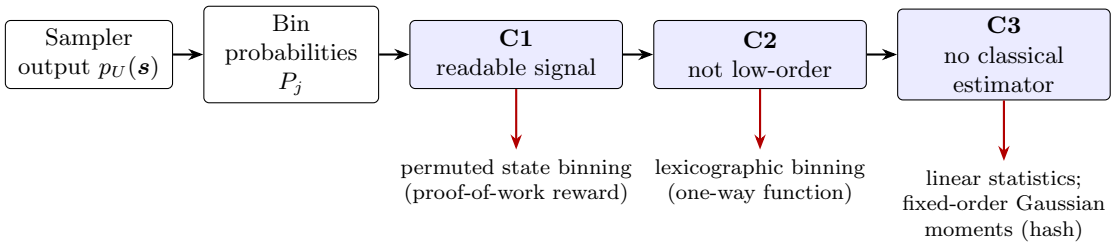


Figure 1. A cryptographic coarse-graining must pass three conditions, and each scheme examined stops at one of them. Outcomes $\boldsymbol{s}$ of a boson sampler with interferometer U are mapped to d bins by a public function g. The bin probabilities must deviate from their Haar mean by an inverse polynomial (C1), the deviation must not be reproduced by single-particle or fixed-order correlations (C2), and no efficient classical algorithm may estimate it (C3). Red arrows mark the first condition that each family of schemes fails.

classically in polynomial time [21]. They left open the schemes that permute the outcome space before binning [20] or that group outcomes nonlinearly [18], and pointed to quadratic functions of the occupation numbers as a possible hard case.

However, for the coarse-grainings outside the reach of that algorithm it has not been quantified whether the bin probabilities still depend measurably on the interferometer, nor whether the dependence that remains is produced by many-photon interference rather than by single-particle transition probabilities, so that the absence of a classical algorithm cannot by itself be read as evidence that a quantum device is needed.

In this study, we formulate three conditions that a coarse-graining must satisfy to support a cryptographic use: a signal of inverse-polynomial size, a signal that is not exhausted by low-order correlations, and the absence of an efficient classical estimator. We then test existing schemes against each condition by proof and by exact computation up to 20 photons. We found that a generic permutation of the outcome space removes the signal altogether, that fixed-order correlations of Gaussian states are classically exact, and that lexicographic binning keeps a polynomially large signal which a product formula for distinguishable particles largely reproduces. The result shows that each scheme examined misses at least one condition (Fig. 1 and Table 1), and that the open region is narrower than the absence of classical algorithms suggests.

The three conditions apply beyond cryptography. They are the same requirements that a coarse-grained observable must meet to certify many-photon interference, and they give experiments a way to budget samples: the number needed to see the signal and the number needed to see its interference part differ by a factor that grows with the number of modes.

2 Framework

2.1 Coarse-grained boson sampling

N single photons enter the first N of M modes of a linear interferometer described by a unitary U. An outcome is an occupation list $\boldsymbol{s} = (s_1, \ldots, s_M)$ with $\sum_j s_j = N$, or equivalently the sorted list of occupied modes $i_1 \le i_2 \le \cdots \le i_N$. The outcome space Y has $|Y| = \binom{M+N-1}{N}$ elements and

$$p_U(\boldsymbol{s}) = \frac{|\mathrm{Per}\, U_{\boldsymbol{s}}|^2}{\prod_j s_j!}, \tag{1}$$

where $U_{\boldsymbol{s}}$ is the $N \times N$ matrix built from the first N columns of U and from row j repeated s_j times [3]. A coarse-graining is a publicly known map $g : Y \to \{1, \ldots, d\}$ with d polynomial in N and M. Its bin probabilities are $P_j = \sum_{g(\boldsymbol{s})=j} p_U(\boldsymbol{s})$, and the most probable bin (MPB) is the label that the schemes of Refs. [14, 15, 16] use as the function value. We write $M = N^2$ for the dilute regime and $M = 2N$ for the dense one.

2.2 Haar mean and signal

The N-photon states span the symmetric subspace of $(\mathbb{C}^M)^{\otimes N}$, on which the interferometers act irreducibly. Averaging the output state over Haar-random U therefore gives the maximally mixed state on that subspace, and

$$\mathbb{E}_U\, p_U(\boldsymbol{s}) = \frac{1}{|Y|} \quad \text{for every } \boldsymbol{s}, \qquad \mathbb{E}_U\, P_j = w_j := \frac{|g^{-1}(j)|}{|Y|}. \tag{2}$$

Table 1. Coarse-grainings used in cryptographic proposals based on boson sampling and, for those examined here, the condition each fails.

