arXiv:2405.01786v5 [quant-ph] 11 Mar 2026

On computational complexity and average-case hardness of shallow-depth Boson Sampling

Byeongseon Go[1], Changhun Oh[2], and Hyunseok Jeong[1]

[1]NextQuantum Innovation Research Center, Department of Physics and Astronomy, Seoul National University, Seoul 08826, Republic of Korea

[2]Department of Physics, Korea Advanced Institute of Science and Technology, Daejeon 34141, Republic of Korea

Boson Sampling, a computational task believed to be classically hard to simulate, is expected to hold promise for demonstrating quantum computational advantage using near-term quantum devices. However, noise in experimental implementations poses a significant challenge, potentially rendering Boson Sampling classically simulable and compromising its classical intractability. Numerous studies have proposed classical algorithms that can efficiently simulate Boson Sampling under various noise models, particularly as noise rates increase with circuit depth. To address this challenge, we investigate the viability of achieving quantum computational advantage through Boson Sampling implemented with shallow-depth linear optical circuits. In particular, as the average-case hardness of estimating output probabilities of Boson Sampling is a crucial ingredient in demonstrating its classical intractability, we make progress on establishing the average-case hardness of Boson Sampling confined to logarithmic-depth regimes. We also obtain the average-case hardness for logarithmic-depth Fock-state Boson Sampling subject to lossy environments and for the logarithmic-depth Gaussian Boson Sampling. By providing complexity-theoretical backgrounds for the classical simulation hardness of logarithmic-depth Boson Sampling, we expect that our findings will mark a crucial step towards a more noise-tolerant demonstration of quantum advantage with shallow-depth Boson Sampling.

Contents

1 Introduction **2**
- 1.1 Backgrounds and motivation . . . 2
- 1.2 Our results: Average-case hardness of shallow-depth Boson Sampling . . . 3
- 1.3 Paper organization . . . 5

2 Our settings **5**
- 2.1 Notations . . . 5
- 2.2 Definitions for random circuit and outcome ensemble . . . 8

3 Worst-case hardness of output probability estimation **9**

Changhun Oh: changhun0218@gmail.com
Hyunseok Jeong: h.jeong37@gmail.com

4 Average-case hardness of output probability estimation **10**

5 Classical simulation hardness of shallow-depth boson sampling **13**

6 Average-case hardness for shallow-depth Gaussian Boson Sampling **15**

7 Hardness analyses in lossy environments **18**

8 Concluding remarks **19**

A Previous foundations on average-case hardness **26**

B The bosonic birthday paradox for the shallow average-case circuits **27**

C Proof of Theorem 3 (worst-case hardness) **29**

D Proof of Lemma 7: Encoding constant-depth worst-case circuit in Lemma 5 to the circuit architecture $\mathcal{BB}^*$ **31**

D.1 Mapping worst-case circuit in Lemma 5 to a grid linear optical circuit . . . 32

D.2 Encoding a grid linear optical circuit to the butterfly circuit architecture $\mathcal{B}$ 33

D.3 Summary: Encoding the worst-case circuit in Lemma 5 to $\mathcal{BB}^*$ 37

E About Cayley transform **38**

F Proof of Theorem 4 (average-case hardness) **40**

G Proof of Lemma 3 **43**

H On average-case hardness including collision outcomes **44**

I Numerical simulation for the collision patterns of Boson Sampling for the local random circuit over $(\mathcal{BB}^*)^q$ **47**

1 Introduction

1.1 Backgrounds and motivation

Boson Sampling is a computational task that is complexity-theoretically proven to be hard to classically simulate under plausible assumptions [1–3]. Accordingly, Boson Sampling has gathered significant attention, as it would possibly play a key role in the experimental demonstration of quantum computational advantage using near-term quantum devices. However, the implementation of Boson Sampling in experimental settings with near-term quantum devices is inevitably subject to various sources of noise [4–7]. The problem is that those noises would possibly rule out the classical intractability of Boson Sampling, and thus potentially hinder the experimental demonstration of quantum advantage with Boson Sampling. Indeed, both for finite-size near-term experiments and asymptotic limits as system size scales, numerous studies [8–26] have proposed efficient classical simulation algorithms of Boson Sampling under various noise models, such as photon loss, partial distinguishability, random Gaussian noise, etc. Their results indicate that as the noise rate of Boson Sampling increases, it eventually renders such a noisy sampler classically simulable. Moreover, as the noise is typically accumulated with each circuit depth, the quantum signal for classical intractability exhibits exponential decay with increasing circuit

depth. Hence, circuits with polynomially increasing depth with system size would suffer from significantly enlarged noise rates, posing substantial challenges to achieving quantum advantage in such settings.

A viable alternative to preclude the classical simulability due to the inevitable noise is to consider Boson Sampling with *shallow-depth* linear optical circuits, where the noise rate can be highly reduced. Specifically, among the shallow-depth regime, our primary focus is on investigating the simulation hardness for *logarithmic* depth circuits; the intuition behind investigating logarithmic depth circuits lies in the potential to offer a "sweet-spot" regime for the hardness of Boson Sampling. Namely, this depth regime may avoid significant increases in noise rates to prevent classical simulability, while still being sufficiently large to generate quantum correlations and uphold classical intractability. Despite such intuitive understanding, the hardness argument of Boson Sampling in this shallow-depth regime, particularly from a complexity-theoretical perspective, has been less studied so far and thus remains widely open. Hence, our goal is to establish the *complexity-theoretical foundations* of the classical hardness of shallow-depth Boson Sampling, to suppress the classical simulability by noise in a rigorous manner and obtain a more noise-tolerant demonstration of quantum advantage with Boson Sampling.

In this work, we investigate the classical simulation hardness of Boson Sampling in shallow linear optical circuits. Specifically, as the average-case #P-hardness of estimating output probabilities of Boson Sampling is a crucial ingredient to demonstrate the classical intractability of Boson Sampling, we make progress on establishing the average-case #P-hardness confined in shallow-depth regimes. Similarly, we obtain the average-case hardness result in the shallow-depth regime for the Gaussian Boson Sampling scheme. Finally, since noise is our main motivation for investigating shallow-depth Boson Sampling, we generalize our average-case hardness result to noisy Boson Sampling subject to a photon loss channel.

To avoid confusion, we note that the allowed imprecision level of our average-case #P-hardness result is not sufficient to fully demonstrate the classical intractability of Boson Sampling in shallow-depth regimes. However, to the best of our knowledge, the complexity-theoretical analysis on the average-case hardness of shallow-depth Boson Sampling has not yet been investigated. Hence, we believe that our hardness result in shallow-depth regimes will provide a first step toward a stronger hardness result and, ultimately, toward the full demonstration of the classical intractability of shallow-depth Boson Sampling.

1.2 Our results: Average-case hardness of shallow-depth Boson Sampling

We set our goal as proving the hardness of classical simulation of Boson Sampling in the shallow-depth regime, specifically for *approximate* simulation within total variation distance error. Two key ingredients for the current hardness proof of the approximate simulation of Boson Sampling are (i) average-case #P-hardness of output probability approximation up to sufficiently large additive imprecision ϵ, and (ii) hiding property. Informally, average-case hardness means that approximating the output probability of Boson Sampling with high probability over randomly chosen circuits (i.e., on average over circuits) is #P-hard. Here, by choosing random circuit ensembles that have the hiding property (i.e., symmetry over outcomes), one can reduce the average-case instances for the hardness from circuit instances to outcome instances, which is a crucial step to prove the hardness of approximate simulation within total variation distance error (See Appendix A for more details).

Most of the current theoretical foundations of the average-case hardness of Boson Sampling rely on global Haar random unitary circuits [1–3, 27–30], as they almost satisfy the two conditions described above. Namely, the outcome instances can be effectively hidden

by global Haar random unitaries, and approximating the output probability within sufficiently large ϵ on average over global Haar random circuit instances is #P-hard under some conjectures. However, the problem for our case is that implementing a global Haar random unitary requires at least polynomially large circuit depth (e.g., see [31, 32]), and thus is not implementable in sub-polynomial circuit depths. This poses a challenge in proving the average-case hardness of Boson Sampling in the shallow-depth regime—specifically, the lack of the hiding property. Due to the absence of hiding in the shallow-depth regime, the random outcome instances of Boson Sampling cannot be hidden by random circuit instances. Accordingly, even if we establish average-case hardness over randomly chosen circuits for a fixed outcome, it does not directly imply classical simulation hardness via Stockmeyer's reduction, which requires average-case hardness over randomly chosen outcomes [1].

Another problem is that there already exist efficient classical algorithms that can approximately simulate shallow-depth Boson Sampling in certain circumstances, which directly rule out the classical simulation hardness in the shallow-depth regime for those cases. While exact simulation of Boson Sampling is classically hard even for constant-depth circuits [33], approximate simulation becomes easy for 1-dimensional local log-depth circuits [34, 35] and also for more general dimension local circuits under some constraints [36–38]. According to their results, if we use circuits composed of only geometrically local gates, a polynomial circuit depth is generally required for a sufficiently large correlation to achieve approximate simulation hardness. Those results indicate that we cannot expect the hardness results in the most general case of shallow-depth circuits composed of local gates only.

To deal with those problems, we take the following approach: First, we consider shallow linear optical circuit architectures that fully employ geometrically *non-local* gates. In fact, the implementation of non-local gates is promising for near-term experimental settings; for example, experiments of linear optical systems based on trapped ions [39, 40] and photonic architecture [6] implemented long-range interactions. Because of the absence of hiding, we derive the average-case hardness over *both* randomly chosen outcomes and randomly chosen circuits over the shallow circuit architectures, by establishing a worst-to-average-case reduction for both outcome and circuit instances. In short, this reduction process can be done in two steps: (i) from a given fixed outcome to a randomly chosen collision-free outcome, and (ii) from a worst-case circuit to a randomly chosen circuit over a local random circuit distribution.

For random circuit ensembles, we employ a local random circuit ensemble in the shallow circuit architecture, inspired by the hardness results of random circuit sampling [28, 41–44]. Here, the local random distribution in this context means that each gate composing the circuit is independently chosen Haar random gate; we note that recent experimental setups of Boson Sampling [4–7] follow a similar circuit distribution, but with geometrically local architectures. Additionally, due to the absence of symmetry, the local random circuit ensemble itself does not guarantee the bosonic birthday paradox, which is crucial for the hardness of Boson Sampling when using only collision-free outcomes [1]. Hence, we also include a random permutation circuit at the input of our random circuit to guarantee the bosonic birthday paradox for our random circuit ensemble (see Appendix B). In fact, this random permutation circuit can be implemented within the shallow depth regime (see, e.g., Lemma 1), and can be experimentally realized by preparing a random input configuration instead.

To sum up, we show the average-case hardness over outcomes and circuit instances for shallow circuit architectures composed of geometrically non-local gates and employing the local random circuit ensemble. We informally present here our average-case hardness

result of Boson Sampling in the logarithmic depth regime, for photon number N and mode number $M \propto N^\gamma$ with $\gamma \geq 2$.

Theorem 1 (Informal)**.** *Consider a Boson Sampling problem with N photons over $M = \Omega(N^\gamma)$ mode, where $\gamma \geq 2$. There exists an $O(\log N)$-depth linear optical circuit architecture $\mathcal{A}$, consisting of $O(N^\gamma \log N)$ number of geometrically non-local gates, such that estimating the output probability to within additive error $2^{-O(N^{\gamma+1}(\log N)^2)}$, with high probability over the choice of outcomes and circuits in $\mathcal{A}$, is #P-hard under $\mathrm{BPP}^{\mathrm{NP}}$ reduction.*

Also, since our average-case hardness result regards both the random outcomes and random circuits, it is not straightforward to show the classical simulation hardness of shallow-depth Boson Sampling as in the original Boson Sampling proposal [1]. Accordingly, we show how our average-case hardness result over both the randomly chosen outcomes and circuits leads to the classical simulation hardness argument. This implies that improving the additive imprecision for our average-case hardness result is the only remaining problem for the fully theoretically guaranteed classical intractability of shallow-depth Boson Sampling.

Theorem 2 (Informal)**.** *For $M \propto N^\gamma$ with $\gamma \geq 2$, if the allowed additive imprecision for the problem in Theorem 1 to be #P-hard is improved to $\epsilon = 2^{-(\gamma-1)N\log N - O(N)}$, the approximate Boson Sampling for the shallow-depth circuit in Theorem 1, up to constant total variation distance, is classically hard to simulate.*

1.3 Paper organization

Our paper is organized as follows, which is outlined in Fig. 1. We first define in Sec. 2 a shallow-depth circuit architecture $(\mathcal{BB}^*)^q$ composed of non-local gates, which we will use throughout our results. Next, in Sec. 3, we prove the worst-case #P-hardness of approximating output probability $p_{\boldsymbol{s}}(C)$ of a fixed outcome $\boldsymbol{s}$ of Boson Sampling, for any circuit C in the shallow circuit architecture previously defined. In Sec. 4, we prove the average-case #P-hardness of approximating output probability $p_{\boldsymbol{s}}(U)$ for randomly chosen outcome $\boldsymbol{s}$ and randomly chosen circuit U in the shallow circuit architecture, by establishing worst-to-average-case reduction. We prove in Sec. 5 how our average-case hardness results over both the random outcomes and random circuits lead to the classical simulation hardness. We also extend our average-case hardness result to the Gaussian Boson Sampling scheme in Sec. 6, and to the lossy Boson Sampling subject to photon loss channels in Sec. 7. In Sec. 8, we conclude with several remarks.

2 Our settings

2.1 Notations

Let us define the total mode number as M, where we set M as a power of 2 for simplicity. We set the output photon number N polynomially related to M as $M = c_0 N^\gamma$ for a constant c_0 and $\gamma \geq 2$, to ensure that collision-free outcomes dominate the probability weight over all possible N-photon outcomes over M modes (i.e., by bosonic birthday paradox; see Refs. [1, 45] and Appendix B). We use the notation $\boldsymbol{s}$ as an M-dimensional output configuration vector representing N-photon outcome over M modes, such that each element s_i ($i \in [M]$) of $\boldsymbol{s}$ denotes the occupied number of photons in the ith mode. Namely, $\boldsymbol{s}$ can be expressed as $\boldsymbol{s} = (s_1, \ldots, s_M)$ where each $s_i \in [N]$ with the constraint $\sum_{i=1}^M s_i = N$. Also, let $\boldsymbol{t}$ be an M-dimensional input configuration vector corresponding to an N-photon

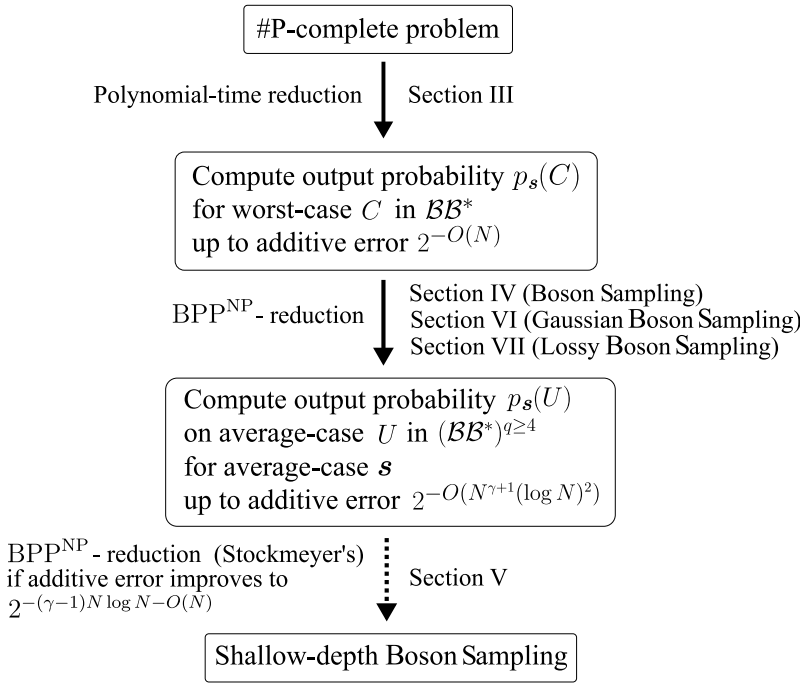

Figure 1: Outlines of our result

input state over M modes, such that $\boldsymbol{t} = (t_1, \ldots, t_M)$ where now each $t_i \in \{0, 1\}$ with $\sum_{i=1}^{M} t_i = N$.

Hereafter, we use the terminology Boson Sampling representing Boson Sampling with Fock-state inputs as in the original Boson Sampling proposal [1]; we will explicitly use the terminology Gaussian Boson Sampling for Boson Sampling with Gaussian state inputs [2]. We now define $p_{\boldsymbol{s}}(C)$ as the output probability of Boson Sampling with a linear optical circuit (unitary) matrix C to obtain an outcome $\boldsymbol{s}$ from a fixed input configuration $\boldsymbol{t}$. More precisely, $p_{\boldsymbol{s}}(C)$ can be represented as [1]

$$p_{\boldsymbol{s}}(C) = \frac{1}{\prod_{i=1}^{M} s_i!} |\mathrm{Per}(C_{\boldsymbol{s},\boldsymbol{t}})|^2, \tag{1}$$

where $C_{\boldsymbol{s},\boldsymbol{t}}$ is an N by N matrix obtained by taking s_i copies of the ith row and t_j copies of the jth column of the matrix C. Here, we omit the $\boldsymbol{t}$ dependence in the output probability $p_{\boldsymbol{s}}(C)$ for simplicity because we will fix the input configuration throughout the paper.

We note that an M-mode linear optical circuit can be represented by an M by M unitary matrix in $\mathrm{U}(M)$, which unitarily transforms mode operators. Specifically, we can represent a single-mode gate (i.e., a phase shifter) as a $\mathrm{U}(1)$ matrix to the mode, and a two-mode gate (i.e., a beam splitter) as a $\mathrm{U}(2)$ matrix along the modes. Also, the parallel application of gates can be represented as a unitary matrix with a block matrix form, and the serial application of gates can be represented as matrix multiplication of the unitary matrices. Accordingly to this correspondence, throughout this work, we will interchangeably use the terminology '(linear optical) circuit' and '(unitary) matrix'.

We first define the linear optical circuit architecture for a more rigorous analysis of the hardness proof.

Definition 1 (Linear optical circuit architecture)**.** *The linear optical circuit architecture $\mathcal{A}$ is a linear optical circuit with fixed type (i.e., single- or two-mode) and fixed location of gates, where the coefficients of each gate are not specified. If the coefficients of each gate*

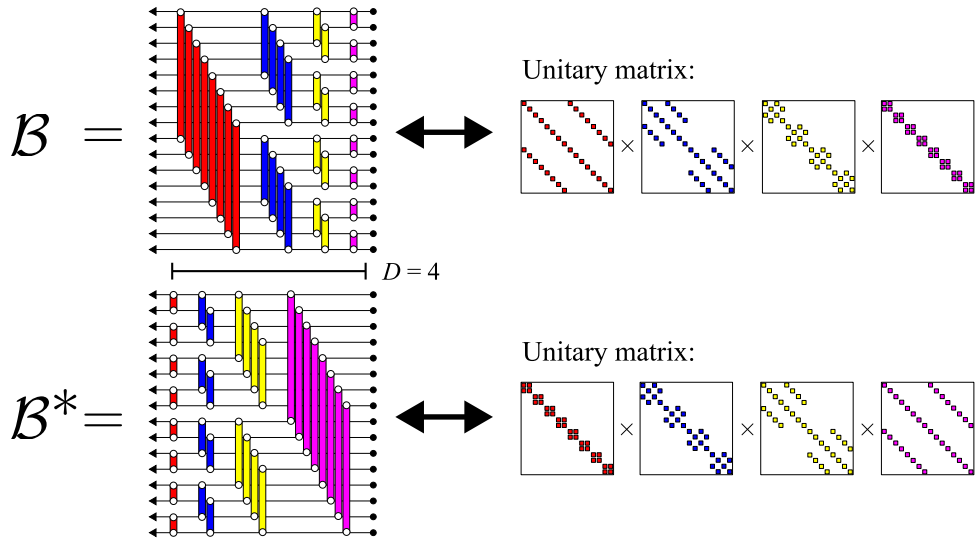

Figure 2: Schematics of the butterfly circuit architectures in Definition 2 and their unitary matrix form, for mode number $M = 2^4 = 16$.

are specified with unitary matrices (in $\mathrm{U}(1)$ *or* $\mathrm{U}(2)$*), then the circuit and the corresponding unitary matrix are specified.*

Given the definition of the circuit architecture $\mathcal{A}$, we denote m as the number of gates in the circuit architecture. Also, we denote D as the depth of the circuit architecture, where we define a unit depth as a single step of *parallel* application of gates, such that any serial application of gates is not allowed in the unit depth. Here, we emphasize that single-mode gates represented by $\mathrm{U}(1)$ matrices do not contribute to the depth count since they can be absorbed by the nearest two-mode gates represented by $\mathrm{U}(2)$ matrices.

In this work, we consider the shallow-depth circuit architecture, specifically in logarithmic depth regime $D = O(\log N)$. As we only consider the binary logarithm throughout this work, we will omit the base of the logarithm such that $\log(\cdot)$ implies $\log_2(\cdot)$ hereafter. We now define the shallow linear optical circuit architecture of circuit depth $D = \log M$ with the number of gates $m = \frac{M}{2}\log M$, using the convention used in [46, 47].

Definition 2 (Butterfly circuit architecture)**.** *We define butterfly circuit architecture* $\mathcal{B}$ *as follows: for each layer* $L = 1, 2, \ldots, D = \log M$ *of the circuit architecture, allocate two-mode gate between mode number* $2^L(j-1)+k$ *and* $2^L(j-1)+k+2^{L-1}$*, for all* $j = 1, 2, \ldots, 2^{D-L}$ *and* $k = 1, 2, \ldots, 2^{L-1}$*. Also, we define inverse butterfly circuit architecture* $\mathcal{B}^*$ *as a butterfly circuit architecture with the inverse sequence of gate application along the depth.*

We emphasize that in Definition 2, gates are applied in parallel for all $j = 1, 2, \ldots, 2^{D-L}$ and $k = 1, 2, \ldots, 2^{L-1}$. Hence, by definition of the circuit depth, the total depth for $\mathcal{B}^*$ (and thus also $\mathcal{B}^*$) is given by $D = \log M$. We illustrate in Fig. 2 the circuit architecture $\mathcal{B}$ and $\mathcal{B}^*$ for $2^4 = 16$ modes (such that the circuit depth is $D = 4$ each), and the form of their corresponding unitary matrix.

Next, we define the Kaleidoscope circuit architecture proposed in [46], using the butterfly circuit architecture defined above.