Scheme	Coarse-graining	Fails	Evidence
One-way function and related schemes [14, 15, 16, 17]	contiguous bins in lexicographic order	C2, C3	Sec. 5; Ref. [21]
Proof-of-work, validation step [20]	groups of modes (linear)	C3	by design [20, 9]
Proof-of-work, reward step [20]	permutation of Y, then equal bins	C1	Sec. 3, for a generic permutation
Photonic hash [19]	digit of a three-mode photon-number correlation of a Gaussian state	C3	Sec. 4
Hash with grouped coarse-graining [18]	collision-dependent grouping	—	outside the tests of this work; open in Ref. [21]

The weights w_j are fixed by counting and carry no information about U. We call $\delta_j = P_j - w_j$ the signal of bin j. A sampler that ignores U and draws outcomes uniformly from Y reproduces w_j exactly, so a scheme can test or exploit the device only through δ_j.

2.3 Three conditions

C1, readable signal. For typical U, $\max_j |\delta_j|$ is at least an inverse polynomial in N. Otherwise polynomially many samples cannot distinguish the device from the uniform sampler, and the scheme tests nothing.

C2, signal not exhausted by low-order correlations. The feature that the scheme uses, such as the MPB label, is not determined by quantities that involve the interference of a fixed number of photons. The reference case is fully distinguishable particles, whose bin probabilities depend on U only through the single-particle transition probabilities $|U_{jk}|^2$. Two-mode correlators already separate bosons from distinguishable particles [22], and the contributions in which only a fixed number of photons interfere can be computed classically, which underlies the simulation of partially distinguishable photons [23]; a signal that such quantities reproduce does not require a many-photon interferometer.

C3, no efficient classical estimator. No classical algorithm estimates the P_j to inverse-polynomial additive error in polynomial time.

C3 is the condition usually discussed. C1 and C2 are necessary for it to be meaningful: without C1 there is nothing to estimate, and where C2 fails the signal is accessible without interference even when no general-purpose estimator has been written down. Failing C2 is weaker than failing C3, since it concerns a specific predictor rather than all classical algorithms, and we state it separately for that reason. We do not attempt to establish C3 for any coarse-graining, which would be a complexity-theoretic statement; it enters our results only where it fails.

Table 1 lists the coarse-grainings of the existing schemes and, for those we examined, the condition each one fails.

3 Permuted binning carries no signal

In the reward step of the proof-of-work protocol of Ref. [20], a public random beacon is mapped to a permutation π of the whole outcome space, and the permuted list is cut into d consecutive blocks of equal size. Miners are rewarded when the peak bin probability of their own samples agrees, within a tolerance ϵ, with that of the pooled samples of the network. The protocol assigns to this step the role of requiring a boson sampler, on the ground that no efficient classical simulation of the state-binned distribution is known. The map from the beacon to π is not specified. We analyse the case in which π behaves as a uniformly random permutation.

Lemma 1 (Flatness)**.** *Let p be any probability distribution on a set Y and let d divide $|Y|$. Partition Y uniformly at random into d bins of size $|Y|/d$ and let P_j be the probability of bin j. Then $\mathbb{E}[P_j] = 1/d$ and*

$$\mathrm{Var}[P_j] = \frac{d-1}{d^2}\,\frac{|Y|}{|Y|-1}\Big(\sum_{\boldsymbol{s}} p(\boldsymbol{s})^2 - \frac{1}{|Y|}\Big). \tag{3}$$

Proof. Write $K = |Y|$ and $k = K/d$, and let $\chi_{\boldsymbol{s}}$ indicate that $\boldsymbol{s}$ falls in bin j. The bin is a uniformly random k-subset, so $\mathrm{Var}[\chi_{\boldsymbol{s}}] = k(K-k)/K^2$ and $\mathrm{Cov}[\chi_{\boldsymbol{s}}, \chi_{\boldsymbol{t}}] = -k(K-k)/\big(K^2(K-1)\big)$ for $\boldsymbol{s} \neq \boldsymbol{t}$.

Using $\sum_{\boldsymbol{s}\neq\boldsymbol{t}} p(\boldsymbol{s})p(\boldsymbol{t}) = 1 - \sum_{\boldsymbol{s}} p(\boldsymbol{s})^2$,

$$\mathrm{Var}\Big[\sum_{\boldsymbol{s}} p(\boldsymbol{s})\chi_{\boldsymbol{s}}\Big] = \frac{k(K-k)}{K^2}\Big(\sum_{\boldsymbol{s}} p(\boldsymbol{s})^2 - \frac{1-\sum_{\boldsymbol{s}} p(\boldsymbol{s})^2}{K-1}\Big) = \frac{k(K-k)}{K(K-1)}\Big(\sum_{\boldsymbol{s}} p(\boldsymbol{s})^2 - \frac{1}{K}\Big),$$

and $k = K/d$ gives Eq. (3). □

Corollary 1. *By Chebyshev's inequality and a union bound over the bins,*

$$\Pr\Big[\max_j \big|P_j - \tfrac{1}{d}\big| \geq \varepsilon\Big] < \frac{1}{\varepsilon^2}\sum_{\boldsymbol{s}} p(\boldsymbol{s})^2. \tag{4}$$

The lemma holds for every distribution; the only property of boson sampling that enters is the collision probability $\sum_{\boldsymbol{s}} p_U(\boldsymbol{s})^2$. Its Haar average, normalised as $P_2 = |Y|\,\mathbb{E}_U \sum_{\boldsymbol{s}} p_U(\boldsymbol{s})^2$, is known in closed form for all N and M [24]: it grows linearly with N in the dilute regime and is asymptotically $M/N + 1$ when M grows more slowly than N^2. Evaluating that expression, we find $P_2 < 0.72\,(N+1)$ for $M = N^2$ and $P_2 < 3.5$ for $M = 2N$ for all $N \leq 40$; the value $N+1$ is the one given by the fourth moment of the permanent of a Gaussian matrix, $(N+1)(N!)^2$ [3]. Our exact output distributions with 3 to 8 photons give $|Y|\sum_{\boldsymbol{s}} p_U(\boldsymbol{s})^2$ between 2.6 and 4.9, each within 1.1 standard errors of the closed-form average for its size (2.6 to 4.8). By Markov's inequality the collision probability of an interferometer exceeds a given multiple of this average only with correspondingly small probability, so for typical interferometers the deviation from $1/d$ is at most of order $(N/|Y|)^{1/2}$, exponentially small in N. Figure 2(a) confirms this: under a random permutation the measured variance agrees with Eq. (3) to within 2%, and the mean largest deviation falls as $|Y|^{-0.46\pm0.01}$, whereas lexicographic bins of the same size keep deviations between 0.06 and 0.12.

To see what this means for the protocol we simulated the reward criterion alone, with the sample size $1.8 \times 10^5 d^{7/2}$ and the tolerance $\epsilon = (0.05/d)^{5/4}$ prescribed in Ref. [20]. A miner who holds one tenth of the pooled samples and draws them uniformly from Y is rewarded in 6% of trials at $N = 4$, $M = 16$ (between 0% and 12% depending on the interferometer), in 92% at $N = 5$, $M = 25$ (77% to 98%), and in all 10 000 trials at $N = 6$, $M = 36$, for $d = 4$ [Fig. 2(b)]; with $d = 16$ at $N = 6$ the fraction is 99%, and at least 95% for every interferometer. Each case uses 100 trials for each of 100 Haar-random interferometers. An honest miner is rewarded in all trials. With lexicographic bins the uniform miner is never rewarded.

For a generic permutation the reward step therefore fails C1: from about six photons on, the peak bin probability is $1/d$ for any sampler, and the step does not test whether a boson sampler was used. Two limits of this statement should be kept in view. It does not cover a structured map, for instance a permutation of the modes followed by lexicographic bins, which leads back to the case of Sec. 5. And it concerns the reward step in isolation; a miner must first pass the mode-binning validation, which rests on committing samples before the binning is announced, and we make no claim about the protocol as a whole.

4 Fixed-order correlations of Gaussian states are classically exact

The hash function of Ref. [19] sets the squeezing parameter of mode j to the input bit $b_j \in \{0, 1\}$, sends the resulting Gaussian state through a brickwork interferometer of depth equal to the number of modes, and outputs for each mode j one decimal digit of the three-mode correlation $\mu_j = \langle n_j n_{j+1} n_{j+2}\rangle$, reduced modulo 2. Three-mode correlations are chosen there because one- and two-mode correlations are regarded as classically easy.

For a Gaussian state, a photon-number moment of fixed order is a polynomial in the entries of the covariance matrix by Wick's theorem [25]. The covariance matrix of the output state follows from U and the squeezing parameters in $O(M^3)$ operations, and each μ_j is then a sixth-order moment of the quadratures of three modes, obtained in constant time. We evaluated the hash in this way with a standard library routine [26]: 256 modes take 0.27 s on a laptop, with no sampling. The classical value is exact to machine precision, whereas a device estimates μ_j with a statistical error that must be pushed below the extracted digit.

The scheme therefore fails C3. This says nothing about preimage or collision resistance, which we have not examined; it says that the function is a classical hash function, and that the photonic device is not needed to evaluate it. Correlations become hard only when their order grows with the system size, which is the regime of the sampling problem itself.