Definition 3 (Kaleidoscope circuit architecture)**.** *We define Kaleidoscope circuit architecture* $\mathcal{B}\mathcal{B}^*$ *as a serial application of* $\mathcal{B}$ *over* $\mathcal{B}^*$ *in Definition 2. We define q-Kaleidoscope*

circuit architecture $(\mathcal{BB}^*)^q$ *as a repeat of the Kaleidoscope circuit architecture, for a repetition number* $q \in \mathbb{N}$.

Here, the circuit depth of the q-Kaleidoscope architecture is $D = 2q \log M$, which is indeed a logarithmic depth $D = O(\log N)$ for $q = O(1)$. Also, the number of gates for the q-Kaleidoscope architecture is $m = qM \log M = O(N^\gamma \log N)$ for $q = O(1)$. Note that, because all the gates are geometrically non-local except for the gates in the first and the last depth for each $\mathcal{BB}^*$, the number of non-local gates also scales $O(N^\gamma \log N)$ for $q = O(1)$. Throughout this work, we will focus on the q-Kaleidoscope circuit architecture with $q = O(1)$ to demonstrate the hardness results for shallow-depth circuits. The motivation for employing this linear optical circuit architecture is that it enjoys a useful property that is crucial for our analysis, which is that for M a power of 2, any M-mode permutation circuit can be implemented within $\mathcal{BB}^*$:

Lemma 1 (Dao et al [46])**.** *Let* $\boldsymbol{P}$ *be an arbitrary* $M \times M$ *permutation matrix with* M *a power of 2. Then* $\boldsymbol{P}$ *can be efficiently implemented in* $\mathcal{BB}^*$ *using* $M \log M$*-number of two-mode permutation gates, i.e.,* $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ *and* $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$.

2.2 Definitions for random circuit and outcome ensemble

In this work, we show the average-case hardness over both randomly chosen outcomes and randomly chosen circuits by establishing a worst-to-average-case reduction for both outcomes and circuit instances. Here, for the average-case instances, it is crucial to choose a proper circuit ensemble and outcome ensemble to successfully establish the worst-to-average-case reduction.

For the random circuit instances, we consider the *product* of two random circuit ensembles. The first is the local random circuit ensemble, which is the circuit ensemble with each gate being independently drawn from the local Haar measure.

Definition 4 (Local random circuit ensemble)**.** *Let* $\mathcal{A}$ *be the circuit architecture with* m *number of gates. We define* $\mathcal{H}_\mathcal{A}$ *as the distribution over circuits with architecture* $\mathcal{A}$*, whose gates are independently distributed local Haar random matrices* $\{H_i\}_{i=1}^m$.

Indeed, this circuit ensemble is similar to the random circuit ensemble employed by the previous random circuit sampling proposals [28, 41–44] which employ gate-wise Haar random unitary circuits. Next, we introduce our second random circuit ensemble, which is the random permutation circuit ensemble over the fixed shallow-depth circuit architecture $\mathcal{BB}^*$.

Definition 5 (Random permutation circuit ensemble)**.** *We define* $\mathcal{P}$ *as the uniform distribution over* M*-mode permutation circuits over the fixed Kaleidoscope circuit architecture* $\mathcal{BB}^*$ *composed of* $M \log M$ *number of two-mode permutation gates* $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ *and* $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$. *Each* $\boldsymbol{P} \sim \mathcal{P}$ *represents an* M *by* M *permutation matrix, uniformly chosen from all possible permutation matrices.*

Note that this random permutation circuit ensemble can be implemented in the shallow-depth circuit architecture $\mathcal{BB}^*$ by Lemma 1. Before proceeding, we briefly remark on the motivation for including the permutation circuit ensemble in our random circuit construction. Specifically, for any M-mode circuit architecture $\mathcal{A}$, including $\mathcal{A} = (\mathcal{BB}^*)^q$ considered in this work, a circuit drawn from the product circuit ensemble $\mathcal{H}_\mathcal{A}\mathcal{P}$ guarantees the *bosonic*

birthday paradox [1,45]. The bosonic birthday paradox ensures that the probability of obtaining only collision outcomes is sufficiently large given that the photon number N is sufficiently smaller than the mode number M, as stated in the following lemma.

Lemma 2 (Bosonic birthday paradox for the product circuit ensemble)**.** *Let $M \geq 2N^2$, and let $B_{M,N}$ be a set of collision outcomes for M modes and N photons. Then, for a randomly chosen unitary circuit $U \sim \mathcal{H}_{\mathcal{A}}\mathcal{P}$, the probability to obtain collision outcomes is upper bounded by*

$$\mathop{\mathbb{E}}_{U\sim\mathcal{H}_{\mathcal{A}}\mathcal{P}}\left[\Pr_{\mathcal{D}_U}\left[\boldsymbol{s} \in B_{M,N}\right]\right] < \frac{2N^2}{M}, \tag{2}$$

where $\mathcal{D}_U$ is the output distribution according to $p_{\boldsymbol{s}}(U)$ defined in Eq. (1).

The proof of Lemma 2 is provided in Appendix B. Accordingly, the total probability weight of collision outcomes in the Boson Sampling output distribution can be made sufficiently small in the dilute regime $M = \Omega(N^2)$, which is the primary regime of interest in this work. Therefore, when using the circuit ensemble $\mathcal{H}_{\mathcal{A}}\mathcal{P}$, the bosonic birthday paradox in Lemma 2 ensures that collision-free outcomes dominate the probability weight, thereby enabling us to confine the outcome space only to collision-free outcomes.

It is worth emphasizing that, for the bosonic birthday paradox, the circuit ensemble need not be restricted to $\mathcal{H}_{\mathcal{A}}$. In fact, regardless of the underlying circuit distribution, the initial random permutation layer $\mathcal{P}$ prior to the circuit distribution ensures the bosonic birthday paradox. From an experimental perspective, rather than physically implementing the permutation circuit $\mathcal{P}$, one may instead prepare a uniformly random input configuration, which is operationally equivalent.

Remark 1 (Conditions for the bosonic birthday paradox)**.** *The bosonic birthday paradox is ensured in the dilute regime as long as the input collision-free configuration is prepared uniformly random by any means. This observation opens the possibility of improving hardness results for shallow-depth Boson Sampling by considering more general circuit ensembles beyond $\mathcal{H}_{\mathcal{A}}$.*

For the random outcome instances, as we can now only focus on collision-free outcomes for the hardness argument, we also define the random outcome ensemble as the uniform distribution over collision-free outcomes.

Definition 6 (Random collision-free outcome ensemble)**.** *We define $\mathcal{G}_{M,N}$ as the uniform distribution over $\binom{M}{N}$ number of collision-free outcomes of Boson Sampling with M modes and N photons. Each outcome $\boldsymbol{s} \sim \mathcal{G}_{M,N}$ is an M-dimensional output configuration vector representing the collision-free outcome, such that $\boldsymbol{s} = (s_1, \ldots, s_M)$ where each $s_i \in \{0, 1\}$ with $\sum_{i=1}^{M} s_i = N$.*

3 Worst-case hardness of output probability estimation

In this section, we find the worst-case #P-hardness of output probability estimation of shallow-depth linear optical circuits in the circuit architecture $\mathcal{BB}^*$ defined in Definition 3, for a fixed input $\boldsymbol{t}$ and a fixed collision-free output $\boldsymbol{s}_0$, to within a certain additive imprecision. Our worst-case hardness result for our shallow linear optical circuit architecture can be represented as follows.

Theorem 3 (Worst-case hardness)**.** *Given $M = \Omega(N)$, for a fixed collision-free outcome $\boldsymbol{s}_0$, approximating an output probability $p_{\boldsymbol{s}_0}(C_0)$ to within additive imprecision $2^{-O(N)}$ for any C_0 over linear optical circuit architecture $\mathcal{BB}^*$ is #P-hard in the worst case.*

We briefly sketch the proof of our worst-case hardness result for the shallow circuit architecture $\mathcal{BB}^*$. The proof is based on the result by [33], which showed the simulation hardness of exact Boson Sampling with constant-depth linear optical circuits. Specifically, as shown in [33], based on measurement-based quantum computation (MBQC) scheme, there exists a constant-depth linear optical circuit that can simulate an arbitrary given quantum circuit using post-selection. To extend this result to our case, we show that the constant-depth worst-case MBQC circuit in [33] can be embedded in our circuit architecture $\mathcal{BB}^*$. Combining these arguments, additively approximating the output probability of any quantum circuit can be reduced to additively approximating the output probability of any circuit in $\mathcal{BB}^*$, with imprecision blowup up to the inverse of post-selection probability. Using the fact that the additive approximation of any quantum circuit is #P-hard for certain additive imprecision [43], we can obtain the worst-case hardness of our shallow circuit architecture $\mathcal{BB}^*$.

Proof of Theorem 3. See Appendix C. □

We conclude this section by noting that, it is not necessary to derive worst-case hardness of shallow-depth Boson Sampling exclusively for the circuit architecture $\mathcal{BB}^*$. That is, embedding the constant-depth worst-case MBQC circuit [33] into the logarithmic-depth $\mathcal{BB}^*$ circuit, as done in our proof of Theorem 3, is not essential for the worst-case hardness of shallow-depth Boson Sampling in principle (and, as will be discussed later, not even for the worst-to-average-case reduction). Although not mandatory, however, since we later make use of $\mathcal{BB}^*$ to implement the global permutation during the worst-to-average-case reduction (due to the absence of hiding), we establish worst-case hardness within $\mathcal{BB}^*$ so as to simplify the overall circuit architecture to the repeated structure $(\mathcal{BB}^*)^q$ defined in Definition 3.

4 Average-case hardness of output probability estimation

In this section, we introduce our average-case hardness result of approximating output probabilities of shallow-depth Boson Sampling. We focus on q-Kaleidoscope circuit architecture $(\mathcal{BB}^*)^q$ in Definition 3 with $q = O(1)$, whose circuit depth is $D = 2q \log M$ and gate number is $m = qM \log M$. By Theorem 3, $\mathcal{BB}^*$ has the worst-case hardness, and thus $(\mathcal{BB}^*)^q$ with $q \geq 1$ also has the worst-case hardness by adding trivial gates (identity gates) over the worst-case circuit in $\mathcal{BB}^*$.

As introduced, our main strategy for the average-case hardness is the worst-to-average-case reduction, using the previous worst-case hardness result in Theorem 3. In other words, we show that if we can well-estimate the output probability *on average*, we can also well-estimate the worst-case output probability in Theorem 3, which means that well-estimating the average-case output probability is also #P-hard. Here, as we have stated previously, our average-case approximation regards both the randomly chosen outcome and the randomly chosen circuit.

We now formally state our average-case hardness result for shallow-depth linear optical circuit architecture $(\mathcal{BB}^*)^{q\geq 4}$ with $q = O(1)$.

Theorem 4 (Average-case hardness)**.** *The following problem is #P-hard under a* $\mathrm{BPP}^{\mathrm{NP}}$ *reduction: for any constant* $\delta, \eta \geq 0$ *with* $\delta + \eta < \frac{1}{4}$*, on input an* $O(\log N)$*-depth random circuit* $U \sim \mathcal{H}_{\mathcal{A}}\mathcal{P}$ *with* $\mathcal{A} = (\mathcal{BB}^*)^{q\geq 3}$ *and a random collision-free outcome* $\boldsymbol{s} \sim \mathcal{G}_{M,N}$*, compute the output probability* $p_{\boldsymbol{s}}(U)$ *within additive imprecision* $\epsilon = 2^{-O(N^{\gamma+1}(\log N)^2)}$*, with probability at least* $1-\delta$ *over the choice of* U *for at least* $1-\eta$ *over the choice of* $\boldsymbol{s}$*.*

We sketch the proof of Theorem 4 in the following by briefly describing the worst-to-average-case reduction process; we leave in Appendix F a detailed proof of Theorem 4. Now, we clarify how to establish the reduction: The goal is to well-estimate worst-case output probability $p_{\boldsymbol{s}_0}(C_0)$ for a fixed collision-free outcome $\boldsymbol{s}_0$ and worst-case circuit C_0 in $(\mathcal{BB}^*)^{q\geq 1}$ (Theorem 3), given oracle access to well-estimated values $p_{\boldsymbol{s}}(U)$ for average-case collision-free outcomes over $\boldsymbol{s} \sim \mathcal{G}_{M,N}$ and average-case circuits over $U \sim \mathcal{H}_{\mathcal{A}}\mathcal{P}$ with $\mathcal{A} = (\mathcal{BB}^*)^{q\geq 3}$ (in other words, U is a randomly chosen circuit over $(\mathcal{BB}^*)^{q\geq 4}$).

First, since our average-case argument regards both outcomes and circuits, we begin by describing how to *fix* the outcome in establishing the worst-to-average-case reduction. The remaining problem is then to establish a "circuit-level" worst-to-average-case reduction, from a worst-case circuit to randomly chosen average-case circuits for a fixed output probability, similarly to previous hardness results of random circuit sampling [28, 41–44]. To fix the outcome, our strategy is to randomly *permute* both a given worst-case circuit C_0 in $(\mathcal{BB}^*)^{q\geq 1}$ and a given collision-free outcome s_0. That is, we sample random M-mode permutation $\boldsymbol{P}_0$ uniformly, and permute the worst-case circuit C_0 and the fixed outcome $\boldsymbol{s}_0$ in Theorem 3 equally with $\boldsymbol{P}_0$. Then, the permuted outcome $\boldsymbol{s} = \boldsymbol{P}_0\boldsymbol{s}_0$ essentially follows the random outcome ensemble $\boldsymbol{s} \sim \mathcal{G}_{M,N}$, and the permuted circuit $\boldsymbol{P}_0 C_0$ is in $(\mathcal{BB}^*)^{q\geq 2}$ by Lemma 1. Note that the fixed worst-case output probability $p_{\boldsymbol{s}_0}(C_0)$ is now equivalent to $p_{\boldsymbol{s}}(\boldsymbol{P}_0 C_0)$, and thus our new goal is to estimate $p_{\boldsymbol{s}}(\boldsymbol{P}_0 C_0)$ value, given estimated average-case values of $p_{\boldsymbol{s}}(U)$.

Next, we describe how to estimate $p_{\boldsymbol{s}}(\boldsymbol{P}_0 C_0)$ given average-case estimation values $p_{\boldsymbol{s}}(U)$ for $U \sim \mathcal{H}_{\mathcal{A}}\mathcal{P}$ with $\mathcal{A} = (\mathcal{BB}^*)^{q\geq 3}$. For the permutation part $\mathcal{P}$ in the average-case circuit ensemble $\mathcal{H}_{\mathcal{A}}\mathcal{P}$, we newly sample an M-mode random permutation $\boldsymbol{P}_1 \sim \mathcal{P}$ again, and using the sampled $\boldsymbol{P}_1$ we set the circuit $\boldsymbol{P}_0 C_0 \boldsymbol{P}_1^{-1}$ as a "revised" worst-case circuit in $(\mathcal{BB}^*)^{q\geq 3}$. For the local random circuit part $\mathcal{H}_{\mathcal{A}}$ in the average-case circuit ensemble $\mathcal{H}_{\mathcal{A}}\mathcal{P}$, we sample a random circuit V from $\mathcal{H}_{\mathcal{A}}$ and *perturb* each gate of V by each gate of $\boldsymbol{P}_0 C_0 \boldsymbol{P}_1^{-1}$ parameterized by a constant $\theta \in [0, 1]$ using Cayley transform method, and obtain a perturbed circuit $V(\theta)$ (see Appendix E for details). Here, by Cayley transform, $V(\theta)$ lies on a continuous path over unitary matrices interpolating $V(0)$ and $V(1)$, where $\theta = 0$ corresponds to the random circuit drawn from local random circuit ensemble $V(0) \sim \mathcal{H}_{(\mathcal{BB}^*)^{q\geq 3}}$ and $\theta = 1$ corresponds to the worst-case circuit $V(1) = \boldsymbol{P}_0 C_0 \boldsymbol{P}_1^{-1}$. Notably, $V(0)\boldsymbol{P}_1$ follows the average-case circuit distribution $\mathcal{H}_{\mathcal{A}}\mathcal{P}$ with $\mathcal{A} = (\mathcal{BB}^*)^{q\geq 3}$, and thus, one can expect that the distribution for $V(\theta)\boldsymbol{P}_1$ for sufficiently small θ would not be largely deviated from the average-case circuit distribution. Accordingly, one can obtain well-estimated values of $p_{\boldsymbol{s}}(V(\theta)\boldsymbol{P}_1)$ for sufficiently small θ with high probability, given oracle access to the average-case output probabilities (i.e., Theorem 4).

Importantly, we find that $p_{\boldsymbol{s}}(V(\theta)\boldsymbol{P}_1)$ can be expressed as a finite-degree rational function $\frac{P(\theta)}{Q(\theta)}$ for $\theta \in [0, 1]$, where $Q(\theta)$ is efficiently computable. Hence, for sufficiently small θ values, one can also obtain the well-estimated values for the finite degree *polynomial* $P(\theta)$, by multiplying the well-estimated values of $p_{\boldsymbol{s}}(V(\theta)\boldsymbol{P}_1)$ with $Q(\theta)$. Then, given the estimated values of $P(\theta)$ for different θ values, one can use *polynomial interpolation* for the polynomial $P(\theta)$, and obtain the estimate of $P(1) = p_{\boldsymbol{s}}(V(1)\boldsymbol{P}_1)Q(1)$. Therefore, one can finally infer the value $p_{\boldsymbol{s}}(V(1)\boldsymbol{P}_1)$ by multiplying $Q(1)^{-1}$ to the estimate of $P(1)$, which corresponds to the worst-case output probability $p_{\boldsymbol{s}}(V(1)\boldsymbol{P}_1) = p_{\boldsymbol{s}}(\boldsymbol{P}_0 C_0 \boldsymbol{P}_1^{-1}\boldsymbol{P}_1) = p_{\boldsymbol{s}}(\boldsymbol{P}_0 C_0) = p_{\boldsymbol{s}_0}(C_0)$, concluding the worst-to-average-case reduction.

To sum up, for the worst-case output probability $p_{\boldsymbol{s}_0}(C_0)$ in Theorem 3 and the average-case output probability $p_{\boldsymbol{s}}(U)$ in Theorem 4, the overall reduction process can be schema-

tized by

$$p_{\boldsymbol{s}_0}(C_0) = p_{\boldsymbol{s}}(\boldsymbol{P}_0 C_0) = p_{\boldsymbol{s}}(\boldsymbol{P}_0 C_0 \boldsymbol{P}_1^{-1} \boldsymbol{P}_1) \xleftarrow[\substack{V \to \boldsymbol{P}_0 C_0 \boldsymbol{P}_1^{-1} \\ \text{Cayley transform} \\ \& \\ \text{polynomial} \\ \text{interpolation}}]{} p_{\boldsymbol{s}}(V \boldsymbol{P}_1) = p_{\boldsymbol{s}}(U) \approx \mathcal{O}(\boldsymbol{s}, U), \tag{3}$$

where

$\mathcal{O}$: Oracle for average-case output probabilities (Theorem 4)

C_0 : Worst-case circuit in $(\mathcal{BB}^*)^{q \geq 1}$

$\boldsymbol{s}_0$: A fixed collision-free outcome

$\boldsymbol{s} \sim \mathcal{G}_{M,N}$: Random collision-free outcome (Definition 6)

$V \sim \mathcal{H}_{(\mathcal{BB}^*)^{q \geq 3}}$: Local random circuit over $(\mathcal{BB}^*)^{q \geq 3}$ (Definition 4)

$\boldsymbol{P}_0$, $\boldsymbol{P}_1 \sim \mathcal{P}$: Random permutation circuit in $\mathcal{BB}^*$ (Definition 5).

So far, we have outlined the proof of Theorem 4, describing how to establish worst-to-average-case reduction. In Appendix E, we introduce in detail the Cayley transform, the circuit perturbation method for the worst-to-average-case reduction in our work. Subsequently, in Appendix F, we provide a detailed proof of Theorem 4, including the imprecision analysis we have missed out in the main text (i.e., why the average-case imprecision level $\epsilon = 2^{-O(N^{\gamma+1}(\log N)^2)}$ is necessary to achieve the worst-case imprecision level $2^{-O(N)}$).

Proof of Theorem 4. See Appendix F □

It is worth mentioning that, although we have concentrated on the dilute regime such that $\gamma \geq 2$, our average-case hardness result does not rule out the case $1 \leq \gamma < 2$ because it is derived from the worst-to-average-case reduction and our worst-case hardness result allows $M = \Omega(N)$. However, since our average-case hardness result only considers collision-free outcomes, it does not directly lead to the sampling hardness result (i.e., quantum advantage claim for shallow-depth Boson Sampling) by only increasing the allowed additive imprecision for the average-case hardness because collision-free outcomes are expected to occupy a small portion of outcomes when $1 \leq \gamma < 2$. Therefore, to lead to the sampling hardness argument when $1 \leq \gamma < 2$, our average-case hardness result should be extended to also contain collision outcomes for the average-case outcome instances (i.e., instead of $\mathcal{G}_{M,N}$, now uniformly choose over all possible outcomes). For a more complete analysis, in Appendix H, we present a promising approach that can be extended to the average-case hardness result that includes collision outcomes.

We also remark again that one could further reduce the circuit depth by taking the worst-case circuit C_0 as the constant-depth MBQC circuit in [33] directly, without encoding into the logarithmic-depth architecture $(\mathcal{BB}^*)^{q \geq 1}$ as in Theorem 3. In this case, the overall circuit architecture $\mathcal{A}$ for the average-case hardness consists of the constant-depth circuit (for MBQC) sandwiched between the $\mathcal{BB}^*$ architectures used for implementing required permutations (i.e., $\boldsymbol{P}_0$ and $\boldsymbol{P}_1$), yielding a slightly shallower circuit architecture than $\mathcal{A} = (\mathcal{BB}^*)^{q \geq 3}$ used in our case. Nevertheless, since we are already in the logarithmic-depth regime (by implementing permutations with $\mathcal{BB}^*$), we adopt $\mathcal{A} = (\mathcal{BB}^*)^{q \geq 3}$ as our overall circuit architecture for the ensemble $\mathcal{H}_{\mathcal{A}}$ in order to provide a unified description of the circuit architecture required for the hardness of shallow-depth Boson Sampling, as well as a clearer layout for experimental implementations.

5 Classical simulation hardness of shallow-depth boson sampling

As we have previously discussed, since our average-case hardness result considers both the random outcomes and random circuits, it is not straightforward to show the classical simulation hardness of shallow-depth Boson Sampling as in the original Boson Sampling proposal [1]. Therefore, in this section, we provide a self-contained analysis of how our average-case hardness result leads to the classical simulation hardness arguments of shallow-depth Boson Sampling. Specifically, we show that if the allowed additive error in Theorem 4 for the hardness is improved to a certain imprecision level, an efficient classical algorithm that can approximately simulate the shallow-depth Boson Sampling is unlikely to exist. This emphasizes that improving the imprecision level of the average-case hardness in Theorem 4 is a crucial step for the classical intractability of shallow-depth Boson Sampling.