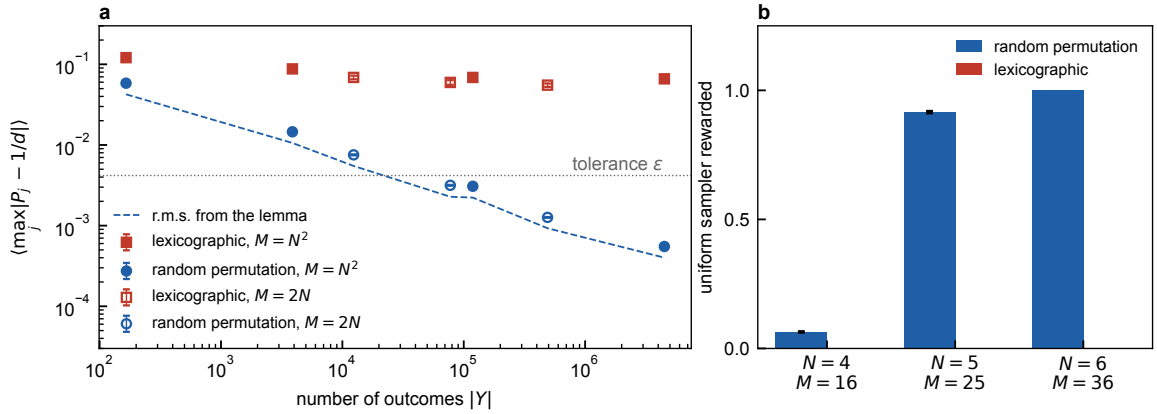


Figure 2. A random permutation of the outcome space erases the signal exponentially, whereas lexicographic order retains it. (a) Largest deviation of a bin probability from $1/d$ for $d = 4$, averaged over 40 permutations and over 100 to 200 Haar-random interferometers per size, against the number of outcomes; $N = 3$ to 8; error bars are the standard error over interferometers and are mostly smaller than the symbols. The dashed line is the root-mean-square deviation predicted by Eq. (3) from the measured collision probability, and the dotted line is the tolerance of the reward test. (b) Fraction of trials in which a miner drawing uniformly random outcomes passes the reward test of Ref. [20]; 100 trials for each of 100 interferometers per size, $d = 4$. Error bars are the standard error over interferometers, 0.003 at $N = 4$ and 0.005 at $N = 5$; they are absent where all 10 000 trials gave the same outcome.

5 Lexicographic binning: a single-particle signal

5.1 Exact bin probabilities

Order the outcomes lexicographically by their sorted mode lists $(i_1, \dots, i_N)$ and cut the ordered list into contiguous bins, as in Refs. [14, 15, 16]. This coarse-graining is a linear statistic and falls under the algorithm of Ref. [21], so C3 fails. We ask about C1 and C2, for which exact bin probabilities at large N are needed.

If the bin boundaries are placed at prefixes of the mode list, the tail probabilities are coefficients of a single permanent. Let V be the $M \times N$ matrix of the first N columns of U, $\Pi_{\geq a}$ the projector on modes $a, \dots, M$ and Π_x the projector on mode x. Then

$$\Pr[i_1 \geq a] = \mathrm{Per}\big(V^\dagger \Pi_{\geq a} V\big), \tag{5}$$

and, for $x < y < a$, the probability of finding k photons in mode x, l in mode y and the rest in modes $\geq a$ is the coefficient of $z_1^k z_2^l$ in $\mathrm{Per}\big(V^\dagger(\Pi_{\geq a} + z_1\Pi_x + z_2\Pi_y)V\big)$. Boundaries defined by the first three entries (i_1, i_2, i_3) need only $k + l \leq 2$. In Glynn's formula [27] each term is a product of N factors linear in z_1, z_2, so the required coefficients are accumulated inside the sum at the cost of a few permanents per boundary. For distinguishable particles the same events have the generating function $\prod_{k=1}^{N}\big(q_k(a) + z_1|U_{xk}|^2 + z_2|U_{yk}|^2\big)$ with $q_k(a) = \sum_{j\geq a}|U_{jk}|^2$, which costs $O(NM)$ operations. We refer to it as the product formula. Both computations were checked against full enumeration at four photons.

With three-entry boundaries the finest partition has $d = N^3$ bins in the dilute regime, that is $d = M^{3/2}$, with weights between 0.5 and 1.4 times $1/d$; coarser partitions are obtained by merging contiguous bins. Because the weights are not exactly equal we define the MPB as the bin with the largest ratio P_j/w_j. We call $\delta_j - \delta_j^{\mathrm{dist}}$ the interference part of the signal, where δ_j^{dist} is the deviation of the distinguishable-particle bin probability P_j^{dist} from its own Haar mean w_j^{dist}, which we estimate from 1000 Haar-random interferometers. The product formula ranks the bins by $P_j^{\mathrm{dist}}/w_j^{\mathrm{dist}}$.

5.2 The signal survives and is mostly single-particle

With four bins the signal decays as a power law [Fig. 3(a)]. Fitting $N \geq 8$ gives $\max_j|\delta_j| \propto N^{-0.92\pm0.04}$ in the dilute regime and $N^{-0.54\pm0.07}$ in the dense regime with two bins, both consistent with $M^{-1/2}$. At 20 photons and 400 modes the mean signal is 0.022. Lexicographic binning thus satisfies C1.

The signal is almost entirely reproduced by distinguishable particles. The deviations of bosons and of distinguishable particles from their Haar means are correlated with coefficient above 0.95 for every size, regime and bin number we computed, and with four bins the two MPB labels coincide in 91% to 97% of instances. The interference part decays faster than the signal, as $N^{-1.88\pm0.06}$ or

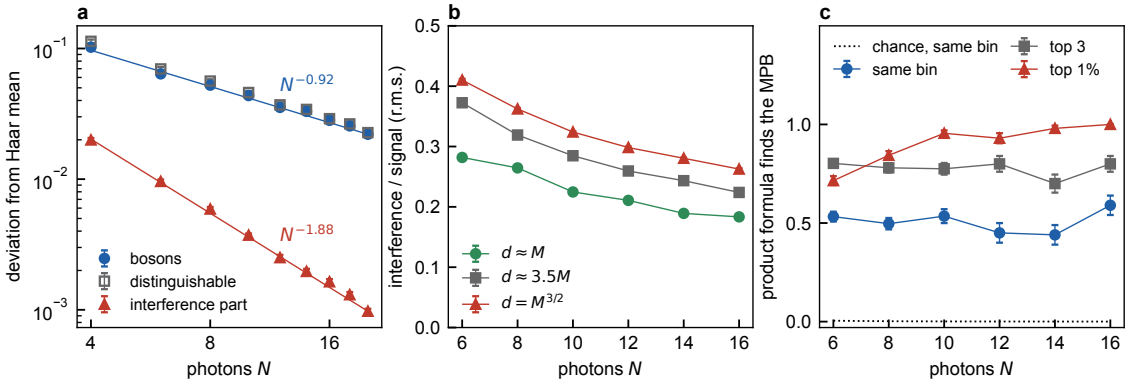


Figure 3. The signal of lexicographic binning is carried by single-particle transition probabilities. Dilute regime, $M = N^2$. (a) Largest deviation of a bin probability from its Haar mean for four bins, for bosons and for distinguishable particles, and the fluctuation of their difference (interference part); lines are power-law fits to $N \geq 8$; 100 to 200 interferometers per size; error bars are the standard error over interferometers, 3% to 5% of each value and smaller than the symbols. (b) Root-mean-square ratio of the interference part to the signal for three bin numbers. (c) Fraction of interferometers for which the distinguishable-particle product formula ranks the true most probable bin (MPB) first, among its top three, or within its top 1%, at $d = M^{3/2}$. Error bars in (b) and (c) are the standard error over interferometers, 100 to 400 per size; in (b) they are smaller than the symbols and exclude the upward bias of 0.001 to 0.003 from the 1000-sample distinguishable-particle Haar mean.

about $1/M$ in the dilute regime, and amounts to 4% of the signal at 20 photons. In the dense regime it decays as $N^{-0.58\pm0.06}$, at the same rate as the signal.

The schemes of Refs. [15, 16] require $d \gtrsim M$, on the argument that interference is then visible in the coarse-grained distribution. Figure 3(b) shows how far this holds. At fixed N the weight of the interference part does grow with the number of bins: at 16 photons its root-mean-square ratio to the signal is 0.18 for $d \approx M$, 0.22 for $d \approx 3.5M$ and 0.26 for $d = M^{3/2} = 16M$. At fixed fineness, however, it decreases with N, from 0.41 at 6 photons to 0.26 at 16 photons for $d = M^{3/2}$ and from 0.28 to 0.18 for $d \approx M$. We find no crossover up to $d = M^{3/2}$.

The practical consequence is in Fig. 3(c). At the finest level the MPB carries 2.7 to 2.9 times its Haar weight, so the signal is large in relative terms. The product formula names the MPB exactly in 44% to 59% of instances, against a chance level between 0.5% and 0.02%, places it among its three highest-ranked bins in 70% to 80%, and for $N \geq 10$ places it within the top 1% of its ranking in at least 93% of instances. Lexicographic binning therefore fails C2: a calculation that involves no interference and takes $O(NM)$ operations per bin reduces the candidates for the function value by two orders of magnitude, and the remaining uncertainty shrinks in relative weight as the system grows.

6 Quadratic coarse-grainings

Reference [21] proposes coarse-graining through a quadratic function $f(\boldsymbol{s}) = \sum_{j,k} J_{jk} s_j s_k$ as a candidate beyond linear statistics, since its characteristic function is a transition amplitude of a circuit with one layer of Kerr interactions. The following consequence of their Theorem 1 delimits the part of this family that is classically estimable.

Proposition 1. *Let $\boldsymbol{\omega}^{(1)}, \ldots, \boldsymbol{\omega}^{(r)}$ be non-negative integer weight vectors with $\boldsymbol{\omega}^{(a)} \cdot \boldsymbol{s} < B$ for all outcomes, and let h be any function of r integers. The bin probabilities of $g(\boldsymbol{s}) = h\big(\boldsymbol{\omega}^{(1)} \cdot \boldsymbol{s}, \ldots, \boldsymbol{\omega}^{(r)} \cdot \boldsymbol{s}\big)$ can be estimated classically to additive error ε in time polynomial in M, B^r and $1/\varepsilon$.*