Similarly to Refs. [1,41], we define an approximate boson sampler as follows.

Definition 7 (Approximate boson sampler)**.** *Approximate boson sampler is a classical randomized algorithm that takes input linear optical circuit C and outputs a sample from the output distribution $\mathcal{D}'_C$ such that*

$$||\mathcal{D}'_C - \mathcal{D}_C|| \leq \beta, \tag{4}$$

where $\mathcal{D}_C$ is the ideal output distribution of the circuit C and $||\cdot||$ represents total variation distance.

Given the total variation distance error, the above approximate sampler can have an arbitrarily large additive error for a fixed output probability. Nevertheless, it still has a comparably small additive error for *most* of the output probabilities due to Markov's inequality. Accordingly, finding the average-case solution of the output probability of the ideal sampler over randomly chosen collision-free outcome $\boldsymbol{s} \sim \mathcal{G}_{M,N}$, up to a certain additive imprecision, is in complexity class $\mathrm{BPP}^{\mathrm{NP}}$ by Stockmeyer's theorem [48].

Lemma 3 (Average-case approximation [1])**.** *If there exists an approximate boson sampler $\mathcal{S}$ with total variation distance β, then for any linear optical circuit C, the following problem is in $\mathrm{BPP}^{\mathrm{NP}^{\mathcal{S}}}$: find the average-case approximate solution $\tilde{p}_{\boldsymbol{s}}(C)$ of $p_{\boldsymbol{s}}(C)$, which satisfies*

$$\Pr_{\boldsymbol{s}\sim\mathcal{G}_{M,N}}\left[|\tilde{p}_{\boldsymbol{s}}(C) - p_{\boldsymbol{s}}(C)| \geq \frac{\kappa}{\binom{M}{N}}\right] \leq \xi, \tag{5}$$

where $\boldsymbol{s}$ is over all collision-free outcomes, and $\kappa, \xi > 0$ are the fixed error parameters satisfying $\beta = \kappa\xi/12$.

We leave the proof of Lemma 3 in Appendix G for a more self-contained analysis. The complexity $\mathrm{BPP}^{\mathrm{NP}}$ is known to be inside the finite level of PH, i.e., $\mathrm{BPP}^{\mathrm{NP}} \subseteq \mathrm{PH}$. Also, by Toda's theorem [49], PH problems can be solved given the ability to solve any #P problem, i.e., $\mathrm{BPP}^{\mathrm{NP}} \subseteq \mathrm{PH} \subseteq \mathrm{P}^{\#\mathrm{P}}$. If finding the average-case solution of output probabilities of sampler $\mathcal{S}$ is #P-hard, then $\mathrm{P}^{\#\mathrm{P}} \subseteq \mathrm{BPP}^{\mathrm{NP}^{\mathcal{S}}}$. Therefore, if an efficient classical algorithm exists that can simulate $\mathcal{S}$, then $\mathrm{P}^{\#\mathrm{P}} \subseteq \mathrm{BPP}^{\mathrm{NP}}$ which implies the collapse of PH. Consequently, under the assumption of the non-collapse of the PH, there is no efficient classical algorithm capable of simulating $\mathcal{S}$.

Based on the above arguments, we show that for the case where the allowed additive imprecision of Theorem 4 for the hardness can be improved, then it is classically hard to simulate shallow-depth Boson Sampling within a constant total variation distance.

Theorem 5. *Suppose that the allowed additive imprecision for the problem in Theorem 4 to be #P-hard can be improved to* $\epsilon = 2^{-(\gamma-1)N\log N - O(N)}$*. Then, in the dilute regime such that* $\gamma \geq 2$*, the efficient classical simulation of approximate boson sampler* $\mathcal{S}$ *with respect to circuits from the shallow architecture* $\mathcal{A} = (\mathcal{BB}^*)^{q\geq 4}$ *implies the collapse of* PH.

Proof. We establish a reduction from the problem in Theorem 4 with allowed additive error $\epsilon = 2^{-(\gamma-1)N\log N - O(N)}$ to the problem in Lemma 3. Let $\mathcal{O}$ be an oracle that solves the problem in Lemma 3, i.e., on input a circuit C and a randomly chosen collision-free outcome $\boldsymbol{s} \sim \mathcal{G}_{M,N}$, $\mathcal{O}$ outputs an estimate of output probability $p_{\boldsymbol{s}}(C)$ up to additive imprecision $\kappa\binom{M}{N}^{-1}$, with probability at least $1-\xi$ over outcomes for any circuit. For convenience, we refer to the outcomes which $\mathcal{O}$ estimates with error larger than $\kappa\binom{M}{N}^{-1}$ as *bad* outcomes, such that the portion of bad outcomes over possible collision-free outcomes is at most ξ for any circuit.

Note that, Lemma 3 holds for a randomly chosen input circuit drawn from an arbitrary circuit ensemble, such that the success probability of $\mathcal{O}$ is at least $1-\xi$ over randomly chosen outcomes and circuits. However, for the randomly chosen circuit input, bad outcomes can vary with the circuit instances, as the sampler $\mathcal{S}$ has the freedom to choose its error distribution according to the input circuit (e.g., C and $\boldsymbol{P}C$ would have different bad outcomes for permutation $\boldsymbol{P}$). Nevertheless, no matter how the bad outcomes vary with circuit instances, $\mathcal{O}$ succeeds at least $1-\frac{\xi}{\eta}$ fraction over circuit instances for at least $1-\eta$ fraction of the outcomes, for any η satisfying $\xi < \eta < 1$. Otherwise, the failure probability over outcomes and circuits would exceed ξ, which contradicts the proposition that the success probability of $\mathcal{O}$ is at least $1-\xi$ over randomly chosen outcomes and circuits. As the above argument holds for any random circuit families, we choose the random circuit distribution as $\mathcal{H}_{\mathcal{A}}\mathcal{P}$ with the shallow architecture $\mathcal{A} = (\mathcal{BB}^*)^{q\geq 3}$.

To sum up, on input a random circuit $H \sim \mathcal{H}_{\mathcal{A}}$ and a random outcome $\boldsymbol{s} \sim \mathcal{G}_{M,N}$, the oracle $\mathcal{O}$ estimates the output probability $p_{\boldsymbol{s}}(H)$ up to imprecision $\kappa\binom{M}{N}^{-1}$, with probability at least $1-\frac{\xi}{\eta}$ over the choice of H for at least $1-\eta$ over the choice of $\boldsymbol{s}$. Here, the additive imprecision can be bounded as

$$\kappa\binom{M}{N}^{-1} = \kappa\frac{N!(M-N)!}{M!} = 2^{-(\gamma-1)N\log N - O(N)}, \tag{6}$$

for constant β and so as κ, where we used the relation $M = c_0 N^{\gamma}$ with a constant c_0 and $\gamma \geq 1$. By setting η and ξ small constant satisfying $\eta + \frac{\xi}{\eta} < \frac{1}{4}$, we can solve the problem in Theorem 4 up to additive imprecision $\epsilon = 2^{-(\gamma-1)N\log N - O(N)}$ using the oracle $\mathcal{O}$. Hence, assuming that the above problem is #P-hard under $\mathrm{BPP}^{\mathrm{NP}}$ reduction, we can obtain the complexity-theoretical relation $\mathrm{P}^{\#\mathrm{P}} \subseteq \mathrm{BPP}^{\mathrm{NP}^{\mathcal{S}}}$, which implies the collapse of PH if $\mathcal{S}$ with respect to shallow-depth circuit architecture $(\mathcal{BB}^*)^{q\geq 4}$ can be done in classical polynomial time. This completes the proof. □

We lastly remark that, because our average-case hardness result in Sec. 4 solely relies on collision-free outcomes, it must be ensured that these collision-free outcomes occupy a sufficient portion of all possible outcomes of Boson Sampling (for more details, see Appendix B). Hence, we focus on the dilute regime $\gamma \geq 2$ for the classical hardness result in this section, to ensure the dominance of collision-free outcomes. To similarly obtain the classical hardness result in the saturated regime $1 \leq \gamma < 2$, extending our average-case hardness result for the collision outcomes is crucial.

6 Average-case hardness for shallow-depth Gaussian Boson Sampling

In this section, we show that our hardness results of the shallow-depth Boson Sampling can be generalized to the Gaussian Boson Sampling scheme [2]. Our specific setup for the Gaussian Boson Sampling is as follows. Let the total mode number M of the circuit be a power of 2, and now the input state is an M product of single-mode squeezed vacuum (SMSV) state $|\text{SMSV}\rangle^{\otimes M}$ with equal squeezing parameter r and equal squeezing direction. Also, let us define the output mean photon number as an integer N (i.e., $N = M\sinh^2 r$) where M and N are polynomially related as $M = c_1 N^\gamma$ for a constant c_1 and $\gamma \geq 2$. We define $q_{\boldsymbol{s}}(C)$ as an output probability of the Gaussian Boson Sampling, for an N-photon outcome $\boldsymbol{s}$ from an M-mode linear optical circuit matrix C on input M SMSV states. For collision-free outcome $\boldsymbol{s}$, $q_{\boldsymbol{s}}(C)$ can be expressed as [2]

$$q_{\boldsymbol{s}}(C) = \left|\langle \boldsymbol{s}|\,\hat{\mathcal{U}}(C)\,|\text{SMSV}\rangle^{\otimes M}\right|^2 = \frac{\tanh^N r}{\cosh^M r}|\text{Haf}((CC^T)_{\boldsymbol{s}})|^2, \tag{7}$$

where $|\boldsymbol{s}\rangle$ is an M-mode Fock-state corresponding to the outcome $\boldsymbol{s}$, $\hat{\mathcal{U}}(C)$ is a unitary operator corresponding to the circuit C, and $(CC^T)_{\boldsymbol{s}}$ is an N by N matrix obtained by taking s_i copies of the ith row and column of the matrix CC^T.

Using the above settings, we first prove the worst-case hardness of Gaussian Boson Sampling for a fixed outcome $\boldsymbol{s}$, with the shallow-depth circuit architecture $\mathcal{BB}^*$.

Theorem 6. *Approximating the output probability $q_{\boldsymbol{s}_0}(C)$ of Gaussian Boson Sampling to within additive error $2^{-\frac{\gamma-1}{2}N\log N - O(N)}$ for any C over linear optical circuit architecture $\mathcal{BB}^*$ is #P-hard in the worst case.*

Proof. We establish a reduction from the worst-case hardness of Boson Sampling in Theorem 3 to the problem in Theorem 6. Let $p_{\boldsymbol{s}_0}(C_0)$ be the output probability of a fixed input and output $\boldsymbol{s}_0$ of Boson Sampling in Theorem 3, for mode number M_0, photon number N_0, and the circuit C_0 in M_0 mode circuit architecture $\mathcal{BB}^*$. In the following, we show that $p_{\boldsymbol{s}_0}(C_0)$ can be efficiently reduced to the output probability $q_{\boldsymbol{s}}(C)$ of Gaussian Boson Sampling, for mode number $M = 2M_0$ and mean photon number $N = 2N_0$, with output $\boldsymbol{s}$ and circuit C determined by $\boldsymbol{s}_0$ and C_0 each.

Our strategy is to employ the scheme in Ref. [50], which used M_0 product of equally squeezed two-mode squeezed vacuum (TMSV) state as an input state to perform M_0 mode Boson Sampling task. Specifically, a single TMSV state with squeezing parameter r can be represented as

$$|\text{TMSV}\rangle = \frac{1}{\cosh r}\sum_{n=0}^{\infty}\tanh^n r\,|n\rangle\,|n\rangle\,, \tag{8}$$

and thus M_0 product of the TMSV state is

$$\begin{aligned}|\text{TMSV}\rangle^{\otimes M_0} &= \frac{1}{\cosh^{M_0} r}\left(\sum_{n=0}^{\infty}\tanh^n r\,|n\rangle_{(1)}\,|n\rangle_{(2)}\right)^{\otimes M_0}\\ &= \frac{1}{\cosh^{M_0} r}\sum_{n=0}^{\infty}\tanh^n r\sum_{\boldsymbol{s}_n}|\boldsymbol{s}_n\rangle_{(1)}\,|\boldsymbol{s}_n\rangle_{(2)}\,,\end{aligned} \tag{9}$$

where the summation of $\boldsymbol{s}_n$ is over all possible configurations of Fock-state with a total M_0 mode and n photon.

For each mode in the given M_0 mode circuit C_0, one-half of the TMSV state (i.e., subscript (2) in Eq. (9)) is input into it, and the other half of each state (i.e., subscript (1)

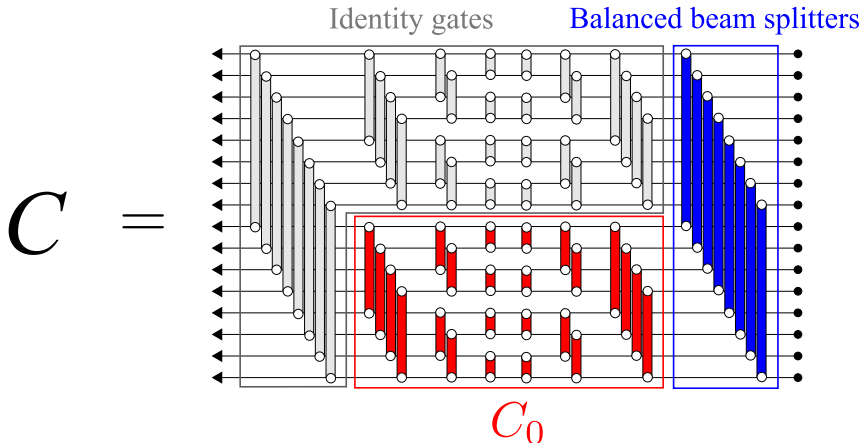

Figure 3: Schematics of an mode number $M = 16$ circuit C in $\mathcal{BB}^*$ which contains a given $M_0 = 8$ mode circuit C_0 in $\mathcal{BB}^*$

in Eq. (9)) is sent directly to a photon counter. By setting each $|\boldsymbol{s}_{\text{in}}\rangle$ and $|\boldsymbol{s}_{\text{out}}\rangle$ as a total M_0 mode and total N_0 photon Fock-state, the output probability can be represented as

$$\left|\langle \boldsymbol{s}_{\text{in}}|_{(1)} \langle \boldsymbol{s}_{\text{out}}|_{(2)} \hat{\mathcal{U}}_{(2)}(C_0) \left|\text{TMSV}\right\rangle^{\otimes M_0}\right|^2 = \frac{\tanh^{2N_0} r}{\cosh^{2M_0} r} \left|\langle \boldsymbol{s}_{\text{in}}| \hat{\mathcal{U}}(C_0) \left|\boldsymbol{s}_{\text{out}}\right\rangle\right|^2, \tag{10}$$

which is the output probability of M_0 mode and N_0 photon Boson Sampling in circuit C_0, with an additional multiplicative factor.

Note that two M_0 mode $\mathcal{BB}^*$ architecture can be embedded in the middle of an $M = 2M_0$ mode $\mathcal{BB}^*$ architecture. Accordingly, we define a circuit C in an M-mode $\mathcal{BB}^*$ by embedding the given M_0 mode circuit C_0 in one side of $\mathcal{BB}^*$, setting gates located right in front of the input ports as balanced beam splitters, and setting the remaining gates as identity gates; we leave in Fig. 3 an illustration of $M = 16$ mode circuit C for more clarity. Here, the input M SMSV states with squeezing parameter r combined with the balanced beam splitters at the front becomes M_0 TMSV states with squeezing parameter r. Therefore, our overall setup exactly follows the scheme in Ref. [50], such that the first M_0 mode is the photon counter sector to determine the input configuration of Boson Sampling, and the last M_0 mode is to simulate the Boson Sampling for the given circuit C_0.

We also define an M-dimensional vector $\boldsymbol{s}$ as a serial concatenation of two $\boldsymbol{s}_0$ vectors, so that $\boldsymbol{s}$ represents an N-photon outcome for M modes. Then the output probabilities $p_{\boldsymbol{s}_0}(C_0)$ and $q_{\boldsymbol{s}}(C)$ are related as

$$\begin{aligned} q_{\boldsymbol{s}}(C) &= \left|\langle \boldsymbol{s}| \hat{\mathcal{U}}(C) \left|\text{SMSV}\right\rangle^{\otimes M}\right|^2 \\ &= \left|\langle \boldsymbol{s}_0|_{(1)} \langle \boldsymbol{s}_0|_{(2)} \hat{\mathcal{U}}_{(2)}(C_0) \left|\text{TMSV}\right\rangle^{\otimes M_0}\right|^2 \\ &= \frac{\tanh^{2N_0} r}{\cosh^{2M_0} r} p_{\boldsymbol{s}_0}(C_0). \end{aligned} \tag{11}$$

Hence, approximating the output probability $p_{\boldsymbol{s}_0}(C_0)$ can be reduced to approximating the output probability $q_{\boldsymbol{s}}(C)$ of Gaussian Boson Sampling, with a blowup in the additive imprecision. The size of the additive imprecision blowup is

$$\frac{\cosh^{2M_0} r}{\tanh^{2N_0} r} = \left(\frac{M+N}{M}\right)^{M_0+N_0} \left(\frac{M}{N}\right)^{N_0} = 2^{\frac{\gamma-1}{2} N \log N + O(N)}, \tag{12}$$

using the relation $N = M \sinh r$ and $M \propto N^\gamma$. Since the allowed additive error for the worst-case hardness of $p_{\boldsymbol{s}_0}(C_0)$ is $2^{-O(N)}$, the allowed additive error for the reduction is $2^{-O(N)} 2^{-\frac{\gamma-1}{2} N \log N - O(N)} = 2^{-\frac{\gamma-1}{2} N \log N - O(N)}$. This completes the proof. □

Using the results of Theorem 6 and the previous proof of the average-case hardness of Boson Sampling in Theorem 4, it is straightforward to find the average-case hardness of Gaussian Boson Sampling, for a randomly chosen N-photon outcomes $\boldsymbol{s} \sim \mathcal{G}_{M,N}$ and a randomly chosen circuit $U \sim \mathcal{H}_{\mathcal{A}}$ in shallow-depth architecture $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^{q \geq 2}$.

Theorem 7. *The following problem is #P-hard under a* $\mathrm{BPP}^{\mathrm{NP}}$ *reduction: for any constant* $\delta, \eta \geq 0$ *with* $\delta + \eta < \frac{1}{4}$*, on input a* $O(\log N)$*-depth random circuit* $U \sim \mathcal{H}_{\mathcal{A}}$ *with* $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^{q \geq 2}$ *and a random collision-free outcome* $\boldsymbol{s} \sim \mathcal{G}_{M,N}$*, compute the output probability* $q_{\boldsymbol{s}}(U)$ *of Gaussian Boson Sampling within additive imprecision* $\epsilon = 2^{-O(N^{\gamma+1}(\log N)^2)}$*, with probability at least* $1 - \delta$ *over the choice of* U *for at least* $1 - \eta$ *over the choice of* $\boldsymbol{s}$*.*

It is worth emphasizing that for our setup of the Gaussian Boson Sampling, the bosonic birthday paradox is already guaranteed in the dilute regime $M = \Omega(N^2)$, in contrast to the Boson Sampling case. More specifically, the bosonic birthday paradox requires the permutation symmetry of inputs over the unitary circuits, in the sense that the input pattern should be invariant under the permutations (see Appendix B for more details). Since we are considering "full mode" input, our setup already satisfies this required input symmetry. To see this input symmetry clearly, one can easily check that for any unitary circuit U and permutation matrix $\boldsymbol{P}$, the output probability of Gaussian Boson Sampling for $U\boldsymbol{P}$ is equivalent to

$$q_{\boldsymbol{s}}(U\boldsymbol{P}) = \frac{\tanh^N r}{\cosh^M r} |\mathrm{Haf}((U\boldsymbol{P}\boldsymbol{P}^T U^T)_{\boldsymbol{s}})|^2 = \frac{\tanh^N r}{\cosh^M r} |\mathrm{Haf}((UU^T)_{\boldsymbol{s}})|^2 = q_{\boldsymbol{s}}(U). \tag{13}$$

Therefore, we do not necessarily have to include the random permutation part $\mathcal{P}$ for the average-case hardness result; only the local random circuit for $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^{q \geq 2}$ suffices.

Proof of Theorem 7. The procedure is essentially the same as the proof of Theorem 4, namely, establishing a worst-to-average-case reduction from the problem in Theorem 6 to the problem in Theorem 7: First, given worst-case circuit $C \in (\mathcal{B}\mathcal{B})^{q \geq 1}$ and outcome $\boldsymbol{s}_0$ depicted in Theorem 6, sample permutation $\boldsymbol{P}$, such that the revised worst-case circuit and outcome become $C_p = \boldsymbol{P}C \in (\mathcal{B}\mathcal{B})^{q \geq 2}$ and $\boldsymbol{s} = \boldsymbol{P}\boldsymbol{s}_0 \sim \mathcal{G}_{M,N}$ respectively. Then, we sample $V(\theta)$ from local random circuit ensemble with circuit architecture $\mathcal{A} = (\mathcal{B}\mathcal{B})^{q \geq 2}$ and perturb by $C_p = \boldsymbol{P}C$ via Cayley transform, such that $V(\theta)$ follows the distribution $\mathcal{H}^{C_p}_{\mathcal{A},\theta}$. Given access to the oracle that solves the problem in Theorem 7, we infer the output probability $q_{\boldsymbol{s}}(V(\theta))$ using different values of θ, and thus obtain the worst-case output probability $q_{\boldsymbol{s}}(V(1)) = q_{\boldsymbol{s}}(C_p) = q_{\boldsymbol{s}_0}(C)$ (for more details, see Appendix F, where we prove the average-case hardness for the Boson Sampling case).

Hence, the only different part for the Gaussian Boson Sampling case is the functional form of the output probability $q_{\boldsymbol{s}}(V(\theta))$ parameterized by θ via Cayley transformation. Accordingly, we show that $q_{\boldsymbol{s}}(V(\theta))$ can also be represented as a degree $(4mN, 4mN)$ rational function in θ, the same degree as the Boson Sampling case in Lemma 10.