Proof. Set $\boldsymbol{\Omega} = \sum_a B^{a-1} \boldsymbol{\omega}^{(a)}$. The base-$B$ digits of $\boldsymbol{\Omega} \cdot \boldsymbol{s} < B^r$ are the values $\boldsymbol{\omega}^{(a)} \cdot \boldsymbol{s}$. By Theorem 1 of Ref. [21] the cumulative distribution $S(x) = \Pr[\boldsymbol{\Omega} \cdot \boldsymbol{s} \leq x]$ is estimated to additive error ε' in time polynomial in M, $r \log B$ and $1/\varepsilon'$. Each of the at most B^r cell probabilities is $S(x) - S(x-1)$, with error $2\varepsilon'$, and a bin probability is a sum of cells; $\varepsilon' = \varepsilon/(2B^r)$ suffices. □

For $r = O(1)$ and B polynomial the estimate is efficient. This covers $f = \sum_a \lambda_a (\boldsymbol{\omega}^{(a)} \cdot \boldsymbol{s})^2$ with a constant number of polynomially bounded integer vectors, and every J that is constant on a constant number of blocks of modes, for which f depends only on the block photon numbers. The method stops working when the number of blocks grows with N; the extreme case is the collision statistic $\sum_j s_j^2$.

Exact enumeration up to 6 photons in the dilute regime and 8 in the dense regime (Appendix A) indicates that the high-rank members of the family are constrained by C1 and C2 rather than by C3. In the dilute regime the signal of the collision statistic decays as $N^{-2.38\pm0.03}$, somewhat faster than $1/M$, and a predictor matched to its first three moments accounts for 88% of it at 5 photons and 90% at 6. In the dense regime 20% to 24% of the signal remains after three moments at 8 photons, but the same residual is found for a block statistic that the proposition makes classically estimable, so the residual does not discriminate between easy and potentially hard statistics. The dense regime is undecided.

7 Discussion

The three results have a common structure. A coarse-graining must break the symmetry of the Haar ensemble to carry a signal, and the ways of doing so that we examined either break it by a structure that single-particle probabilities already see (an ordering of the modes) or do not break it at all (a generic permutation). The lemma makes the second statement general: any coarse-graining that is uncorrelated with the output distribution is flat up to the collision probability. The search for a hard coarse-graining is thus a search for a public function g that is correlated with p_U for typical U without being correlated with its low-order marginals.

Where could such a function lie? The proposition excludes statistics that factor through a constant number of linear ones. In the dilute regime the signal of the collision statistic decays at least as fast as $1/M$, and low moments account for about 90% of the signals of the quadratic statistics we examined at 6 photons. What remains is the dense regime, $M = O(N)$, where sampling hardness has recently been put on a firmer footing [28], where collisions are frequent, and where we find that the interference part of the lexicographic signal does not decay faster than the signal itself. High-rank nonlinear statistics in that regime are the candidates that our tests leave open.

The conditions also bear on the relation between verification and cryptography. Chia *et al.* argue that a sampling-based quantum advantage is either efficiently verifiable or yields cryptographic primitives [29]. Coarse-grained schemes aim at a middle ground in which bin probabilities are both checkable and hard. Our results indicate that, for the binnings examined here, checkability has been bought with classical accessibility: the bin probabilities that samples can resolve are those that low-order correlations determine.

For experiments, the two decay rates translate into sample budgets. Resolving a deviation δ of a bin probability takes of order $1/\delta^2$ samples. For four lexicographic bins at 20 photons and 400 modes this is about 2×10^3 samples for the signal and about 10^6 for its interference part, and the ratio grows in proportion to M in the dilute regime. A validation test based on such bins is thus sensitive mainly to the single-particle transfer matrix and to the interferometer-independent offset between the Haar means of bosons and of distinguishable particles; the interference contribution specific to an interferometer becomes visible only when the sample size is scaled with M^2. Photon loss and partial distinguishability reduce the interference part further and leave linear statistics classically estimable [21, 23], so they tighten the conditions rather than relax them.

Our numerical statements have limits that should be stated plainly. The bound on the collision probability is a Haar average taken from Ref. [24] and applies to typical, not to all, interferometers; we checked it directly up to 8 photons. The exponents quoted for lexicographic binning are fits over 8 to 20 photons and are not derived analytically. The bin numbers reach $M^{3/2}$; for $d \geq M^2$ boundaries on a fourth entry are needed, and although the trend with N at fixed fineness is downward at every level we computed, a crossover at much finer binning is not excluded. C2 is tested with one predictor, the product formula, and its failure is a statement about that predictor.

8 Conclusion

We formulated three conditions under which a coarse-grained boson sampling distribution can support a cryptographic function, namely a readable signal, a signal beyond low-order correlations and the absence of a classical estimator, and we tested published schemes against them by a general lemma and by exact computation up to 20 photons and 4096 bins. To our knowledge this is the first quantitative separation of the single-particle and interference contributions to coarse-grained bin probabilities at these sizes.

Two trends emerge. First, permuting the outcome space removes the signal: the deviation of every bin probability from the uniform value is bounded by the collision probability and falls as $|Y|^{-1/2}$, whereas Ref. [20] relies on this step to require a quantum device. Second, ordering the outcome space lexicographically keeps a signal of order $M^{-1/2}$, but its interference part is of order $1/M$ for coarse bins in the dilute regime and stays below 42% of the signal for every bin number up

to $M^{3/2}$; Ref. [21] established that these bin probabilities are classically estimable, and we add that most of the signal does not involve interference in the first place.

These findings place cryptographic coarse-graining in the sequence that vibronic spectra and graph problems have followed, in which a proposed use of boson sampling was matched classically once the quantity actually read out was identified [12, 13]. If new schemes are proposed without checking the three conditions, they will keep describing functions for which the photonic device is either untested or unnecessary. For photonic quantum information processing, the practical effect is a test that can be run before an experiment is built: compute the Haar mean, the distinguishable-particle prediction and the low-order moments of the proposed coarse-graining, and ask what is left. The remaining open case, high-rank nonlinear statistics in the dense regime, calls for a quantification of high-order interference that exact enumeration cannot reach.

Data availability

The scripts that produce every number and figure in this article, together with the numerical results, are openly available at https://github.com/deeptell-inc/cgbs-constraints-repro.

A Quadratic statistics: numerical results

We computed the exact distribution $P_k = \Pr[f(\boldsymbol{s}) = k]$ by full enumeration for five statistics: the number of photons in the first half of the modes (linear, for reference); $\sum_b N_b^2$ for $r = 4$ blocks of modes with block photon numbers N_b; the collision statistic $\sum_j s_j^2$; the neighbour statistic $\sum_j s_j s_{j+1}$ with periodic boundary; and $\sum_{j<k} J_{jk} s_j s_k$ with $J_{jk} \in \{0, 1\}$ drawn uniformly and independently for each size. The signal is the total variation distance between P and its Haar mean w, which follows from Eq. (2).

The qth moment of a quadratic f is a photon-number correlator of order $2q$, and correlators of fixed order are classically computable in polynomial time: the normally ordered correlator of output modes $j_1, \ldots, j_k$ equals $\sum_S |\mathrm{Per}\, U_{J,S}|^2$, where $U_{J,S}$ has rows $j_1, \ldots, j_k$ and columns S and the sum runs over the $\binom{N}{k}$ sets S of k input photons. As a predictor that uses only such quantities we take the exponential tilt of the Haar mean,

$$\hat{P}_k^{(p)} \propto w_k \exp\bigl(a_1 k + \cdots + a_p k^p\bigr), \tag{6}$$

with coefficients chosen so that the first p moments equal those of P, and we report the residual $\mathrm{TV}(P, \hat{P}^{(p)})$ divided by the signal. The moments were taken from the enumerated distribution.

Figure 4 shows the results. In the dilute regime the signal of the collision statistic falls from 0.044 at 3 photons to 0.0085 at 6, as $N^{-2.38\pm0.03}$, while the linear statistic falls as $N^{-1.15\pm0.02}$; the quoted uncertainties are those of the fit. The residual after three moments is 0.12, 0.12 and 0.10 of the signal at 4, 5 and 6 photons for the collision statistic and 0.30, 0.20 and 0.09 for the random-graph statistic. In the dense regime the signals of the collision and neighbour statistics decay as $N^{-1.88\pm0.09}$ and $N^{-1.74\pm0.07}$ for $N = 4$ to 8. At 8 photons the residual after three moments is 0.22 for the collision statistic, 0.20 for the neighbour statistic, 0.21 for the random graph, and 0.24 for the block statistic with $r = 4$, which is classically estimable by the proposition of Sec. 6; each value has a standard error between 0.01 and 0.02. From 5 to 8 photons the residual decreases for every statistic, the block statistic included. Each size uses 100 to 800 interferometers; the range of N is too narrow to extrapolate.

References

[1] Zhong H S, Wang H, Deng Y H, Chen M C, Peng L C, Luo Y H *et al.* 2020 *Science* **370** 1460–1463

[2] Madsen L S, Laudenbach F, Askarani M F, Rortais F, Vincent T, Bulmer J F F *et al.* 2022 *Nature* **606** 75–81

[3] Aaronson S and Arkhipov A 2011 The computational complexity of linear optics *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* pp 333–342

[4] Hamilton C S, Kruse R, Sansoni L, Barkhofen S, Silberhorn C and Jex I 2017 *Phys. Rev. Lett.* **119** 170501

[5] Hangleiter D and Eisert J 2023 *Rev. Mod. Phys.* **95** 035001

[6] Huh J, Guerreschi G G, Peropadre B, McClean J R and Aspuru-Guzik A 2015 *Nat. Photonics* **9** 615–620