From Eq. (7), the output probability $q_{\boldsymbol{s}}(V(\theta))$ has the form of

$$q_{\boldsymbol{s}}(V(\theta)) = \tanh^N r \, \mathrm{sech}^M r \left| \sum_{\mu \in \mathrm{PMP}} \prod_{j=1}^{N/2} \left[(V(\theta)V(\theta)^T)_{\boldsymbol{s}} \right]_{\mu(2j-1),\mu(2j)} \right|^2, \tag{14}$$

where μ is along all possible perfect matching permutations over N modes. From the proof of Lemma 10, using reduction to the common denominator for all of the $m = qM \log M$ gates, $[V(\theta)]_{j,k}$ can be represented as $(2m, 2m)$ rational function in θ with the common denominator $\prod_{i=1}^{m} q_i(\theta)$. Using this fact, one can easily check that $\prod_{j=1}^{N/2} \left[(V(\theta)V(\theta)^T)_{\boldsymbol{s}}\right]_{\mu(2j-1),\mu(2j)}$ can be represented as $(2mN, 2mN)$ rational function in θ, with the common denominator $[\prod_{i=1}^{m} q_i(\theta)]^N$ which does not change with μ. Therefore, the output probability can be represented as $q_{\boldsymbol{s}}(V(\theta)) = \frac{P(\theta)}{Q(\theta)}$, with each $Q(\theta) = [\prod_{i=1}^{m} |q_i(\theta)|^2]^N$ and $P(\theta)$ a degree $4mN$ polynomial function in θ.

Given that $q_{\boldsymbol{s}}(V(\theta))$ can be represented as a degree $(4mN, 4mN)$ rational function with the same denominator $Q(\theta) = [\prod_{i=1}^{m} |q_i(\theta)|^2]^N$ from the Boson Sampling case in Lemma 10, we can repeat all the steps identically to the proof of Theorem 4 and obtain the same result. □

7 Hardness analyses in lossy environments

In this section, we generalize our hardness results for *lossy* environments, namely, shallow-depth linear optical circuits suffering from photon loss channels after each gate implementation. The reason we consider such a noise channel is that photon loss is indeed a major source of error in optical systems [4–7]. Also, photon loss ruins the classical intractability of Boson Sampling, as there exist many efficient classical algorithms that can simulate lossy Boson Sampling within a constant total variation distance [15–17]. Therefore, we mainly deal with the photon loss error here; our goal is to provide evidence for the hardness of the approximate simulation of Boson Sampling in lossy shallow circuits within total variation distance error. For simplicity, we do not consider any photon gain error here, such as thermal radiation noise subjected to the circuits.

To proceed, we start with a brief review of the results presented by Ref. [51], which shows the hardness of simulating noisy quantum circuits. Specifically, one can simulate a noiseless circuit using a larger noisy circuit up to the desired imprecision, by establishing error-detecting code in the noisy circuit and post-selecting null syndrome measurements. Therefore, given the probability to post-select the no-error syndromes, one can approximate the output probability of the noiseless circuit from the output probability of the noisy circuit. Based on this argument, Ref. [28] demonstrates the average-case hardness of approximating output probabilities of noisy quantum circuits, under some plausible assumptions of the noise model. This result gives evidence of the approximate simulation hardness of noisy quantum circuits, within total variation distance error.

The main strategy of the above hardness results is approximating ideal output probabilities by post-selecting error-free results from noisy circuits. Here, we can directly apply their strategy to our case, i.e., lossy shallow-depth linear optical circuits. The crucial observation is that considering photon loss error on Boson Sampling, the error syndrome is the output photon number *itself*. Specifically, if the output photon number is the same as the input photon number, this implies that no loss occurred throughout the circuit. Therefore, by post-selecting the event that the measured output photon number is the same as the input photon number, we can infer ideal output probabilities.

For a more detailed analysis, we set the loss model as follows. Let the photon loss model $\mathcal{N}$ be local and stochastic. Specifically, $\mathcal{N}$ is a set of loss channels $\{\mathcal{N}_i\}_{i=1}^{l}$, such that after each unitary gate is applied, loss channel $\mathcal{N}_i$ acts on each mode participated in the unitary gate. Hence, the number of loss channels is $l = O(m)$ for gate number m in a

given circuit architecture. We can decompose each noise channel $\mathcal{N}_i$ as follows:

$$\mathcal{N}_i = (1-\rho_i)\mathcal{I} + \rho_i \mathcal{E}_i, \tag{15}$$

where $\mathcal{I}$ is identity, $\mathcal{E}_i$ is an CPTP map representing photon loss, and ρ_i is a loss rate for each channel satisfying $\rho_i \leq \rho$ for a constant ρ. The validity of such modeling for photon loss channel is represented in [3, 15].

To simplify, we assume that we know *a priori* each error rate ρ_i for all $i \in [l]$, and the noise model $\mathcal{N}$ is fixed so that it does not change with random circuit instances. Then we can obtain the hardness of approximating output probabilities of lossy shallow circuits, from our previous hardness proposals. To do so, let $p_{\boldsymbol{s}}(C, \mathcal{N})$ be the output probability of an N-photon outcome $\boldsymbol{s}$ from an M-mode linear optical circuit C which undergoes loss model $\mathcal{N}$ we set. By post-selecting 'no loss event', which can be accomplished by counting the output photon number, the ideal output probability $p_{\boldsymbol{s}}(C)$ can be inferred from $p_{\boldsymbol{s}}(C, \mathcal{N})$ by

$$p_{\boldsymbol{s}}(C) = \frac{p_{\boldsymbol{s}}(C, \mathcal{N})}{\Pr[\text{'no loss event'}]}. \tag{16}$$

From Eq. (15), the probability of 'no loss event' is $\prod_{i=1}^{l}(1-\rho_i)$, which can be efficiently calculated. This implies that approximating $p_{\boldsymbol{s}}(C, \mathcal{N})$ can be reduced from approximating $p_{\boldsymbol{s}}(C)$, with at most $\Pr[\text{'no loss event'}]^{-1} = \prod_{i=1}^{l}(1-\rho_i)^{-1} \leq (1-\rho)^{-l} = 2^{O(\rho m)}$ blowup in the additive imprecision.

Given Eq. (16), we can repeat the same steps from the previous hardness arguments, for the lossy shallow-depth Boson Sampling; the only difference is the imprecision blowup by $2^{O(\rho m)}$. For our shallow-depth architecture $(\mathcal{BB}^*)^q$, the gate number m is $qM \log M$, so the size of imprecision blowup is $2^{O(N^\gamma \log N)}$ in our case. Such imprecision blowup does not affect the allowed additive accuracy $\epsilon = 2^{-O(N^{\gamma+1}(\log N)^2)}$ for our average-case hardness result. Based on the arguments so far, the following corollary is straightforward from Theorem 4.

Corollary 1. *Suppose we have the photon loss model $\mathcal{N}$ with each loss rate $\rho_i \leq \rho$ for a constant ρ. Then the following problem is #P-hard under a $\mathrm{BPP}^{\mathrm{NP}}$ reduction: for any constant $\delta, \eta \geq 0$ with $\delta + \eta < \frac{1}{4}$, on input a $O(\log N)$-depth random circuit $U \sim \mathcal{H}_{\mathcal{A}}\mathcal{P}$ with $\mathcal{A} = (\mathcal{BB}^*)^{q \geq 3}$ and a random collision-free outcome $\boldsymbol{s} \sim \mathcal{G}_{M,N}$, compute the lossy output probability $p_{\boldsymbol{s}}(U, \mathcal{N})$ within additive imprecision $\epsilon = 2^{-O(N^{\gamma+1}(\log N)^2)}$, with probability at least $1-\delta$ over the choice of U for at least $1-\eta$ over the choice of $\boldsymbol{s}$.*

We remark that for our noise model, the imprecision blowup grows exponentially with the gate number m. Therefore, shallow-depth circuits can be more advantageous in this perspective, since they are likely to have less gate number and thus have small imprecision blowup. For example, the current hardness results are based on M by N submatrices of M-dimensional Haar random unitaries, and the implementation of such matrices requires gate number $m = \Omega(N^{\gamma+1})$. This arouses the imprecision blowup at least $2^{O(N^{\gamma+1})}$, which restricts the allowed additive error for the average-case hardness at most $2^{-O(N^{\gamma+1})}$.

8 Concluding remarks

Here we provide a few remarks about our overall results and related open questions.

1. Our result demonstrates the average-case hardness for additive imprecision $2^{-O(N^{\gamma+1}(\log N)^2)}$. Indeed, there still remains a gap to the desired additive imprecision for the simulation hardness $2^{-(\gamma-1)N\log N - O(N)}$ in Theorem 5. Hence, closing this gap would be an ultimate

challenge to the full achievement of classical intractability; more advanced proof techniques are required to reduce this gap. Here, one can take the following approach: finite-size numerical experiments suggest that the output distributions of local random circuits in our circuit architecture (i.e., Definition 2) are close enough to those of global Haar random circuits (see, e.g., Ref. [47] and Appendix I). Accordingly, if one can analytically prove that the distance between those output distributions is close enough, we can directly obtain a better imprecision level $2^{-O(N \log N)}$ by results in [3, 28–30], which employed the global Haar random circuits.

Another possible approach for reducing the imprecision gap is to perturb a random circuit matrix in a different way from the Cayley transform (i.e., Definition 8), i.e., as depicted in [29]. Specifically, instead of perturbing each random gate, one can perturb a submatrix X of our random circuit matrix U with a worst-case matrix A as $X(\theta) = (1-\theta)X + \theta A$ for $\theta \in [0, 1]$. Here, a degree of polynomial $|\mathrm{Per}(X(\theta))|^2$ is $2N$, which is lower than ours derived by the Cayley transform. Therefore, if one can prove that $X(\theta)$ is distributed similarly to X for small θ, we expect that we can also obtain a better imprecision level by using the same interpolation method. However, the above approach requires one to figure out a global circuit distribution generated by the convolution of local circuit distributions. Although we believe that this problem can be resolved using techniques from random matrix theories, we have not yet developed a complete analysis. Hence, we leave it as an open question.

2. Another important challenge that should be addressed is the extension of the result in Sec. 7, i.e., finding the classical simulation hardness of noisy Gaussian Boson Sampling with photon loss, or Boson Sampling for more general types of noise. As described in [3, 28, 51], employing the threshold theorem would be a viable choice for this goal, which requires efficient error detection codes using only linear optical elements for these setups. However, to the best of our knowledge, such an error detection code does not yet exist. That is, in contrast to the Boson Sampling case, we cannot post-select the so-called 'no loss' event in the Gaussian Boson Sampling case because the input photon number is not fixed in this case. This is also for the Boson Sampling with general types of noise because we cannot post-select no error event by using only photon number measurements. Hence, constructing any of these error detection codes would be a crucial step toward the classical hardness of noisy (Gaussian) Boson Sampling, which will contribute to a more noise-tolerant demonstration of quantum advantage with (Gaussian) Boson Sampling. We leave this problem as another open question.

3. One of the existing gaps between theoretical Boson Sampling hardness arguments and Boson Sampling experiments to date is the ratio between mode number M and photon number N. Generally, most of the current theoretical Boson Sampling hardness arguments (including our current work) consider only collision-free outcomes and rule out collision outcomes due to the convenience of mathematical analysis [1, 3, 27, 28, 30]. This calls for the "dilute" regime $M = \Omega(N^2)$ to ensure the dominance of collision-free outcomes over the probability weight. In contrast, all the experiments to date [4–7] operate in the "saturated regime" where the number of modes is comparable to the number of photons. Therefore, to reduce this gap, extending our hardness result of shallow-depth Boson Sampling to the saturated regime $M = O(N^\gamma)$ for $1 \leq \gamma < 2$ is crucial, similarly to [29] for the general Boson Sampling case. To do so, one also needs to consider collision outcomes in the analysis because collision outcomes now dominate the probability weight in this regime. Here, given that our worst-case hardness result in Theorem 3 holds for a linear number of modes $M = \Omega(N)$, this opens the possibility of establishing average-case hardness in the saturated regime by extending the worst-to-average-case reduction from Theorem 3

to include collision outcomes. However, since the collision outcomes do not interchange with each other generally by permutations (in contrast to collision-free outcomes), our proof technique in Sec. 4 using random permutation does not work for general collision outcomes. Therefore, more advanced proof techniques are required to address this challenge. In Appendix H, we present an approach that could offer a useful insight toward extending our average-case hardness result to include collision outcomes. Also, another possible way is to analytically prove that the local random circuit for our circuit architecture $(\mathcal{BB}^*)^q$ is close enough to those of global Haar random circuits (numerically shown in [47] and Appendix I). If this can be done, we can employ the previous argument [29], which used global Haar random circuits, for our circuit architecture $(\mathcal{BB}^*)^q$ in the saturated regime.

4. We finally remark that if we only use the local random circuit ensemble $\mathcal{H}_{\mathcal{A}}$ for our shallow-depth circuit architecture $\mathcal{A} = (\mathcal{BB}^*)^q$ as defined in Definition 4, the bosonic birthday paradox is not guaranteed because of the absence of translation invariance property (in contrast to the global Haar random unitary circuit). Hence, for the sake of analytical completeness, we introduced the additional (shallow-depth) random permutation circuit at the input to guarantee the bosonic birthday paradox and suppress the occurrence of collision outcomes. In fact, from an experimental perspective, this setup can be realized by constructing the local random circuit $\mathcal{H}_{\mathcal{A}}$ and preparing a random input configuration. Furthermore, if one can analytically verify that the local random circuit ensemble $\mathcal{H}_{\mathcal{A}}$ has a permutation invariance property for $\mathcal{A} = (\mathcal{BB}^*)^q$, then we can safely leave out this random permutation circuit part, thereby simplifying the average-case circuits and reducing the required circuit depth for the average-case hardness. We leave in Appendix I numerical analysis for the bosonic birthday paradox for our local random circuit ensemble $\mathcal{H}_{\mathcal{A}}$ with $\mathcal{A} = (\mathcal{BB}^*)^q$, which gives evidence that only the local random circuit ensemble might be sufficient for the average-case hardness result.

Acknowledgements

We thank Bill Fefferman for insightful discussions. This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (Nos. RS-2024-00413957, RS-2024-00438415, and RS-2023-NR076733) and by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grants funded by the Korea government (MSIT) (IITP-2026-RS-2020-II201606 and IITP-2026-RS-2024-00437191). B.G. was supported by the education and training program of the Quantum Information Research Support Center, funded through the National Research Foundation of Korea (NRF) by the Ministry of Science and ICT (MSIT) of the Korean government (No.2021M3H3A1036573). C.O. was supported by the National Research Foundation of Korea Grants (No. RS-2024-00431768 and No. RS-2025-00515456) funded by the Korean government (Ministry of Science and ICT (MSIT)) and the Institute of Information & Communications Technology Planning & Evaluation (IITP) Grants funded by the Korea government (MSIT) (No. IITP-2025-RS-2025-02283189 and IITP-2025-RS-2025-02263264). This work was supported by Global Partnership Program of Leading Universities in Quantum Science and Technology (RS-2025-08542968) through the National Research Foundation of Korea (NRF) funded by the Korean government (Ministry of Science and ICT(MSIT)).

References

[1] Scott Aaronson and Alex Arkhipov. "The computational complexity of linear optics". In Proceedings of the forty-third annual ACM symposium on Theory of computing.

Pages 333–342. (2011). doi: 10.1145/1993636.1993682.

[2] Craig S Hamilton, Regina Kruse, Linda Sansoni, Sonja Barkhofen, Christine Silberhorn, and Igor Jex. "Gaussian boson sampling". Physical review letters **119**, 170501 (2017). doi: 10.1103/PhysRevLett.119.170501.

[3] Abhinav Deshpande, Arthur Mehta, Trevor Vincent, Nicolás Quesada, Marcel Hinsche, Marios Ioannou, Lars Madsen, Jonathan Lavoie, Haoyu Qi, Jens Eisert, et al. "Quantum computational advantage via high-dimensional gaussian boson sampling". Science advances **8**, eabi7894 (2022). doi: 10.1126/sciadv.abi7894.

[4] 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, et al. "Quantum computational advantage using photons". Science **370**, 1460–1463 (2020). doi: 10.1126/science.abe8770.

[5] 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". Physical review letters **127**, 180502 (2021). doi: 10.1103/PhysRevLett.127.180502.

[6] Lars S Madsen, Fabian Laudenbach, Mohsen Falamarzi Askarani, Fabien Rortais, Trevor Vincent, Jacob FF Bulmer, Filippo M Miatto, Leonhard Neuhaus, Lukas G Helt, Matthew J Collins, et al. "Quantum computational advantage with a programmable photonic processor". Nature **606**, 75–81 (2022). doi: 10.1038/s41586-022-04725-x.

[7] Yu-Hao Deng, Yi-Chao Gu, Hua-Liang Liu, Si-Qiu Gong, Hao Su, Zhi-Jiong Zhang, Hao-Yang Tang, Meng-Hao Jia, Jia-Min Xu, Ming-Cheng Chen, et al. "Gaussian boson sampling with pseudo-photon-number-resolving detectors and quantum computational advantage". Physical review letters **131**, 150601 (2023). doi: 10.1103/PhysRevLett.131.150601.

[8] Jelmer Renema, Valery Shchesnovich, and Raul Garcia-Patron. "Classical simulability of noisy boson sampling" (2018). doi: 10.48550/arXiv.1809.01953. arXiv:1809.01953.

[9] Jelmer J Renema, Adrian Menssen, William R Clements, Gil Triginer, William S Kolthammer, and Ian A Walmsley. "Efficient classical algorithm for boson sampling with partially distinguishable photons". Physical review letters **120**, 220502 (2018). doi: 10.1103/PhysRevLett.120.220502.

[10] Alexandra E Moylett, Raúl García-Patrón, Jelmer J Renema, and Peter S Turner. "Classically simulating near-term partially-distinguishable and lossy boson sampling". Quantum Science and Technology **5**, 015001 (2019). doi: 10.1088/2058-9565/ab5555.

[11] Jelmer J Renema. "Simulability of partially distinguishable superposition and gaussian boson sampling". Physical Review A **101**, 063840 (2020). doi: 10.1103/PhysRevA.101.063840.

[12] Junheng Shi and Tim Byrnes. "Effect of partial distinguishability on quantum supremacy in gaussian boson sampling". npj Quantum Information **8**, 54 (2022). doi: 10.1038/s41534-022-00557-9.

[13] SN van den Hoven, E Kanis, and JJ Renema. "Efficient classical algorithm for simulating boson sampling with inhomogeneous partial distinguishability" (2024). doi: 10.48550/arXiv.2406.17682. arXiv:2406.17682.

[14] Valery S Shchesnovich. "Noise in boson sampling and the threshold of efficient classical simulatability". Physical Review A **100**, 012340 (2019). doi: 10.1103/PhysRevA.100.012340.

[15] Michał Oszmaniec and Daniel J Brod. "Classical simulation of photonic linear optics with lost particles". New Journal of Physics **20**, 092002 (2018). doi: 10.1088/1367-2630/aadfa8.

[16] Raúl García-Patrón, Jelmer J Renema, and Valery Shchesnovich. "Simulating boson sampling in lossy architectures". Quantum **3**, 169 (2019). doi: 10.22331/q-2019-08-05-169.

[17] Haoyu Qi, Daniel J Brod, Nicolás Quesada, and Raúl García-Patrón. "Regimes of classical simulability for noisy gaussian boson sampling". Physical review letters **124**, 100502 (2020). doi: 10.1103/PhysRevLett.124.100502.

[18] Daniel Jost Brod and Michał Oszmaniec. "Classical simulation of linear optics subject to nonuniform losses". Quantum **4**, 267 (2020). doi: 10.22331/q-2020-05-14-267.

[19] Benjamin Villalonga, Murphy Yuezhen Niu, Li Li, Hartmut Neven, John C Platt, Vadim N Smelyanskiy, and Sergio Boixo. "Efficient approximation of experimental gaussian boson sampling" (2021). doi: 10.48550/arXiv.2109.11525. arXiv:2109.11525.

[20] Changhun Oh, Kyungjoo Noh, Bill Fefferman, and Liang Jiang. "Classical simulation of lossy boson sampling using matrix product operators". Physical Review A **104**, 022407 (2021). doi: 10.1103/PhysRevA.104.022407.

[21] 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). doi: 10.1126/sciadv.abl9236.

[22] Minzhao Liu, Changhun Oh, Junyu Liu, Liang Jiang, and Yuri Alexeev. "Simulating lossy gaussian boson sampling with matrix-product operators". Physical Review A **108**, 052604 (2023). doi: 10.1103/PhysRevA.108.052604.

[23] Changhun Oh, Liang Jiang, and Bill Fefferman. "On classical simulation algorithms for noisy boson sampling" (2023). doi: 10.48550/arXiv.2301.11532. arXiv:2301.11532.

[24] Changhun Oh, Minzhao Liu, Yuri Alexeev, Bill Fefferman, and Liang Jiang. "Classical algorithm for simulating experimental gaussian boson sampling". Nature PhysicsPages 1–8 (2024). doi: 10.1038/s41567-024-02535-8.

[25] Changhun Oh. "Classical simulability of constant-depth linear-optical circuits with noise". npj Quantum Information **11**, 126 (2025). doi: 10.1038/s41534-025-01041-w.

[26] Changhun Oh. "Recent theoretical and experimental progress on boson sampling". Current Optics and Photonics **9**, 1–18 (2025). doi: 10.3807/COPP.2025.9.1.1.

[27] Daniel Grier, Daniel J Brod, Juan Miguel Arrazola, Marcos Benicio de Andrade Alonso, and Nicolás Quesada. "The complexity of bipartite gaussian boson sampling". Quantum **6**, 863 (2022). doi: 10.22331/q-2022-11-28-863.

[28] Adam Bouland, Bill Fefferman, Zeph Landau, and Yunchao Liu. "Noise and the frontier of quantum supremacy". In 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS). Pages 1308–1317. IEEE (2022). doi: 10.1109/FOCS52979.2021.00127.

[29] Adam Bouland, Daniel Brod, Ishaun Datta, Bill Fefferman, Daniel Grier, Felipe Hernandez, and Michal Oszmaniec. "Complexity-theoretic foundations of boson-sampling with a linear number of modes" (2023). doi: 10.48550/arXiv.2312.00286. arXiv:2312.00286.

[30] Adam Bouland, Ishaun Datta, Bill Fefferman, and Felipe Hernandez. "On the average-case hardness of bosonsampling" (2024). doi: 10.48550/arXiv.2411.04566. arXiv:2411.04566.

[31] Karol Zyczkowski and Marek Kus. "Random unitary matrices". Journal of Physics A: Mathematical and General **27**, 4235 (1994). doi: 10.1088/0305-4470/27/12/028.

[32] Nicholas J Russell, Levon Chakhmakhchyan, Jeremy L O'Brien, and Anthony Laing. "Direct dialling of haar random unitary matrices". New journal of physics **19**, 033007 (2017). doi: 10.1088/1367-2630/aa60ed.