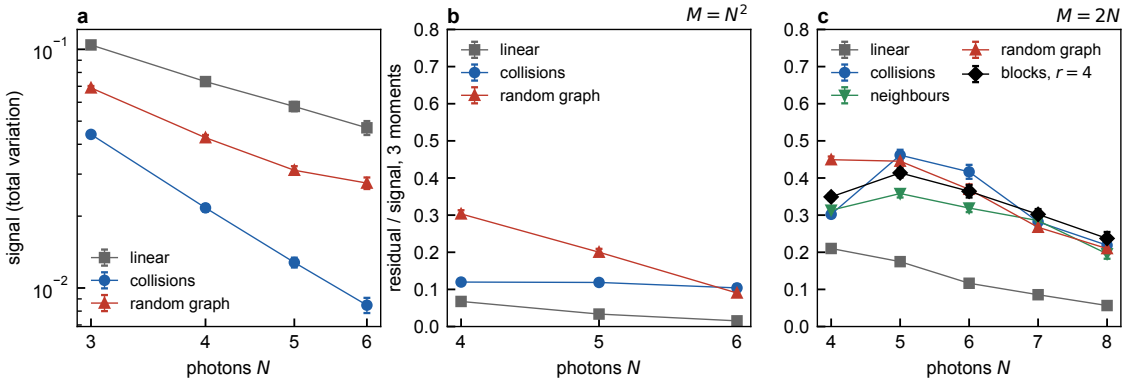


Figure 4. In the dilute regime the signal of quadratic statistics is weak and largely accounted for by low moments. (a) Total variation distance between the distribution of each statistic and its Haar mean, $M = N^2$; error bars are the standard error over 100 to 800 interferometers. (b), (c) Residual distance after matching three moments, divided by the signal, in the dilute and dense regimes; error bars are the standard error of the ratio. The block statistic with $r = 4$ is classically estimable and shows the same residual as the high-rank statistics.

[7] Arrazola J M and Bromley T R 2018 *Phys. Rev. Lett.* **121** 030503

[8] Wang S T and Duan L M 2016 Certification of boson sampling devices with coarse-grained measurements (*Preprint* arXiv:1601.02627)

[9] Seron B, Novo L, Arkhipov A and Cerf N J 2024 *Quantum* **8** 1479

[10] Bressanini G, Seron B, Novo L, Cerf N J and Kim M S 2025 *Phys. Rev. A* **112** 012610

[11] Correa Anguita M, Camillini A, Marzban S, Robbio M, Seron B, Novo L and Renema J J 2025 *Quantum Sci. Technol.* **10** 035062

[12] Oh C, Lim Y, Wong Y, Fefferman B and Jiang L 2024 *Nat. Phys.* **20** 225–231

[13] Oh C, Fefferman B, Jiang L and Quesada N 2024 *PRX Quantum* **5** 020341

[14] Nikolopoulos G M and Brougham T 2016 *Phys. Rev. A* **94** 012315

[15] Nikolopoulos G M 2019 *Quantum Inf. Process.* **18** 259

[16] Wang X, Zhou W, Fu Y, Gao J, Lu Y, Chang Y, Qiao L, Ren R, Jiang Z, Jiao Z, Nikolopoulos G M and Jin X 2023 *Phys. Rev. Lett.* **130** 060802

[17] Nikolopoulos G M 2023 *Phys. Scr.* **98** 014001

[18] Shi J, Lu Y, Feng Y, Huang D, Lou X, Li Q and Shi R 2022 *Quantum Inf. Process.* **21** 73

[19] Hatanaka T, Fushio R, Watanabe M, Munro W J, Ikeda T N and Sugiura S 2024 A quantum-resistant photonic hash function (*Preprint* arXiv:2409.19932)

[20] Singh D, Muraleedharan G, Fu B, Cheng C M, Roussy Newton N, Rohde P P and Brennen G K 2025 *Quantum Sci. Technol.* **10** 025020

[21] Seron B, Thomas H, Araujo E, Arkhipov A, Oh C and Novo L 2026 Efficient classical algorithm for estimating linear statistics of boson sampling (*Preprint* arXiv:2609.35447)

[22] Walschaers M, Kuipers J, Urbina J D, Mayer K, Tichy M C, Richter K and Buchleitner A 2016 *New J. Phys.* **18** 032001

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

[24] Mhiri H, Thomas H, Monbroussou L, Chabaud U, Holmes Z and Kashefi E 2026 Boson sampling beyond the dilute regime: second moments and anti-concentration (*Preprint* arXiv:2604.14323)

[25] Cardin Y and Quesada N 2024 *Quantum* **8** 1521

[26] Gupt B, Izaac J and Quesada N 2019 *J. Open Source Softw.* **4** 1705

[27] Glynn D G 2010 *Eur. J. Combin.* **31** 1887–1891

[28] Bouland A, Brod D, Datta I, Fefferman B, Grier D, Hernández F and Oszmaniec M 2026 *Phys. Rev. X* **16** 021059

[29] Chia N H, Fu H, Song F and Yao P 2025 *ACM Trans. Quantum Comput.* **7** 1–20