[33] Daniel J Brod. "Complexity of simulating constant-depth bosonsampling". Physical Review A **91**, 042316 (2015). doi: 10.1103/PhysRevA.91.042316.

[34] Guifré Vidal. "Efficient classical simulation of slightly entangled quantum computations". Physical review letters **91**, 147902 (2003). doi: 10.1103/PhysRevLett.91.147902.

[35] Haoyu Qi, Diego Cifuentes, Kamil Brádler, Robert Israel, Timjan Kalajdzievski, and Nicolás Quesada. "Efficient sampling from shallow gaussian quantum-optical circuits with local interactions". Physical Review A **105**, 052412 (2022). doi: 10.1103/PhysRevA.105.052412.

[36] Abhinav Deshpande, Bill Fefferman, Minh C Tran, Michael Foss-Feig, and Alexey V Gorshkov. "Dynamical phase transitions in sampling complexity". Physical review letters **121**, 030501 (2018). doi: 10.1103/PhysRevLett.121.030501.

[37] Changhun Oh, Youngrong Lim, Bill Fefferman, and Liang Jiang. "Classical simulation of boson sampling based on graph structure". Physical Review Letters **128**, 190501 (2022). doi: 10.1103/PhysRevLett.128.190501.

[38] Zoltán Kolarovszki, Ágoston Kaposi, Tamás Kozsik, and Zoltán Zimborás. "Simulating sparse and shallow gaussian boson sampling". In International Conference on Computational Science. Pages 209–223. Springer (2023). doi: 10.1007/978-3-031-36030-5_17.

[39] Chao Shen, Zhen Zhang, and L-M Duan. "Scalable implementation of boson sampling with trapped ions". Physical review letters **112**, 050504 (2014). doi: 10.1103/PhysRevLett.112.050504.

[40] Wentao Chen, Yao Lu, Shuaining Zhang, Kuan Zhang, Guanhao Huang, Mu Qiao, Xiaolu Su, Jialiang Zhang, Jing-Ning Zhang, Leonardo Banchi, et al. "Scalable and programmable phononic network with trapped ions". Nature PhysicsPages 1–7 (2023). doi: 10.1038/s41567-023-01952-5.

[41] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. "On the complexity and verification of quantum random circuit sampling". Nature Physics **15**, 159–163 (2019). doi: 10.1038/s41567-018-0318-2.

[42] Ramis Movassagh. "The hardness of random quantum circuits". Nature PhysicsPages 1–6 (2023). doi: 10.1038/s41567-023-02131-2.

[43] Yasuhiro Kondo, Ryuhei Mori, and Ramis Movassagh. "Quantum supremacy and hardness of estimating output probabilities of quantum circuits". In 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS). Pages 1296–1307. IEEE (2022). doi: 10.1109/FOCS52979.2021.00126.

[44] Hari Krovi. "Average-case hardness of estimating probabilities of random quantum circuits with a linear scaling in the error exponent" (2022). doi: 10.48550/arXiv.2206.05642. arXiv:2206.05642.

[45] Alex Arkhipov and Greg Kuperberg. "The bosonic birthday paradox". Geometry & Topology Monographs **18**, 10–2140 (2012). doi: 10.2140/gtm.2012.18.1.

[46] Tri Dao, Nimit S Sohoni, Albert Gu, Matthew Eichhorn, Amit Blonder, Megan Leszczynski, Atri Rudra, and Christopher Ré. "Kaleidoscope: An efficient, learnable representation for all structured linear maps" (2020). doi: 10.48550/arXiv.2012.14966. arXiv:2012.14966.

[47] Byeongseon Go, Changhun Oh, Liang Jiang, and Hyunseok Jeong. "Exploring shallow-depth boson sampling: Toward a scalable quantum advantage". Physical Review A **109**, 052613 (2024). doi: 10.1103/PhysRevA.109.052613.

[48] Larry Stockmeyer. "On approximation algorithms for# p". SIAM Journal on Computing **14**, 849–861 (1985). doi: 10.1137/0214060.

[49] Seinosuke Toda. "Pp is as hard as the polynomial-time hierarchy". SIAM Journal on Computing **20**, 865–877 (1991). doi: 10.1137/0220053.

[50] Austin P Lund, Anthony Laing, Saleh Rahimi-Keshari, Terry Rudolph, Jeremy L O'Brien, and Timothy C Ralph. "Boson sampling from a gaussian state". Physical review letters **113**, 100502 (2014). doi: 10.1103/PhysRevLett.113.100502.

[51] Keisuke Fujii. "Noise threshold of quantum supremacy" (2016). doi: 10.48550/arXiv.1610.03632. arXiv:1610.03632.

[52] Robert Raussendorf and Hans J Briegel. "A one-way quantum computer". Physical review letters **86**, 5188 (2001). doi: 10.1103/PhysRevLett.86.5188.

[53] Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. "Universal blind quantum computation". In 2009 50th Annual IEEE Symposium on Foundations of Computer Science. Pages 517–526. IEEE (2009). doi: 10.1109/FOCS.2009.36.

[54] Andrew M Childs, Debbie W Leung, and Michael A Nielsen. "Unified derivations of measurement-based schemes for quantum computation". Physical Review A **71**, 032318 (2005). doi: 10.1103/PhysRevA.71.032318.

[55] Emanuel Knill, Raymond Laflamme, and Gerald J Milburn. "A scheme for efficient quantum computation with linear optics". nature **409**, 46–52 (2001). doi: 10.1038/35051009.

[56] Emanuel Knill. "Quantum gates using linear optics and postselection". Physical Review A **66**, 052306 (2002). doi: 10.1103/PhysRevA.66.052306.

[57] Michael J Bremner, Ashley Montanaro, and Dan J Shepherd. "Average-case complexity versus approximate simulation of commuting quantum computations". Physical review letters **117**, 080501 (2016). doi: 10.1103/PhysRevLett.117.080501.

[58] Dominik Hangleiter, Juan Bermejo-Vega, Martin Schwarz, and Jens Eisert. "Anticoncentration theorems for schemes showing a quantum speedup". Quantum **2**, 65 (2018). doi: 10.22331/q-2018-05-22-65.

[59] Juan Bermejo-Vega, Dominik Hangleiter, Martin Schwarz, Robert Raussendorf, and Jens Eisert. "Architectures for quantum simulation showing a quantum speedup". Physical Review X **8**, 021010 (2018). doi: 10.1103/PhysRevX.8.021010.

[60] Alexander M Dalzell, Nicholas Hunter-Jones, and Fernando GSL Brandão. "Random quantum circuits anticoncentrate in log depth". PRX Quantum **3**, 010333 (2022). doi: 10.1103/PRXQuantum.3.010333.

[61] Sheldon M Ross. "Introduction to probability models". Academic press. (2014). doi: 10.1016/C2012-0-03564-8.

[62] Peter Clifford and Raphaël Clifford. "The classical complexity of boson sampling". In Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. Pages 146–155. SIAM (2018). doi: 10.1137/1.9781611975031.10.

A Previous foundations on average-case hardness

In this appendix, we introduce the existing proof technique used to establish the classical hardness of Boson Sampling, specifically in the context of approximate simulation within a bounded total variation distance error. [1]. The current state-of-the-art proof technique for the hardness of sampling problems like Boson Sampling essentially builds upon Stockmeyer's algorithm about approximate counting [48]. Specifically, given a classical sampler that outputs a sample from a given output distribution, Stockmeyer's algorithm enables one to multiplicatively estimate a fixed output probability of the sampler, within complexity class $\mathrm{BPP}^{\mathrm{NP}}$.

Now suppose there exists an approximate classical sampler capable of simulating ideal Boson Sampling up to total variation distance error, as in Definition 7. This approximate sampler can have a large additive error for a fixed output probability, but have a comparably small additive error for *most* of the output probabilities due to Markov's inequality. Then, Stockmeyer's algorithm, combined with the approximate sampler, can well approximate the ideal output probability of Boson Sampling within a certain additive error, with a high probability over randomly chosen outcomes (See Lemma 3 for more details). For convenience, let us refer to this computational task as an *average-case approximation* problem of Boson Sampling. If the complexity of the average-case approximation problem is outside the Polynomial Hierarchy (PH), it implies the collapse of PH, since the complexity of Stockmeyer's algorithm is indeed inside the finite level of PH.

Here, *average-case hardness* comes into the proof of the classical simulation hardness argument, which means that approximating the ideal output probability of Boson Sampling with high probability over randomly chosen outcomes is #P-hard. More precisely, if the average-case hardness holds up to the imprecision level of the average-case approximation problem, this comes down to the classical simulation hardness of the approximate sampling unless PH collapses, by the complexity-theoretical foundation $\mathrm{PH} \subseteq \mathrm{P}^{\#\mathrm{P}}$ [49].

Moreover, by choosing random circuit instances that have symmetry over the outcomes, one can reduce the average-case instances for the hardness from *outcome* instances to *circuit* instances, which is called the *hiding* property. For the Boson Sampling case, global Haar random unitary (i.e., unitary matrix drawn from Haar measure on U(M), for mode number M) satisfies this condition. In detail, instead of randomly choosing the outcome, we can fix the outcome by applying a random permutation to the global Haar random unitary distribution, which is still Haar distributed due to its translation invariance property. This hiding property plays an important role in the current proofs of the average-case hardness, as it enables one to establish worst-to-average-case reduction in terms of circuit instances. Specifically, as the output probability of Boson Sampling can be written as a low-degree polynomial of input circuit (matrix) values, it allows one to infer the value of a worst-case instance from the output probability of many average-case circuit instances. Hence, the

average-case hardness argument for Boson Sampling is typically used in this context, i.e., average-case hardness over random circuit instances, for a fixed outcome [1, 3, 28].

Accordingly, the crucial step for the classical simulation hardness of approximate sampling is to prove the average-case hardness for the desired imprecision level. While there have been many impressive results and technical developments about the average-case hardness of Boson Sampling [1, 3, 27–30], the average-case hardness for the desired imprecision level is not yet fully demonstrated. Still, there exists a gap between the imprecision level of average-case hardness in the strongest existing results and the imprecision level of average-case approximation problem. Hence, closing this imprecision gap remains the ultimate challenge for the complete theoretical demonstration of the classical hardness of approximate Boson Sampling.

B The bosonic birthday paradox for the shallow average-case circuits

Throughout our work, we consider only collision-free outcomes in our average-case hardness arguments instead of all possible outcomes. In this appendix, we examine whether the bosonic birthday paradox [1] holds for our average-case circuit ensemble $\mathcal{H}_{\mathcal{A}}\mathcal{P}$. More specifically, we investigate whether collision-free outcomes are observed with high probability in the so-called "dilute" regime, where the number of modes satisfies $M = \Omega(N^2)$.

To claim the classical hardness and quantum computational advantage of Boson Sampling based solely on collision-free outcomes, it must be ensured that these collision-free outcomes occupy at least a not-a-very-small portion of the probability weight over all possible outcomes. To understand why, let us denote the probability weight corresponding to the collision-free outcomes by ε. For the probability weight of collision-free outcomes ε, the "adversarial" classical sampler that outputs only collision outcomes is also a good sampler, if the allowed total variation distance error for the classical sampler is larger than ε. Using this adversarial sampler combined with Stockmeyer's algorithm that multiplicative approximates the output probability of the sampler, the obtained estimates of output probabilities of collision-free outcomes are always "0", since this sampler does not output collision-free outcomes. Therefore, even if we derive the average-case #P-hardness of estimating output probabilities for collision-free outcomes, we cannot solve this #P-hard problem by Stockmeyer's reduction if the given sampler is adversarial; thereby, we cannot derive the classical hardness result of sampling (which allows total variation distance error) from the average-case #P-hardness.

Accordingly, to avoid this issue and claim the classical hardness result using only collision-free outcomes, the bosonic birthday paradox is crucial, which suppresses the occurrence of collision outcomes over the experiments. This property holds for global Haar random circuits in the dilute regime $M = \Omega(N^2)$ [1], but not guaranteed for any other random circuit ensemble, including our circuit ensemble. In fact, collision-free outcomes dominate the outcome space in the regime $M = \Omega(N^2)$ since the portion of collision-free outcomes over all possible outcomes can be bounded as $\binom{M}{N}/\binom{M+N-1}{N} > 1 - \frac{N^2}{M}$. Accordingly, one might expect that the bosonic birthday paradox automatically holds for any circuit ensemble. However, even in this regime, the probability weight might be highly concentrated on collision outcomes for some circuit ensembles (i.e., anticoncentration is not guaranteed for arbitrary circuit ensembles), such that the probability weight of the collision-free outcomes is still too small. Hence, it is crucial to show that the bosonic birthday paradox for our circuit ensemble, such that our circuit ensemble outputs collision-free outcomes with high probability.

Based on this understanding, we now argue the bosonic birthday paradox for the circuit

ensemble $\mathcal{H}_{\mathcal{A}}\mathcal{P}$ in the dilute regime $M = \Omega(N^2)$. We first review a lemma given in Ref. [1, Lemma C.1], which is indeed a crucial ingredient for the bosonic birthday paradox.

Lemma 4 (Unitary Pigeonhole Principle [1])**.** *Partition a finite set $[K]$ into a "good part" G and "bad part" $B = [K]\backslash G$. Also, let $\hat{U}$ be an arbitrary K by K unitary matrix. Suppose we choose an element $x \in G$ uniformly at random, apply $\hat{U}$ to $|x\rangle$, and then measure $\hat{U}\,|x\rangle$ in the standard basis. Then, for measurement outcome y, we have* $\Pr[y \in B] \leq |B|/|G|$.

Here, let us consider the finite set $[K]$ in the lemma as a set of all possible outcomes of Boson Sampling for N photons and M modes, i.e., the set of possible $\boldsymbol{s} = (s_1, \ldots, s_M)$ where each $s_i \in [N]$ with the constraint $\sum_{i=1}^{M} s_i = N$ (such that, $|K| = \binom{M+N-1}{N}$). Also, we consider $\hat{U}$ in the lemma as a unitary evolution over the Hilbert space of N photons and M modes, characterized by a given unitary circuit of Boson Sampling.

We now consider the good part G in the lemma as a set of collision-free $\boldsymbol{s} = (s_1, \ldots, s_M)$, such that each $s_i \in \{0, 1\}$ instead of $s_i \in [N]$, which we will denote as $G_{M,N}$. Likewise, let $B_{M,N}$ be the bad part in the lemma, i.e., a set of collision $\boldsymbol{s}$. Then, observe that $|G_{M,N}| = \binom{M}{N}$, $|B_{M,N}| = \binom{M+N-1}{N} - \binom{M}{N}$, and thus we have

$$\frac{|B_{M,N}|}{|G_{M,N}|} = \frac{\binom{M+N-1}{N} - \binom{M}{N}}{\binom{M}{N}} < \frac{\frac{N^2}{M}}{1 - \frac{N^2}{M}} \leq \frac{2N^2}{M}, \tag{17}$$

where we used the fact that $\binom{M}{N}/\binom{M+N-1}{N} > 1 - \frac{N^2}{M}$ and assumed $M \geq 2N^2$ in the last inequality for simplicity. Therefore, by Lemma 4, if we uniformly choose collision-free input (i.e., $\boldsymbol{t}$ in Eq. (1)), then the probability of observing collision outcomes is bounded by $\frac{2N^2}{M}$, for any unitary circuit.

Now, instead of using uniformly chosen input and fixed unitary circuit, we can rather fix the collision-free input $\boldsymbol{t}$ and draw random unitary circuits from $\mathcal{H}_{\mathcal{A}}\mathcal{P}$ (see Definition 4 and Definition 5), which is the random circuit ensemble for the average-case hardness in the shallow-depth regime. Then, the random permutation circuit part from $\mathcal{P}$ permutes the collision-free input $\boldsymbol{t}$ uniformly random, and effectively makes input uniform distribution over possible collision-free configurations $G_{M,N}$. Hence, by choosing circuits from $\mathcal{H}_{\mathcal{A}}\mathcal{P}$, then the probability to observe collision outcomes is bounded by $\frac{2N^2}{M}$, and thus we have the following theorem.

Theorem 8 (Bosonic birthday paradox for the product circuit ensemble $\mathcal{H}_{\mathcal{A}}\mathcal{P}$)**.** *For a randomly chosen unitary circuit $U \sim \mathcal{H}_{\mathcal{A}}\mathcal{P}$, the probability to obtain collision outcomes is upper bounded by*

$$\mathop{\mathbb{E}}_{U \sim \mathcal{H}_{\mathcal{A}}\mathcal{P}} \left[\Pr_{\mathcal{D}_U} \left[\boldsymbol{s} \in B_{M,N} \right] \right] < \frac{2N^2}{M}, \tag{18}$$

where $\mathcal{D}_U$ is the output distribution according to $p_{\boldsymbol{s}}(U)$ defined in Eq. (1).

We have now obtained the desired result; the above theorem implies that for most of the random circuits over $\mathcal{H}_{\mathcal{A}}\mathcal{P}$, the sampler outputs a collision-free outcome with high probability in the dilute regime $M = \Omega(N^2)$.

Crucially, the proof of Theorem 8 does not rely on the circuit ensemble $\mathcal{H}_{\mathcal{A}}$; rather, the essential ingredient for the bosonic birthday paradox is the initial random permutation layer $\mathcal{P}$. Therefore, the bosonic birthday paradox is guaranteed in the dilute regime $M = \Omega(N^2)$ as long as the as long as the input (collision-free) configuration is uniformly random by any means. In practice, this can also be achieved experimentally by directly preparing

a uniformly random input configuration. Nevertheless, we chose to include an explicit random permutation circuit $\mathcal{P}$ for mathematical clarity, since we fix the input configuration for Boson Sampling.

We additionally note that this bosonic birthday paradox is not guaranteed solely by the local random circuit ensemble in Definition 4 itself, because of the absence of translation invariance property (in contrast to the global Haar random unitary circuit). Hence, the additional random permutation circuit at the input is crucial to guarantee the bosonic birthday paradox mathematically. Nevertheless, if one can prove that the local random circuit ensemble has a permutation-invariance property for a specific circuit architecture, then we can leave out this random permutation circuit part at the input. We leave in Appendix I numerical evidence of the bosonic birthday paradox for solely the local random circuit ensemble $\mathcal{H}_{\mathcal{A}}$ with our shallow-depth circuit architecture $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^q$.

C Proof of Theorem 3 (worst-case hardness)

To prove Theorem 3, we first introduce a crucial lemma, which states that any BQP circuit can be simulated by a constant-depth linear optical circuit under post-selection.

Lemma 5 (Brod [33], revised)**.** *For an arbitrary given poly-sized n-qubit quantum circuit Q, there exists a constant depth linear optical circuit C_0 such that for $M = \Omega(N)$ and $N = \mathrm{poly}(n)$,*

$$|\langle J|\hat{\mathcal{U}}(C_0)|I\rangle|^2 = c_Q|\langle 0|^{\otimes n} Q|0\rangle^{\otimes n}|^2, \tag{19}$$

where $c_Q = 2^{-O(N)}$ is a Q-dependent constant which can be efficiently computed, $\hat{\mathcal{U}}(C_0)$ is a unitary operator corresponding to the circuit C_0, and each of $|I\rangle$ and $|J\rangle$ is an M-mode Fock-state composed of N single photon states and vacuum states for the rest modes.

Proof. In this proof, we revise the main result in [33], incorporating a more quantitative analysis on the coefficient c_Q, which later determines the imprecision level allowed for the worst-case hardness of shallow-depth Boson Sampling. To do so, we first go through the main result in [33].

Given an arbitrary poly-sized quantum circuit Q on n qubits, there exists a measurement-based quantum computation (MBQC) scheme using constant depth brickwork graph state composed of maximally $m_Q = \mathrm{poly}(n)$ number of qubits and $\Theta(m_Q)$ number of CZ gates, with an explicit measurement basis sequence for the universal gate set {CX, T, H} to simulate Q [52–54] (see, e.g., Fig. 4 (a)). Since the measurement outcomes for gate implementation are random, one needs adaptive measurements based on these measurement outcomes to implement Q (which is indeed the essence of MBQC). However, if post-selecting desired measurement outcomes are allowed, one can rather post-select "good" outcomes that do not require any adaptive measurements to implement Q.

This brickwork qubit circuit can also be implemented in linear optical systems via the celebrated KLM scheme [55]. In their scheme, they use a dual-rail encoded state ($|01\rangle$ or $|10\rangle$ in the Fock basis) as a qubit, a beam splitter as a one qubit gate, and a 2-depth network of beam splitters for a probabilistic CZ gate which also requires additional two single photon ancillas for post-selecting successful CZ operation. Using this KLM scheme, Ref. [33] proposed a 4-depth linear optical circuit C_0, which implements the brickwork qubit state by post-selecting successful CZ operation (see Fig. 4 (b) or Fig. 3 in Ref. [33] for more details). Here, the required photon number is $N = \Theta(m_Q)$ to implement the brickwork qubit state composed of m_Q number of qubits and $\Theta(m_Q)$ number of CZ gates, because each of qubit and CZ gate requires a constant number of photons in the KLM scheme. For more clarity, we illustrate in Fig. 4 the brickwork qubit circuit and

corresponding linear optical circuit C_0 proposed by [33]. Also, the mode number is given by $M = \Omega(N)$, which is required for the dual-rail encoding in the KLM scheme. We remark that the mode number M of C_0 can be arbitrarily enlarged by considering a circuit that "contains" the circuit in [33], which uses some of the modes for dual-rail encoding, leaving the rest of the modes as vacuum and adding only trivial (i.e., identity) gates at these unwanted modes. As the trivial gates leave the unwanted modes separated, this kind of mode number augmentation does not affect the post-selection probabilities.

To sum up, given an arbitrary quantum circuit Q, there exists a 4-depth linear optical circuit C_0 that can simulate the quantum circuit Q, by post-selecting successful CZ operation for all probabilistic CZ gates and post-selecting good measurement outcomes (that do not require any adaptive measurements) to implement each gate in the universal gate set {CX, T, H} composing the circuit Q. Here, as the explicit measurement basis sequence of brickwork qubit circuit for the universal gate set {CX, T, H} is given, one can find (i) proper coefficients of beam splitters in C_0 to implement the universal gate set via the KLM scheme and (ii) good measurement outcomes in the Fock basis. Hence, for the given quantum circuit Q, we can efficiently construct the constant-depth circuit C_0 (by means of determining its coefficients) and M-mode N-single photon Fock-states $|I\rangle$ and $|J\rangle$ such that Eq. (19) holds. Now, c_Q becomes a post-selection probability to post-select (i) successful CZ operations for all CZ gates and (ii) good measurement outcomes to implement each gate in the universal gate set composing Q via MBQC. In other words, c_Q can be expressed as

$$c_Q = \prod_{k \in \{\mathrm{CZ},\ \mathrm{CX},\ \mathrm{T},\ \mathrm{H}\}} p_k^{\Gamma_k}, \tag{20}$$

where p_k denotes post-selection probability to implement k gate (e.g., p_{CZ} is $2/27$ in [56]), and Γ_k denotes the number of k gate to implement the circuit Q. Hence, given the post-selection probability p_k for each gate, by counting the number of each gate to implement the circuit Q, c_Q can be computed efficiently. To obtain the size of c_Q, note that we set the photon number $N = \Theta(m_Q)$ to implement m_Q qubit brickwork qubit circuit. As this m_Q qubit brickwork qubit circuit can implement at most $O(m_Q)$ number of gates, Q has at most $O(m_Q) = O(N)$ number of gates in the universal gate set {CX, T, H}. Also, the number of CZ gates is given as $O(N)$, and thus the total number of gates to implement Q is at most $\sum_k \Gamma_k = O(N)$. Therefore, c_Q has its amplitude $c_Q = 2^{-O(N)}$, since all the post-selection probabilities p_k are given as constants. This concludes the proof.

□

Next, we use the proposition by [43] which states that estimating the output probability of any BQP circuit is #P-hard for a certain additive imprecision.

Lemma 6 (Kondo et al [43])**.** *It is #P-hard to compute* $|\langle 0|^{\otimes n} Q |0\rangle^{\otimes n}|^2$ *for an arbitrary given quantum circuit* Q *within the additive error less than* 2^{-2n}*.*

By the above two lemmas, one can deduce that there exists a constant-depth linear optical circuit that has the worst-case hardness for a certain level of imprecision. However, this does not directly imply the worst-case hardness for the fixed circuit architecture $\mathcal{BB}^*$ as stated in Theorem 3 because it is not guaranteed that the constant-depth worst-case circuit satisfying Eq. (19) can be implemented in the fixed circuit architecture $\mathcal{BB}^*$.

Hence, we verify that there exists a constant depth worst-case circuit C_0 in Lemma 5 that can be implemented in the fixed circuit architecture $\mathcal{BB}^*$, by properly allocating the mode indices of C_0 in $\mathcal{BB}^*$.

Lemma 7. *There exists a constant-depth circuit C_0 in Lemma 5 that can be implemented in $\mathcal{BB}^*$ under a proper mode index permutation. In other words, for an arbitrary given poly-sized n-qubit quantum circuit Q, there exists a circuit C_0 in $\mathcal{BB}^*$ that satisfies Eq. (19) for $M = \Omega(N)$ and $N = \mathrm{poly}(n)$, where $c_Q = 2^{-O(N)}$, and $|I\rangle$ and $|J\rangle$ are given by M-mode N-single photon Fock-states.*

Proof. See Appendix D. □

We prove Lemma 7 by showing that the constant-depth linear optical circuit C_0 in Lemma 5 can be embedded into a 3-dimensional grid-structured circuit, which can be implemented in the $\mathcal{BB}^*$ architecture under a suitable permutation of mode indices. For interested readers, we provide the full proof of Lemma 7 in Appendix D.

Combining the above results, the proof of Theorem 3 is now straightforward.

Proof of Theorem 3. Given an arbitrary quantum circuit Q, by Lemma 5, we can efficiently construct a constant-depth linear optical circuit C_0 and Fock-states $|I\rangle$ and $|J\rangle$ such that Eq. (19) holds. Also, as depicted in Lemma 7, we can efficiently embed this linear optical circuit C_0 in $\mathcal{BB}^*$ by properly allocating mode indices. This mode index allocation permutes $|I\rangle$ and $|J\rangle$; let us denote the permuted version of $|I\rangle$ and $|J\rangle$ as $|I'\rangle$ and $|J'\rangle$, respectively. Then, we can construct the linear optical circuit C_0 in $\mathcal{BB}^*$ such that

$$c_Q^{-1} p_{\boldsymbol{s}_0}(C_0) = |\langle 0|^{\otimes n} Q |0\rangle^{\otimes n}|^2, \tag{21}$$

for collision-free input $\boldsymbol{t}$ and output $\boldsymbol{s}_0$ corresponding to $|J'\rangle$ and $|I'\rangle$, respectively. Note that, the input and output configuration $|J'\rangle$ and $|I'\rangle$ are fixed and invariant under the quantum circuit Q, because the architecture of C_0 does not depend on Q (more precisely, only the coefficients of C_0 depend on Q). As $c_Q = 2^{-O(N)}$, if we can estimate $p_{\boldsymbol{s}_0}(C_0)$ to within additive imprecision $2^{-(2n+O(N))} = 2^{-O(N)}$, we can estimate the right-hand side of Eq. (21) to within additive imprecision 2^{-2n}, which is #P-hard as given in Lemma 6. Therefore, for collision-free input $\boldsymbol{t}$ and output $\boldsymbol{s}_0$, estimating an output probability $p_{\boldsymbol{s}_0}(C_0)$ to within additive imprecision $2^{-O(N)}$ for any C_0 in $\mathcal{BB}^*$ is #P-hard in the worst case. □

Remark 2. *In principle, encoding the constant-depth MBQC circuit from Lemma 5 to the logarithmic-depth $\mathcal{BB}^*$ circuit, as done in Lemma 7, is not essential for deriving the hardness argument in shallow-depth regimes; one could directly use the constant-depth circuit from Lemma 5 itself to obtain worst-case hardness in such regimes. Nevertheless, since we employ the architecture $\mathcal{BB}^*$ to implement required permutations during the worst-to-average-case reduction, we derive worst-case hardness within $\mathcal{BB}^*$ to simplify the overall circuit architecture to $(\mathcal{BB}^*)^q$ in Definition 3.*

D Proof of Lemma 7: Encoding constant-depth worst-case circuit in Lemma 5 to the circuit architecture $\mathcal{BB}^*$

In the following, we prove Lemma 7, showing that the constant depth worst-case linear optical circuit C_0 argued in [33] and Lemma 5 can be encoded in the Kaleidoscope circuit architecture $\mathcal{BB}^*$ by properly allocating mode indices. To summarize, this can be done in two steps: First, we show how we can map the worst-case linear optical circuit C_0 into a 3-dimensional "grid" linear optical circuit (i.e., a circuit with only nearest neighbor interactions for each dimension, with open boundaries). Next, we show that a single

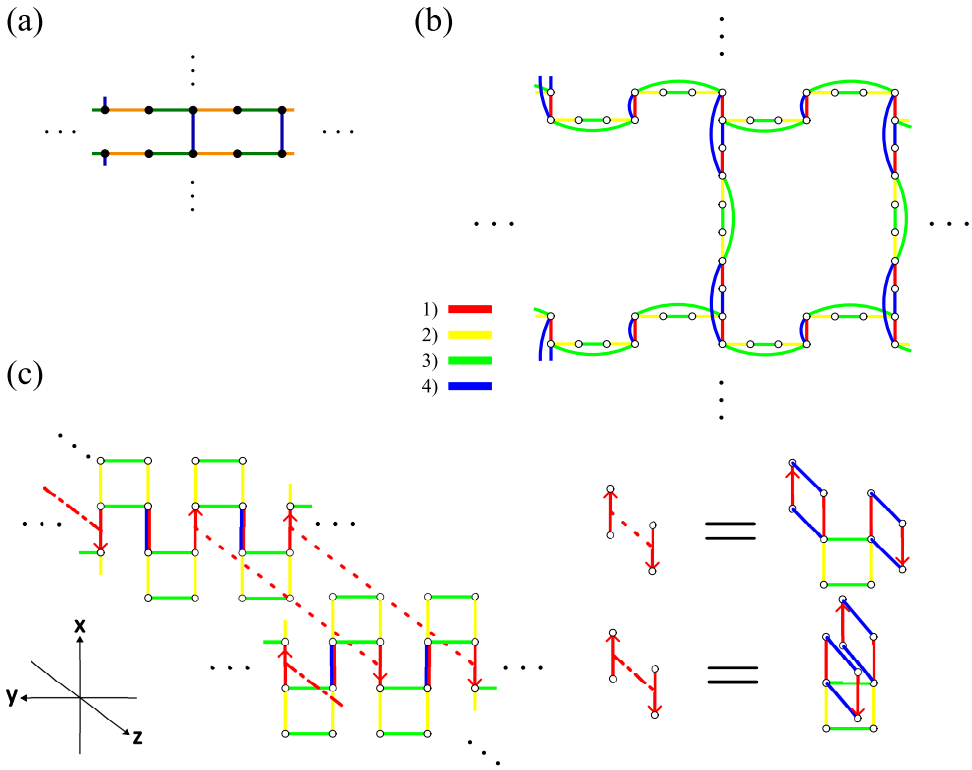

Figure 4: (a) Schematics of the unit brickwork qubit circuit for measurement-based quantum computation [52–54]. Each vertex (i.e., black dot) represents a $|+\rangle$ qubit, and each edge represents the CZ gate applied on two qubits. After CZ gate operations, each vertex is measured with a rotated $|+\rangle$ basis, where the rotation angle is explicitly given for measurement-based quantum computation. (b) Schematics of the 4-depth linear optical circuit proposed by [33] corresponding to the unit brickwork qubit circuit in (a) implemented by the KLM scheme (for more details, see Ref. [33]). Each vertex (i.e., white dot) represents zero or one photon Fock-state, and each edge represents the beam splitter with arbitrarily chosen coefficients. The initial beam splitter layer (red) corresponds to the input qubit (and ancillas), the following two beam splitter layers (yellow and green) correspond to the CZ gate implementation [56], and the last beam splitter layer (blue) corresponds to the qubit rotation required for the measurement-based quantum computation. After the beam splitter operations, each vertex is measured with the Fock-state basis. (c) Schematics of the unit brickwork linear optical circuit in (b) but mapped in 3-dimensional grid linear optical circuit architecture. Here, the beam splitters are applied in parallel along each dimension, where the application sequence is along the direction $x \to y \to z \to x$.

parallel application of any grid-structured linear optical circuit can be implemented in the unit butterfly circuit architecture $\mathcal{B}$ under a mode index permutation. We then combine these arguments, showing that the worst-case circuit C_0 can be encoded in $\mathcal{B}\mathcal{B}^*$.

D.1 Mapping worst-case circuit in Lemma 5 to a grid linear optical circuit

We begin by describing how to map the worst-case linear optical circuit C_0 in Lemma 5 into a 3-dimensional grid-structured linear optical circuit. To do so, we first depict in Fig. 4 (a) the brickwork qubit circuit for the MBQC [52–54], which can simulate any BQP circuit using post-selection. Here, each vertex (i.e., black dot) represents input $|+\rangle$ state, and each edge represents the CZ gate applied on two qubits. After these CZ operations, each vertex is measured with rotated $|+\rangle$ basis, where the rotation angle is explicitly given to implement the universal gate set via MBQC.

We can also construct this brickwork qubit circuit in linear optical systems under

post-selections. As proposed in [33], the brickwork qubit circuit in Fig. 4 (a) can also be implemented in the linear optical circuit by the KLM scheme combined with post-selection, which is indeed the worst-case 4-depth linear optical circuit C_0 in Lemma 5. We illustrate the brickwork of this 4-depth linear optical circuit in Fig. 4 (b). Here, each vertex (i.e., white dot) represents a Fock-state with zero or one photon occupied, and each edge represents the beam splitter with appropriately chosen coefficients. More specifically, the initial beam splitter layer (red) corresponds to the input qubit preparation (and ancillas), the following two beam splitter layers (yellow and green) combined with single photon ancillas correspond to the CZ gate operation [56]. The last beam splitter layer (blue) combined with Fock-state measurement corresponds to the rotated $|+\rangle$ measurement in the KLM scheme, which is required to implement the universal gate set via MBQC.

By appropriately allocating the modes, the worst-case linear optical circuit C_0 depicted in Fig. 4 (b) can be mapped into a 3-dimensional grid linear optical circuit architecture, as depicted in Fig. 4 (c). More specifically, the sequence of beam splitter application in the 3-dimensional architecture to implement the brickwork linear optical circuit is as follows. At first, the red- and the yellow-colored beam splitters are applied in parallel along the x-axis. Next, the greed-colored beam splitters are applied in parallel along the y-axis. Lastly, the blue-colored beam splitters are applied in parallel along the z-axis and then applied in parallel along the x-axis again. Let us consider a single parallel application of a grid linear optical circuit as a sequence of parallel applications of beam splitters along each dimension. Then, one can deduce that the brickwork linear optical circuit in Fig. 4 (b) can be implemented in a single parallel application of 3-dimensional grid linear optical circuit from x- to z-axis, with additional parallel beam splitter application along x-axis. Here, we note that the size of this grid linear optical circuit can be arbitrarily enlarged as long as it contains the desired circuit in Fig. 4 (c), where we can leave the unused modes as vacuums and add only trivial gates at these unused modes. Also, as the trivial gates leave the unused modes separated, this mode number augmentation does not affect the post-selection probabilities of MBQC and KLM scheme, thus leaving the total post-selection probability c_Q given in Lemma 5 unchanged.

D.2 Encoding a grid linear optical circuit to the butterfly circuit architecture $\mathcal{B}$

Second, we show that a single parallel application of a general-dimensional grid-structured linear optical circuit can be implemented in the unit butterfly circuit architecture $\mathcal{B}$, under a proper mode index permutation. Here, as we mentioned previously, a single parallel application of a grid linear optical circuit is a sequence of parallel applications of beam splitters along each dimension. More generally, we prove the following theorem.

Theorem 9. *Suppose that a mode number is given by $M = 2^n$ for $n \in \mathbb{Z}^+$. For an arbitrary M-mode single parallel d-dimensional grid linear optical circuit C with its size $2^{n_1} \times 2^{n_2} \times \cdots \times 2^{n_d}$ for $\sum_{i=1}^{d} n_i = n$, there exist an M-mode circuit C' in $\mathcal{B}$ and an M-mode permutation matrix $\boldsymbol{P}$ that satisfy $C = \boldsymbol{P}C'\boldsymbol{P}^T$*

To prove Theorem 9 we begin by showing how to implement a single parallel 1-dimensional grid circuit (i.e., a 1-dimensional chain, with open boundaries) by $\mathcal{B}$, whose mode number is given by $M = 2^n$. Using this result, we later prove Theorem 9 by extending this result to the general-dimensional case.

Lemma 8. *Suppose that a mode number is given by $M = 2^n$ for $n \in \mathbb{Z}^+$. For an M-mode single parallel 1-dimensional grid linear optical circuit C, there exist an M-mode circuit C' in $\mathcal{B}$ and an M-mode permutation matrix $\boldsymbol{P}$ that satisfy $C = \boldsymbol{P}C'\boldsymbol{P}^T$*

Proof. At first, we constructively establish the permutation matrix $\boldsymbol{P}$ to implement the grid circuit from the circuit in $\mathcal{B}$. Since $\mathcal{B}$ is composed of non-local interactions except for the first depth, we need a proper mode permutation to effectively make the grid structure. In short, this can be done as follows: As we previously defined in Definition 2, for the circuit architecture $\mathcal{B}$, we set the depth index $L \in [n]$ and the corresponding partition index $j \in [2^{n-L}]$ given by integers. Then, in the ascending sequence of j and L, permute the mode indices $[\![2^L(j-1)+2^{L-1}+1, 2^L(j-1)+2^L]\!]$ from the ascending to the descending order, where the square bracket $[\![a,b]\!]$ indicates the integer interval between a and b included. In other words, for each $L \in [n]$ and $j \in [2^{n-L}]$, we permute

$$2^L(j-1)+2^{L-1}+k \;\rightarrow\; 2^L(j-1)+2^L+1-k, \tag{22}$$

for all $k \in \left[2^{L-1}\right]$ (nothing changes in Eq. (22) for $L=1$ so permutation does not occur for $L=1$ case). Let us denote $\boldsymbol{P}_{L,j}$ as the permutation described above, acting non-trivially on the mode indices $[\![2^L(j-1)+2^{L-1}+1, 2^L(j-1)+2^L]\!]$, and acting trivially on the rest modes. Then, we can construct $\boldsymbol{P}$ as $\boldsymbol{P} = \prod_{L=2}^{n}\prod_{j=1}^{2^{n-L}} \boldsymbol{P}_{L,j}$.

Let us explain why this permutation $\boldsymbol{P}$ can make the grid circuit from the circuit in $\mathcal{B}$. First, suppose that the following is true: For given $L \geq 2$ and j, each of the sub-circuit of depth $L-1$ with mode indices $[\![2^{L-1}(2j-2)+1, 2^{L-1}(2j-2)+2^{L-1}]\!]$ and $[\![2^{L-1}(2j-1)+1, 2^{L-1}(2j-1)+2^{L-1}]\!]$ forms a grid circuit with mode number 2^{L-1}; note that, this is trivial for $L=2$ case. By the structure of $\mathcal{B}$, for given $L \geq 2$ and j, initially $(2^L(j-1)+2^{L-1})$-th mode and $(2^L(j-1)+2^L)$-th mode are connected by a gate (i.e., a beam splitter) at depth L. Then, after the ascending-to-descending permutation $\boldsymbol{P}_{L,j}$ for given $L \geq 2$ and j, now $(2^L(j-1)+2^{L-1})$-th mode and $(2^L(j-1)+2^{L-1}+1)$-th mode are connected by the gate. Hence, by setting this gate as a non-trivial gate composing the grid circuit and setting the rest of the gates as trivial gates, we now have a grid circuit with size 2^L for each $L \geq 2$ and j, with mode indices given by $[\![2^L(j-1)+1, 2^L(j-1)+2^L]\!]$. Moreover, because we permute equivalently for all $j \in \left[2^{n-L}\right]$ for given L, these permutations do not affect the structure of $\mathcal{B}$ for larger depth $L+1, L+2, \ldots, n$, such that after the permutations, $(2^{L+1}(j-1)+2^L)$-th mode and $(2^{L+1}(j-1)+2^{L+1})$-th mode are still connected by a gate at depth $L+1$. Therefore, we can inductively repeat this permutation for ascending j and L, thereby obtaining the permutation $\boldsymbol{P} = \prod_{L=2}^{n}\prod_{j=1}^{2^{n-L}} \boldsymbol{P}_{L,j}$ that can implement the grid circuit of size 2^n.

Next, we construct the circuit C' in $\mathcal{B}$. This can be done by choosing the non-trivial gate appropriately in $\mathcal{B}$ and leaving the rest of the gates trivial gates for each depth L and partition j. Firstly, set all the gates in $L=1$ in $\mathcal{B}$ non-trivial gates because they already form a grid structure. Also, as we previously depicted, for each $L \geq 2$ and j, we set the gate between $(2^L(j-1)+2^{L-1})$-th mode and $(2^L(j-1)+2^L)$-th mode as the non-trivial gate, where these modes are permuted from $(2^L(j-1)+2^{L-2}+1)$-th mode and $(2^L(j-1)+2^{L-1}+2^{L-2}+1)$-th mode respectively in the previous step. Hence, given the circuit architecture $\mathcal{B}$, for each $L \geq 2$ and j, we set the gate between $(2^L(j-1)+2^{L-2}+1)$-th mode and $(2^L(j-1)+2^{L-1}+2^{L-2}+1)$-th as the only non-trivial gate, and leave the rest of the gates as trivial gates. For a given single parallel 1-dimensional grid linear optical circuit C, let us denote $\{G_{s,t}\}$ as a set of two-mode gates composing the circuit C, where each gate $G_{s,t}$ is applied on s-th and t-th modes, respectively (for 1-dimensional case, t is given by $s+1$). Because $(2^L(j-1)+2^{L-1})$-th mode and $(2^L(j-1)+2^{L-1}+1)$-th mode are connected by the non-trivial gate after the permutation for each L and j, we set this non-trivial gate as $G_{s,t}$ for $s = 2^L(j-1)+2^{L-1}$ and $t = s+1$, thereby implementing the grid circuit C by repeating this process in the ascending order of j and L.

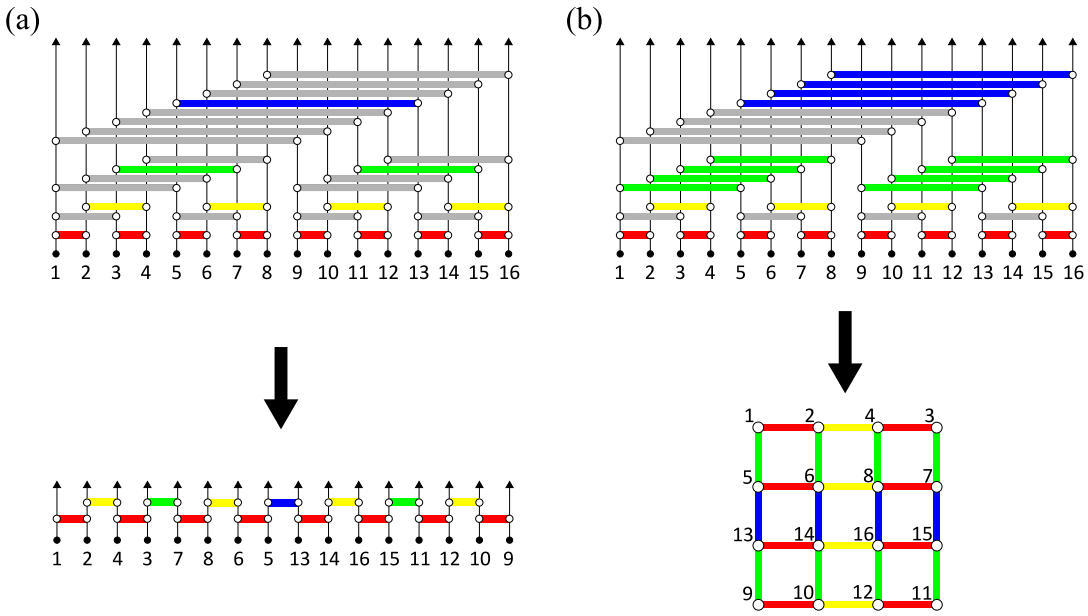

Figure 5: Schematics of $\mathcal{B}$ constructing (a) 1-dimensional grid linear optical circuit and (b) 2-dimensional grid linear optical circuit under a proper mode index permutation, for mode number $M = 16$. The grey-colored edges indicate trivial (i.e., identity) gates, and the other colored edges indicate non-trivial gates that compose the grid linear optical circuit we aim to construct.

Therefore, we can construct the circuit C' in $\mathcal{B}$ as follows. First, for depth $L = 1$ and $j \in [2^{n-1}]$, we set the gate between the modes $2^L(j-1)+2^{L-1}$ and $2^L(j-1)+2^L$ in $\mathcal{B}$ as $G_{s,t}$ for $s = 2^L(j-1)+2^{L-1}$ and $t = s+1$. Next, for each depth $L \in [\![2, n]\!]$ and $j \in [2^{n-L}]$, we set the gate between the modes $2^L(j-1)+2^{L-2}+1$ and $2^L(j-1)+2^{L-1}+2^{L-2}+1$ in $\mathcal{B}$ as $G_{s,t}$ for $s = 2^L(j-1)+2^{L-1}$ and $t = s+1$, and leave trivial for the rest of the gates. Combining the arguments so far, this circuit C' permuted by the proper permutation matrix $\boldsymbol{P}$ described above finally gives the desired grid circuit $C = \boldsymbol{P}C'\boldsymbol{P}^T$ (see, e.g., Fig. 5 (a)).

□

Based on this understanding, we now prove Theorem 9, by extending the analysis in Lemma 8 to the general-dimensional cases.

Proof of Theorem 9. We first emphasize that, by decomposing the depth n of $\mathcal{B}$ with n_i such that $\sum_{i=1}^d n_i = n$, $\mathcal{B}$ can also be represented as d-dimensional circuit with its size 2^{n_i} along the i-th dimension; one can consider the gates of $\mathcal{B}$ between the depth $L \in [\![\sum_{p=1}^{i-1} n_p + 1, \sum_{p=1}^{i} n_p]\!]$ are aligned along i-th dimension for $i \geq 2$, where the square bracket $[\![a, b]\!]$ indicates the integer interval between a and b included. Hence, we naturally extend the analysis in Lemma 8, implementing a grid circuit along each dimension with size 2^{n_i}, for all possible $n_i \in \{n_1, n_2, \ldots, n_d\}$. Here, we assume that $n_i \geq 2$ for all $i \in [d]$ without loss of generality because $\mathcal{B}$ already forms a grid along i-th dimension when $n_i = 1$.

More specifically, we set the size index for each dimension as $n_i \in \{n_1, n_2, \ldots, n_d\}$, along with the depth index $L \in [n]$ and the corresponding partition index $j \in \left[2^{n-L}\right]$ as defined previously. Also, let us denote $\boldsymbol{P}^{(i)}$ as the permutation along the i-th dimension, and $\{G_{s,t}^{(i)}\}$ as the set of gates applied on s-th and t-th modes along the i-th dimension composing the circuit C. Then, along the first dimension up to $L = n_1$, we can similarly

construct the permutation $\boldsymbol{P}^{(1)}$ and encode the gates $\{G_{s,t}^{(1)}\}$ into $C' \in \mathcal{B}$ as depicted in Lemma 8.

In the ascending order of the dimension $i \geq 2$, we note that each mode is departed by $2^{\sum_{p=1}^{i-1} n_p}$ mode index from its nearest-neighborhood along the i-th dimension. By the structure of $\mathcal{B}$, for given $L \in [\![\sum_{p=1}^{i-1} n_p + 1, \sum_{p=1}^{i} n_p]\!]$ and $j \in \left[2^{n-L}\right]$, initially $(2^L(j-1) + 2^{L-1} - 2^{\sum_{p=1}^{i-1} n_p} + l)$-th mode and $(2^L(j-1) + 2^L - 2^{\sum_{p=1}^{i-1} n_p} + l)$-th mode are connected by gates at L for all $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$, where l characterizes the indices along all the previous dimensions. Hence, now our permutation rule is that for each $L \in [\![\sum_{p=1}^{i-1} n_p + 1, \sum_{p=1}^{i} n_p]\!]$ and $j \in \left[2^{n-L}\right]$, we permute mode indices by

$$2^L(j-1) + 2^{L-1} + (k-1)2^{\sum_{p=1}^{i-1} n_p} + l \;\rightarrow\; 2^L(j-1) + 2^L - k2^{\sum_{p=1}^{i-1} n_p} + l, \tag{23}$$

for all $k \in \left[2^{L-1-\sum_{p=1}^{i-1} n_p}\right]$ and $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$. Indeed, this is the ascending-to-descending permutation similar to Eq. (22) but now along the i-th dimension instead. Note that nothing happens in Eq. (23) for $L = \sum_{p=1}^{i-1} n_p + 1$ so permutation does not occur for $L = \sum_{p=1}^{i-1} n_p + 1$ cases. After these permutations, $(2^L(j-1) + 2^{L-1} - 2^{\sum_{p=1}^{i-1} n_p} + l)$-th mode and $(2^L(j-1) + 2^{L-1} + l)$-th mode are connected by the gates for all $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$. Accordingly, we set these gates as non-trivial gates and set the rest of the gates as trivial gates for each L and j, thereby forming a grid structure up to depth L. Also, notice that the permutation in Eq. (23) can be done parallelly for all $k \in \left[\!\left[1, 2^{L-1-\sum_{p=1}^{i-1} n_p}\right]\!\right]$ and $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$. Therefore, by denoting the permutation described in Eq. (23) for all k and x as $\boldsymbol{P}_{L,j}^{(i)}$, we can construct $\boldsymbol{P}^{(i)}$ as

$$\boldsymbol{P}^{(i)} = \prod_{L=\sum_{p=1}^{i-1} n_p+2}^{\sum_{p=1}^{i} n_p} \prod_{j=1}^{2^{n-L}} \boldsymbol{P}_{L,j}^{(i)}, \tag{24}$$

and accordingly, $\boldsymbol{P} = \prod_{i=1}^{d} \boldsymbol{P}^{(i)}$.

We now construct the circuit C' in $\mathcal{B}$ for $i \geq 2$. In the ascending order of $i \geq 2$, for each i, set all the gates in $L = \sum_{p=1}^{i-1} n_p + 1$ non-trivial gates at first, because they already form a grid structure. Also, as we previously discussed, for each $L \in [\![\sum_{p=1}^{i-1} n_p + 2, \sum_{p=1}^{i} n_p]\!]$ and $j \in \left[2^{n-L}\right]$, we set the gates between $(2^L(j-1) + 2^{L-1} - 2^{\sum_{p=1}^{i-1} n_p} + l)$-th mode and $(2^L(j-1) + 2^L - 2^{\sum_{p=1}^{i-1} n_p} + l)$-th mode as the non-trivial gates for all $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$, where these modes are permuted from $(2^L(j-1) + 2^{L-2} + l)$-th mode and $(2^L(j-1) + 2^{L-1} + 2^{L-2} + l)$-th mode respectively in the previous step. Hence, given the circuit architecture $\mathcal{B}$, for each $L \in [\![\sum_{p=1}^{i-1} n_p + 2, \sum_{p=1}^{i} n_p]\!]$ and $j \in \left[2^{n-L}\right]$, we set the gates between $(2^L(j-1) + 2^{L-2} + l)$-th mode and $(2^L(j-1) + 2^{L-1} + 2^{L-2} + l)$-th mode for all $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$ as the only non-trivial gates, and leave the rest of the gates as trivial gates. To specify these non-trivial gates from the gate set $\{G_{s,t}^{(i)}\}$, note that $(2^L(j-1) + 2^{L-1} - 2^{\sum_{p=1}^{i-1} n_p} + l)$-th mode and $(2^L(j-1) + 2^{L-1} + l)$-th mode are connected by the non-trivial gates for all $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$

after the permutation for each $L \in [\![\sum_{p=1}^{i-1} n_p + 2, \sum_{p=1}^{i} n_p]\!]$ and $j \in \left[2^{n-L}\right]$. We also note that the indices along previous dimensions (characterized by l) were permuted by $\prod_{p=1}^{i-1} \boldsymbol{P}^{(p)}$ in the previous dimensions (working on $2^{\sum_{p=1}^{i-1} n_p}$ mode partition). Therefore, to correct these permutations along previous dimensions, we set these non-trivial gates along $(2^L(j-1) + 2^{L-1} - 2^{\sum_{p=1}^{i-1} n_p} + l)$-th mode and $(2^L(j-1) + 2^{L-1} + l)$-th mode (after the permutation) as $G_{s,t}^{(i)}$ for $s = 2^L(j-1) + 2^{L-1} - 2^{\sum_{p=1}^{i-1} n_p} + l'$ and $t = s + 2^{\sum_{p=1}^{i-1} n_p}$, where l' is a permuted version of $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$ by $\prod_{p=1}^{i-1} \boldsymbol{P}^{(p)}$.

To sum up, we can construct the circuit C' in $\mathcal{B}$ as follows.

1. Along the first dimension $i = 1$, for depth $L = 1$ and $j \in \left[2^{n-1}\right]$, we set the gate between the modes $2^L(j-1) + 2^{L-1}$ and $2^L(j-1) + 2^L$ in $\mathcal{B}$ as $G_{s,t}^{(1)}$ for $s = 2^L(j-1) + 2^{L-1}$ and $t = s + 1$. Then, for each $L = [\![2, n_1]\!]$ and $j \in \left[2^{n-L}\right]$, we set the gate between the modes $2^L(j-1) + 2^{L-2} + 1$ and $2^L(j-1) + 2^{L-1} + 2^{L-2} + 1$ in $\mathcal{B}$ as $G_{s,t}^{(1)}$ for $s = 2^L(j-1) + 2^{L-1}$ and $t = s + 1$, and leave trivial for the rest of the gates.

2. In the ascending order of the dimension $i \geq 2$, for $L = \sum_{p=1}^{i-1} n_p + 1$ and $j \in \left[2^{n-L}\right]$, we set the gate between the modes $2^L(j-1) + 2^{L-1} - 2^{\sum_{p=1}^{i-1} n_p} + l$ and $2^L(j-1) + 2^L - 2^{\sum_{p=1}^{i-1} n_p} + l$ in $\mathcal{B}$ as $G_{s,t}^{(i)}$ for $s = 2^L(j-1) + 2^{L-1} - 2^{\sum_{p=1}^{i-1} n_p} + l'$ and $t = s + 2^{\sum_{p=1}^{i-1} n_p}$, for all $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$ and l' permuted from l by $\prod_{p=1}^{i-1} \boldsymbol{P}^{(p)}$. Then, for each $L \in [\![\sum_{p=1}^{i-1} n_p + 2, \sum_{p=1}^{i} n_p]\!]$ and $j \in \left[2^{n-L}\right]$, we set the gates between the modes $2^L(j-1) + 2^{L-2} + l$ and $2^L(j-1) + 2^{L-1} + 2^{L-2} + l$ in $\mathcal{B}$ as $G_{s,t}^{(i)}$ for $s = 2^L(j-1) + 2^{L-1} - 2^{\sum_{p=1}^{i-1} n_p} + l'$ and $t = s + 2^{\sum_{p=1}^{i-1} n_p}$, for all $l \in \left[2^{\sum_{p=1}^{i-1} n_p}\right]$ and l' permuted from l by $\prod_{p=1}^{i-1} \boldsymbol{P}^{(p)}$.

Combining the arguments so far, this circuit C' permuted by the proper permutation matrix $\boldsymbol{P}$ described above finally gives the desired grid circuit $C = \boldsymbol{P}C'\boldsymbol{P}^T$ (see, e.g., Fig. 5 (b)).

□

To help understand the encoding scheme in the proof, we leave some illustrations in Fig 5 for 16-mode circuit architecture $\mathcal{B}$ constructing 1- and 2-dimensional grid linear optical circuits using a proper mode index permutation.

D.3 Summary: Encoding the worst-case circuit in Lemma 5 to $\mathcal{B}\mathcal{B}^*$

To sum up, the worst-case linear optical circuit C_0 in Fig. 4 (b) can be implemented in a single parallel application of 3-dimensional grid linear optical circuit from x- to z-axis, with additional parallel beam splitter application along x-axis. Here, as we stated, the size of the grid linear optical circuit can be arbitrarily enlarged as long as it contains the worst-case circuit C_0, which leaves the total post-selection probability c_Q given in Lemma 5 unchanged. Also, by Theorem 9, one can implement a single parallel application of the 3-dimensional grid linear optical circuit in the circuit architecture $\mathcal{B}$ under a mode index permutation. Combining these arguments, the worst-case circuit C_0 in Lemma 5 can be encoded in $\mathcal{B}^*\mathcal{B}$ under a mode index permutation. Here, note that $\mathcal{B}\mathcal{B}^*$ and $\mathcal{B}^*\mathcal{B}$ are equivalent under

a mode index permutation by reverse-indexing each binary number of mode index (e.g., $10010_2 \to 01001_2$, for mode indices given by $[00000_2, 11111_2]$). Therefore, this implies that C_0 can be encoded in $\mathcal{BB}^*$ under a mode index permutation, thereby finally concluding the proof of Lemma 7.

We remark that our encoding scheme to implement the worst-case circuit C_0 in the circuit architecture $\mathcal{BB}^*$ is suboptimal. Since the worst-case circuit can be implemented in constant depth as depicted in Fig. 4 (b), there would possibly exist easier or more straightforward strategies to encode the worst-case circuit in the log-depth circuit architecture $\mathcal{BB}^*$ by allocating mode indices more carefully. Still, our analysis in this appendix is sufficient for our worst-case hardness result in our current work, i.e., as depicted in Theorem 3.

E About Cayley transform

As we depicted in the main text, our goal is to perturb the average-case circuit with the worst-case circuit, and it is crucial to choose a proper circuit perturbation method to establish the worst-to-average-case reduction successfully. Throughout this work, we use the Cayley path for the perturbation, which was employed in Refs. [28, 42, 43] for the hardness proposals of the random circuit sampling.

Definition 8 (Cayley transform [42])**.** *The Cayley transform of an n by n unitary matrix H parameterized by $\theta \in [0,1]$ is a unitary matrix defined as*

$$H(\theta) := ((2-\theta)H + \theta I_n)(\theta H + (2-\theta)I_n)^{-1}, \tag{25}$$

where I_n is the n by n identity matrix. Also, for the diagonalization of the n by n unitary matrix $H = LDL^\dagger$, with unitary matrix L and diagonal matrix $D = diag(e^{i\phi_1}, \ldots, e^{i\phi_n})$, the equivalent form of the Cayley transform is

$$H(\theta) = \frac{1}{q(\theta)} L \; diag(\{p_j(\theta)\}_{j=1}^n) \, L^\dagger, \tag{26}$$

where

$$q(\theta) = \prod_{j=1}^{n} (1 + i\theta e^{i\frac{\phi_j}{2}} \sin\frac{\phi_j}{2}), \tag{27}$$

$$p_j(\theta) = e^{i\phi_j}(1 - i\theta e^{-i\frac{\phi_j}{2}} \sin\frac{\phi_j}{2}) \prod_{k\in[n]\setminus j} (1 + i\theta e^{i\frac{\phi_k}{2}} \sin\frac{\phi_k}{2}). \tag{28}$$

Using the Cayley transform defined above, we now define the perturbed random circuit distribution, which is the local random circuit ensemble in Definition 4 with each Haar random gate perturbed by the Cayley transform.

Definition 9 (Perturbed local random circuit ensemble)**.** *Let $\mathcal{A}$ be the circuit architecture with m number of gates. For the given circuit C in $\mathcal{A}$ with gates $\{G_i\}_{i=1}^m$, the circuit $V(\theta)$ is defined with each gate of C replaced by $G_i \to H_i(\theta)G_i$, where each $H_i(\theta)$ is a Cayley transform of independently distributed local Haar random gate H_i ($i \in [m]$) parameterized by $\theta \in [0,1]$ in Definition 8. We define $\mathcal{H}^C_{\mathcal{A},\theta}$ as the distribution for such $V(\theta)$. Here, the distribution of the $V(0)$ is the local random circuit ensemble $\mathcal{H}_\mathcal{A}$ in Definition 4, and $V(1) = C$.*

We should make sure that the success probability of average-case approximation over circuits is still large enough after the perturbation described in Definition 9, to establish the

reduction process successfully. This is evident in the case that the total variation distance over circuits induced by the perturbation is small enough, as the success probability over circuits perturbs by, at most, the total variation distance. In fact, Ref. [42] proved that total variation distance between $\mathcal{H}_{\mathcal{A}}$ and $\mathcal{H}^{C}_{\mathcal{A},\theta}$ is small for sufficiently small perturbation θ compared to the inverse gate number of the circuit architecture $\mathcal{A}$.

Lemma 9 (Movassagh [42])**.** *Let $\mathcal{A}$ be the circuit architecture with m number of gates. For $\theta \ll 1$ and for any circuit C in $\mathcal{A}$, total variation distance between $\mathcal{H}^{C}_{\mathcal{A},\theta}$ and $\mathcal{H}_{\mathcal{A}}$ is $O(m\theta)$.*

Therefore, by using small $\theta = O(m^{-1})$, one can upper-bound the total variation distance by an arbitrarily small constant, which implies that the success probability of average-case approximation over circuits (i.e., $1-\delta$ in Theorem 4) also perturbs by at most a small constant.

As we have sketched in Sec. 4, we infer the worst-case output probability $p_{\boldsymbol{s}}(C\boldsymbol{P})$ for some fixed circuit C and permutation matrix $\boldsymbol{P}$ via average-case output probabilities $p_{\boldsymbol{s}}(V(\theta)\boldsymbol{P})$ for perturbed random circuit $V(\theta)$. To investigate the viability of the inference of the worst-case value, we examine the behavior of the function $p_{\boldsymbol{s}}(V(\theta)\boldsymbol{P})$ characterized by the parameter θ. Using Definition 8, we find that $p_{\boldsymbol{s}}(V(\theta)\boldsymbol{P})$ can be represented as a rational function in θ.

Lemma 10. *Let $\mathcal{A}$ be the q-Kaleidoscope circuit architecture $(\mathcal{B}\mathcal{B}^*)^q$ with $m = qM\log M$ number of gates, and $V(\theta) \sim \mathcal{H}^{C}_{\mathcal{A},\theta}$ for any C in $\mathcal{A}$. Then for any collision-free outcome $\boldsymbol{s}$ and M by M permutation $\boldsymbol{P}$, the output probability $p_{\boldsymbol{s}}(V(\theta)\boldsymbol{P})$ can be represented as a degree $(4mN, 4mN)$ rational funtion in θ.*

Proof. For given circuit unitary matrix $V(\theta) \sim \mathcal{H}^{C}_{\mathcal{A},\theta}$ with C composed of $\{G_i\}_{i=1}^{m}$ gates, one can decompose $V(\theta)$ with $m = qM\log M$ product of unitary matrices, such that each matrix element of $V(\theta)$ can be represented as

$$[V(\theta)]_{j,k} = \sum_{l_1=1}^{M}\sum_{l_2=1}^{M}\cdots\sum_{l_{m-1}=1}^{M} V^{(1)}_{j,l_1}V^{(2)}_{l_1,l_2}\cdots V^{(m)}_{l_{m-1},k}, \tag{29}$$

where each $V^{(i)}$ denotes an M-dimensional unitary matrix, with a single gate unitary matrix $H_i(\theta)G_i$ applied to the modes participating in the gate and identity for the rest of the modes. For example, if the ith gate $H_i(\theta)G_i$ is a two-mode gate between the first two modes, $V^{(i)}$ is a block diagonal matrix of $H_i(\theta)G_i$ and identity matrix, namely, $V^{(i)} = H_i(\theta)G_i \bigoplus I_{M-2}$.

For circuit architecture $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^q$ which is composed of only two-mode gates, matrix elements of $V^{(i)}$ can be represented as degree $(2,2)$ rational functions in θ, where the common denominator for the elements is given by $q_i(\theta)$, defined in Eq. (27) but with index i appended for ith random gate $H_i(\theta)$. Using reduction to the common denominator for all of the m gates, $[V(\theta)]_{j,k}$ can be represented as $(2m, 2m)$ rational function in θ with the common denominator $\prod_{i=1}^{m} q_i(\theta)$; note that it does not change with the indices j,k.

Now, by Eq. (1), the output probability $p_{\boldsymbol{s}}(V(\theta)\boldsymbol{P})$ for collision-free outcome $\boldsymbol{s}$ has the form of

$$p_{\boldsymbol{s}}(V(\theta)\boldsymbol{P}) = \left|\sum_{\sigma\in S_N}\prod_{j=1}^{N}[V(\theta)_{\boldsymbol{s},\boldsymbol{t}^*}]_{\sigma_j,j}\right|^2, \tag{30}$$

where we denote $\boldsymbol{t}^*$ as an input configuration vector permuted by $\boldsymbol{P}^{-1}$ from the fixed input configuration $\boldsymbol{t}$ in Eq. (1), and S_N is N-mode permutation group. One can easily

check that the common denominator for $\prod_{j=1}^{N}[V(\theta)_{\boldsymbol{s},\boldsymbol{t}^*}]_{\sigma_j,j}$ is $[\prod_{i=1}^{m} q_i(\theta)]^N$, and it does not change with permutation σ. Let us define $Q(\theta) = [\prod_{i=1}^{m} |q_i(\theta)|^2]^N$, which is a degree $4mN$ polynomial in θ. Then, $Q(\theta)$ serves as the common denominator for the output probability. Hence, the output probability can be represented as $p_{\boldsymbol{s}}(V(\theta)) = \frac{P(\theta)}{Q(\theta)}$, with $P(\theta)$ also a degree $4mN$ polynomial in θ. □

F Proof of Theorem 4 (average-case hardness)

In the following, we provide a proof of Theorem 4. As we have already sketched the proof in the main text, for the sake of completeness, we will provide technical details here. Specifically, we establish the worst-to-average-case reduction: We prove that, for δ, $\eta \geq 0$ with $\delta + \eta < \frac{1}{4}$, the output probability $p_{\boldsymbol{s}_0}(C_0)$ in Theorem 3 can be estimated in additive imprecision $2^{-O(N)}$ in complexity class $\mathrm{BPP}^{\mathrm{NP}}$, given output probability estimations $p_{\boldsymbol{s}}(U)$ within additive imprecision $\epsilon = 2^{-O(N^{\gamma+1}(\log N)^2)}$, with probability at least $1-\delta$ over $U \sim \mathcal{H}_{\mathcal{A}}\mathcal{P}$ with $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^{q\geq 3}$ for at least $1-\eta$ over $\boldsymbol{s} \sim \mathcal{G}_{M,N}$.

Proof of Theorem 4. Let $\mathcal{O}$ be an oracle that solves the problem in Theorem 4, i.e., on input $\boldsymbol{s} \sim \mathcal{G}_{M,N}$ and $U \sim \mathcal{H}_{\mathcal{A}}\mathcal{P}$ with $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^{q\geq 3}$, the oracle outputs $p_{\boldsymbol{s}}(U)$ within additive error ϵ with $1-\delta$ portion of U for $1-\eta$ portion of $\boldsymbol{s}$. In other words, for $1-\eta$ of $\boldsymbol{s} \sim \mathcal{G}_{M,N}$ and $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^{q\geq 3}$, we have

$$\Pr_{U\sim\mathcal{H}_{\mathcal{A}}\mathcal{P}}[|\mathcal{O}(\boldsymbol{s},U) - p_{\boldsymbol{s}}(U)| > \epsilon] < \delta. \tag{31}$$

Let C_0 be the worst-case circuit in $(\mathcal{B}\mathcal{B}^*)^{q_0}$ with $q_0 \geq 1$, and let $\boldsymbol{s}_0$ be the fixed collision-free output given in Theorem 3. In the following, we show that approximating $p_{\boldsymbol{s}_0}(C_0)$ to within additive error $2^{-O(N)}$ (i.e., solving the problem in Theorem 3) is in $\mathrm{BPP}^{\mathrm{NP}^{\mathcal{O}}}$, which implies that the average-case approximation of $p_{\boldsymbol{s}}(U)$ to within ϵ is #P-hard under $\mathrm{BPP}^{\mathrm{NP}}$ reduction.

As we have argued in the main text, we first prepare M by M permutation matrices $\boldsymbol{P}_0$ and $\boldsymbol{P}_1$ each sampled uniformly over the $M!$-number of possible permutations. Using these sampled permutations, we newly define permuted outcome $\boldsymbol{s}_p$ and permuted circuit C_p as

$$\boldsymbol{s}_p \coloneqq \boldsymbol{P}_0\boldsymbol{s}_0, \quad C_p \coloneqq \boldsymbol{P}_0 C_0 \boldsymbol{P}_1^{-1}, \tag{32}$$

where one can easily check that the permuted outcome follows random outcome ensemble $\boldsymbol{s}_p \sim \mathcal{G}_{M,N}$ in Definition 6. Also, the permuted circuit C_p can be implemented in the circuit architecture $(\mathcal{B}\mathcal{B}^*)^{q_0+2}$, because each of the sampled $\boldsymbol{P}_0$ and $\boldsymbol{P}_1$ can be implemented in the architecture $\mathcal{B}\mathcal{B}^*$ by Lemma 1.

Note that, by definition in Eq. (32), the worst-case output probability can also be represented as $p_{\boldsymbol{s}_0}(C_0) = p_{\boldsymbol{s}_p}(\boldsymbol{P}_0 C_0) = p_{\boldsymbol{s}_p}(C_p\boldsymbol{P}_1)$. Hereafter, we consider C_p as a "revised" worst-case circuit, and conduct circuit level worst-to-average-case reduction, from the worst-case circuit C_p to the average-case circuit in $(\mathcal{B}\mathcal{B}^*)^{q_0+2}$.

For the average-case circuit, we sample the perturbed random circuit $V(\theta)$ from $\mathcal{H}_{\mathcal{A},\theta}^{C_p}$ defined in Definition 9, for the circuit architecture $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^q$ with $q = q_0 + 2 \geq 3$ from now on, and the worst-case circuit C_p. This can be done by sampling independently distributed local Haar random gates $\{H_i\}_{i=1}^m$ for gate number $m = qM\log M$, perturbing them by the Cayley transform parameterized by θ according to the gates composing the worst-case circuit C_p. Then, by Definition 9, $V(1) = C_p$ such that $p_{\boldsymbol{s}_p}(V(1)\boldsymbol{P}_1) = p_{\boldsymbol{s}_0}(C_0)$ is the worst-case output probability we aim to estimate.

Given the above $V(\theta)$, we multiply the sampled permutation $\boldsymbol{P}_1$ at the right side of $V(\theta)$. Because $V(0)\boldsymbol{P}_1$ follows the distribution $\mathcal{H}_{\mathcal{A}}\mathcal{P}$, one can expect that the distribution of $V(\theta)\boldsymbol{P}_1$ is close to $\mathcal{H}_{\mathcal{A}}\mathcal{P}$ for sufficiently small θ. Then, one can also expect that the success probability of the oracle $\mathcal{O}$ (i.e., the right-hand side of Eq. (31)) on input a perturbed circuit $V(\theta)\boldsymbol{P}_1$ instead of $V(0)\boldsymbol{P}_1$ would not much be deviated from δ as long as θ is small enough.

Based on this understanding, we input $\boldsymbol{s}_p$ and $V(\theta)\boldsymbol{P}_1$ in the oracle $\mathcal{O}$. Then, by Eq. (31), for at least $1-\eta$ over $\boldsymbol{s}_p$, the failure probability of $\mathcal{O}$ is at most

$$\Pr\left[\left|\mathcal{O}(\boldsymbol{s}_p, V(\theta)\boldsymbol{P}_1) - p_{\boldsymbol{s}_p}(V(\theta)\boldsymbol{P}_1)\right| > \epsilon\right] < \delta + D_{\mathrm{TV}}(\mathcal{H}_{\mathcal{A},\theta}^{C_p}, \mathcal{H}_{\mathcal{A}}), \tag{33}$$

where D_{TV} denotes total variation distance. The inequality in Eq. (33) is evident based on the two facts: We can interpret the total variation distance as the supremum over events of the difference in probabilities of those events (Viz., circuits corresponding to the failure) [41], and the equal permutation on the circuit does not affect the total variation distance over the circuit space. In the right-hand side of Eq. (33), by Lemma 9, $D_{\mathrm{TV}}(\mathcal{H}_{\mathcal{A},\theta}^{C_p}, \mathcal{H}_{\mathcal{A}})$ is $O(m\theta)$. By setting $0 \le \theta \le \Delta$ with $\Delta = O(m^{-1})$, we can upper bound $D_{\mathrm{TV}}(\mathcal{H}_{\mathcal{A},\theta}^{C_p}, \mathcal{H}_{\mathcal{A}})$ by an arbitrarily small constant.

By Lemma 10, $p_{\boldsymbol{s}_p}(V(\theta)\boldsymbol{P}_1)$ is a $(4mN, 4mN)$ degree rational function $\frac{P(\theta)}{Q(\theta)}$, where the denominator is given as $Q(\theta) = \left[\prod_{i=1}^{m} |q_i(\theta)|^2\right]^N$ for $q_i(\theta)$ defined in Definition 8 but with subindex i corresponding to ith random gate $H_i(\theta)$. We note that $Q(\theta)$ can be computed in $\Theta(m)$ time, as it only depends on the constant number of eigenvalues of local gate matrices (i.e., ϕ_j values in Eq. (27) for each local gate matrix H_i). Also, given that $\theta \le \Delta = O(m^{-1})$, $Q(\theta)$ is very close to the unity, because $Q(\theta) \ge 1$ and

$$\begin{aligned} Q(\theta) &= \left[\prod_{i=1}^{m} |q_i(\theta)|^2\right]^N \\ &= \left[\prod_{i=1}^{m}\prod_{j=1}^{2} |(1 + i\theta e^{i\frac{\phi_{i,j}}{2}} \sin\frac{\phi_{i,j}}{2})|^2\right]^N \\ &\le \left[\prod_{i=1}^{m}\prod_{j=1}^{2} (1+\theta^2)\right]^N \\ &\le (1 + O(m^{-2}))^{2mN} \\ &= 1 + O(Nm^{-1}), \end{aligned} \tag{34}$$

where $\phi_{i,j}$ denotes the phase of the jth eigenvalue of the ith gate.

Therefore, $p_{\boldsymbol{s}_p}(V(\theta)\boldsymbol{P}_1)$ is very close to the degree $d = 4mN$ polynomial $P(\theta)$ in $\theta \in [0, \Delta]$, which allows us to use polynomial interpolation technique for $P(\theta)$. Specifically, we obtain estimations of $P(\theta)$ for different values of $\theta \in [0, \Delta]$ by querying the oracle $\mathcal{O}$, use polynomial interpolation for given $P(\theta)$ values to estimate $P(1)$, and infer the value $p_{\boldsymbol{s}_p}(V(1)\boldsymbol{P}_1) = p_{\boldsymbol{s}_0}(C_0)$ by multiplying $Q(1)^{-1}$. However, $Q(1)^{-1}$ becomes arbitrarily large for the case that even a single $\phi_{i,j}$ in Eq. (34) is near $\pm\pi$, which will arbitrarily enlarge the imprecision of the approximation of $p_{\boldsymbol{s}_p}(V(1)\boldsymbol{P}_1)$. To avoid this issue, we employ the strategy from Ref. [28], which only considers the case that all $\phi_{i,j}$ values of randomly chosen gates $\{H_i\}_{i=1}^m$ are in $[-\pi+\zeta, \pi-\zeta]$, and regards the other case as failure. This happens with probability at least $1 - O(m\zeta)$ over the random circuit instances. By setting $\zeta = O(m^{-1})$, we can make $O(m\zeta)$ arbitrarily small constant, and as a result, we can upper bound $Q(1)^{-1}$ (see Eq. (38) below) with high probability over random circuit instances.

Now the problem reduces to approximating degree $d = 4mN$ polynomial $P(\theta)$ with the value $\mathcal{O}(\boldsymbol{s}_p, V(\theta)\boldsymbol{P}_1)Q(\theta)$ in $\theta \in [0, \Delta]$ within additive error smaller than $\epsilon(1+O(Nm^{-1})) \simeq \epsilon$; such approximations will later be used for the estimation of the value $P(1)$ via polynomial interpolation technique. The failure of $\mathcal{O}$ depends on the outcome $\boldsymbol{s}_p \sim \mathcal{G}_{M,N}$ whose failure probability is at most η, and the circuit $V(\theta) \sim \mathcal{H}^{C_p}_{\mathcal{A},\theta}$ whose failure probability is at most $\delta + O(m\Delta)$ from Eq. (33). Also, the probability that at least one $\phi_{i,j}$ of randomly chosen gates $\{H_i\}_{i=1}^m$ is outside of the regime $[-\pi+\zeta, \pi-\zeta]$ is at most $O(m\zeta)$. Putting everything together and applying a simple union bound, the total failure probability of the approximation of $P(\theta)$ is at most

$$\Pr\left[|\mathcal{O}(\boldsymbol{s}_p, V(\theta)\boldsymbol{P}_1)Q(\theta) - P(\theta)| > \epsilon\right] < \eta + \delta + O(m\Delta) + O(m\zeta) \leq \delta', \tag{35}$$

where δ' is an upper bound of $\eta + \delta + O(m\Delta) + O(m\zeta)$, and given $\eta + \delta < \frac{1}{4}$, we can make $\delta' < \frac{1}{4}$ by setting $O(m\Delta)$ and $O(m\zeta)$ arbitrary small constants.

Let $\{\theta_i\}_{i=1}^{O(d^2)}$ be the set of equally spaced points in the interval $[0, \Delta]$. For each θ_i, we obtain the unitary matrix $U(\theta_i)$ using the same random gate $\{H_i\}_{i=1}^m$ and worst-case circuit C_p. Let $y_i = \mathcal{O}(\boldsymbol{s}_p, V(\theta_i)\boldsymbol{P}_1)Q(\theta_i)$. By Eq. (35), each set of points (θ_i, y_i) satisfies

$$\Pr[|y_i - P(\theta_i)| > \epsilon] \leq \delta' < \frac{1}{4}. \tag{36}$$

By using the interpolation algorithm introduced in Theorem 10, we can obtain the additive approximation of $P(1)$ as $\tilde{p}$ with an access to NP oracle, such that

$$\Pr\left[|\tilde{p} - P(1)| > \epsilon'\right] < \frac{1}{3}, \tag{37}$$

where $\epsilon' = \epsilon e^{-d\log\Delta} = \epsilon 2^{O(N^{\gamma+1}(\log N)^2)}$ using $d = 4mN$ and $m = qM\log M$. Note that the failure probability in Eq. (37) can be arbitrarily reduced by taking a polynomial number of trials, and thus we can obtain the estimated value $P(1)$ within additive error ϵ' with arbitrarily high probability.

From the estimated value $P(1)$, we can infer the worst-case output probability value $p_{\boldsymbol{s}_0}(C_0) = p_{\boldsymbol{s}_p}(V(1)\boldsymbol{P}_1) = P(1)/Q(1)$. As the value of $Q(1)$ depends on the values $\phi_{i,j}$ in Eq. (34), the $\phi_{i,j}$ independent lower bound of $Q(1)$ is required to set an upper bound of the additive imprecision of $p_{\boldsymbol{s}_p}(V(1)\boldsymbol{P}_1)$. Since we only consider the case that all of $\phi_{i,j}$ values are in $[-\pi+\zeta, \pi-\zeta]$ with $\zeta = O(m^{-1})$ for all randomly chosen gates $\{H_i\}_{i=1}^m$, we have

$$\begin{aligned}
Q(1) &= \left[\prod_{i=1}^m \prod_{j=1}^2 |(1 + ie^{i\frac{\phi_{i,j}}{2}} \sin\frac{\phi_{i,j}}{2})|^2\right]^N \\
&= \left[\prod_{i=1}^m \prod_{j=1}^2 \left(1 - \sin^2\frac{\phi_{i,j}}{2}\right)\right]^N \\
&\geq \left[\prod_{i=1}^m \prod_{j=1}^2 \left(1 - \sin^2\frac{\pi-\zeta}{2}\right)\right]^N \\
&= \left(O(m^{-2})\right)^{2mN} \\
&= 2^{2mN\log O(m^{-2})}.
\end{aligned} \tag{38}$$

Therefore, the total additive error for estimating $p_{\boldsymbol{s}_p}(V(1)\boldsymbol{P}_1)$ is bounded by $\epsilon' 2^{-2mN\log O(m^{-2})} = \epsilon 2^{O(N^{\gamma+1}(\log N)^2)}$. By setting $\epsilon = 2^{-O(N^{\gamma+1}(\log N)^2)}2^{-O(N)} = 2^{-O(N^{\gamma+1}(\log N)^2)}$, we can estimate the worst-case output probability value $p_{\boldsymbol{s}_p}(V(1)\boldsymbol{P}_1) = p_{\boldsymbol{s}_0}(C_0)$ within additive error $2^{-O(N)}$, and the whole reduction process is in $\mathrm{BPP}^{\mathrm{NP}}$. This completes the proof.

□

For the polynomial interpolation, we employ the Robust Berlekamp-Welch algorithm that requires the NP oracle, recently proposed in Ref. [28].

Theorem 10 (Robust Berlekamp-Welch [28])**.** *Let $P(x)$ be a degree d polynomial in x. Suppose there is a set of points $D = \{(x_i, y_i)\}$ such that $|D| = O(d^2)$ and $\{x_i\}$ is equally spaced in the interval $[0, \Delta]$. Suppose also that each points (x_i, y_i) satisfies*

$$\Pr[|y_i - P(x_i)| \geq \epsilon] \leq \delta, \tag{39}$$

with $\delta < \frac{1}{4}$. Then there exists a P^{NP} algorithm that takes input D and outputs $\tilde{p}$ such that

$$|\tilde{p} - P(1)| \leq \epsilon e^{-d\log\Delta}, \tag{40}$$

with success probability at least $\frac{2}{3}$.

G Proof of Lemma 3

Let $\bar{p}_{\boldsymbol{s}}(C)$ be the output probability distribution from the approximate sampler $\mathcal{S}$ with the given linear optical circuit C. Also, let $\mathcal{C}_{M,N}$ be the set of collision-free outcomes of Boson Sampling, for mode number M and photon number N. Then $\bar{p}_{\boldsymbol{s}}(C)$ satisfies

$$\mathop{\mathbb{E}}_{\boldsymbol{s}\sim\mathcal{G}_{M,N}}[|\bar{p}_{\boldsymbol{s}}(C) - p_{\boldsymbol{s}}(C)|] = \frac{1}{\binom{M}{N}} \sum_{\boldsymbol{s}\in\mathcal{C}_{M,N}} |\bar{p}_{\boldsymbol{s}}(C) - p_{\boldsymbol{s}}(C)| \leq \frac{2\beta}{\binom{M}{N}}. \tag{41}$$

Using Eq. (41) and Markov's inequality, $\bar{p}_{\boldsymbol{s}}(C)$ satisfies

$$\Pr_{\boldsymbol{s}\sim\mathcal{G}_{M,N}}\left[|\bar{p}_{\boldsymbol{s}}(C) - p_{\boldsymbol{s}}(C)| \geq \frac{\beta k}{\binom{M}{N}}\right] \leq \frac{2}{k}, \tag{42}$$

for all $k > 2$. Also, using Stockmeyer's algorithm [48] whose complexity is in $\mathrm{BPP}^{\mathrm{NP}}$, obtaining the estimate $\tilde{p}_{\boldsymbol{s}}(C)$ of $\bar{p}_{\boldsymbol{s}}(C)$ satisfying

$$\Pr\left[|\tilde{p}_{\boldsymbol{s}}(C) - \bar{p}_{\boldsymbol{s}}(C)| \geq \alpha\bar{p}_{\boldsymbol{s}}(C)\right] \leq \frac{1}{2^N}, \tag{43}$$

in polynomial time in N and α^{-1} is in $\mathrm{BPP}^{\mathrm{NP}^{\mathcal{S}}}$. Using $\mathbb{E}_{\boldsymbol{s}\sim\mathcal{G}_{M,N}}[\bar{p}_{\boldsymbol{s}}(C)] = \binom{M}{N}^{-1}\sum_{\boldsymbol{s}\in\mathcal{C}_{M,N}}\bar{p}_{\boldsymbol{s}}(C) \leq \binom{M}{N}^{-1}$,

$$\Pr\left[|\tilde{p}_{\boldsymbol{s}}(C) - \bar{p}_{\boldsymbol{s}}(C)| \geq \frac{\alpha l}{\binom{M}{N}}\right] \leq \Pr\left[\bar{p}_{\boldsymbol{s}}(C) \geq \frac{l}{\binom{M}{N}}\right] + \Pr\left[|\tilde{p}_{\boldsymbol{s}}(C) - \bar{p}_{\boldsymbol{s}}(C)| \geq \alpha\bar{p}_{\boldsymbol{s}}(C)\right] \tag{44}$$

$$\leq \frac{1}{l} + \frac{1}{2^N}, \tag{45}$$

for all $l > 1$. Putting all together, by applying a triangular inequality, finding an average-case approximation $\tilde{p}_{\boldsymbol{s}}(C)$ of $p_{\boldsymbol{s}}(C)$ satisfying

$$\Pr\left[|\tilde{p}_{\boldsymbol{s}}(C) - p_{\boldsymbol{s}}(C)| \geq \frac{\beta k + \alpha l}{\binom{M}{N}}\right] \leq \frac{2}{k} + \frac{1}{l} + \frac{1}{2^N} \tag{46}$$

is in $\mathrm{BPP}^{\mathrm{NP}^{\mathcal{S}}}$. Let κ and ξ be fixed error parameters such that $k/2 = l = 3/\xi$ and $\beta = \kappa\xi/12 = \alpha/2$. As $\beta k + \alpha l = \kappa$ and $\frac{2}{k} + \frac{1}{l} + \frac{1}{2^N} = \frac{2}{3}\xi + \frac{1}{2^N} \leq \xi$, we finally obtain Eq. (5).

H On average-case hardness including collision outcomes

As we have remarked in the main text, extending our hardness result of shallow-depth Boson Sampling to the saturated regime $M = \Theta(N^\gamma)$ for $1 \leq \gamma < 2$ remains a crucial open problem, and to do so, one also needs to consider collision outcomes in the average-case hardness argument. In this appendix, we present a possible approach that can be further improved to the average-case hardness result including collision outcomes.

Our intuition is that while most of the outcomes in the saturated regime are not collision-free, the key point is that most outcomes are expected to be "nearly" collision-free, in the sense that for most of the outcomes, a large portion of the measured modes are occupied only by a single photon. Also, we can establish worst-to-average-case reduction across different system sizes (i.e., mode number & photon number), i.e., the system size for the worst-case hardness result can be smaller at most polynomially than the system size for the average-case hardness result. Hence, by post-selecting outcomes with the number of measured single photons at least as large as the photon number required in the worst-case problem (in Theorem 3, which requires a collision-free outcome), one would possibly establish worst-to-average-case reduction from the randomly chosen outcome over all possible outcomes.

More specifically, let the mode number $M = \Theta(N^\gamma)$ for any $\gamma \geq 1$. Let us define $\mathcal{Q}_{M,N}$ as the uniform distribution over all the possible outcomes of Boson Sampling with M modes and N photons, such that each outcome $\boldsymbol{s} \sim \mathcal{Q}_{M,N}$ is an M-dimensional output configuration vector $\boldsymbol{s} = (s_1, \ldots, s_M)$ with the constraint $\sum_{i=1}^{M} s_i = N$. Also, we denote the number of $s_i = 1$ in $\boldsymbol{s}$ as the number of *singletons* throughout this appendix. Here, let us assume that there exists $N_0 = cN^\lambda$ for a constant c and $0 < \lambda < 1$ such that it is likely that the number of singletons is more than N_0; we later provide the intuition to justify this assumption at the end of this appendix.

Hereafter, we denote each (M_0, N_0) and (M, N) as the system sizes (mode number & photon number) for the worst-case hardness (i.e., Theorem 3) and average-case hardness (we aim to prove), respectively. More specifically, given a worst-case circuit C_0 in $M_0 = \Omega(N_0)$ mode circuit architecture $\mathcal{BB}^*$ and a fixed M_0-mode and N_0-photon collision-free outcome $\boldsymbol{s}_0$ depicted in Theorem 3, by our assumption, one can find a photon number N that satisfies $N_0 = cN^\lambda$ for a constant c and $0 < \lambda < 1$, and mode number $M = \Theta(N^\gamma)$ for $\gamma \geq 1$, such that most of the outcomes over $\mathcal{Q}_{M,N}$ have at least N_0 singletons. In the following, we show a possible approach for worst-to-average-case reduction for general outcomes, by describing how one can estimate $p_{\boldsymbol{s}_0}(C_0)$ to within desired imprecision given oracle access to the well-estimated values $p_{\boldsymbol{s}}(U)$ for average-case outcomes $\boldsymbol{s} \sim \mathcal{Q}_{M,N}$ and M-mode average-case circuits over $U \sim \mathcal{H}_{\mathcal{A}}$ with $\mathcal{A} = (\mathcal{BB}^*)^q$.

At first, for the average-case instances, we sample $\boldsymbol{s} \sim \mathcal{Q}_{M,N}$ that has at least N_0 singletons, and sample $U \sim \mathcal{H}_{\mathcal{A}}$ specifically for M-mode circuit architecture $\mathcal{A} = (\mathcal{BB}^*)^{q \geq 2}$. Subsequently, we sample $M/2$-dimensional configuration vector $\boldsymbol{s}_1$ through the following steps: (1) discard N_0 singletons from $\boldsymbol{s}$; (2) retain only the non-zero elements of $\boldsymbol{s}$ and discard all the zeros; (3) append zeros as needed to construct $M/2$-dimensional vector $\boldsymbol{s}_1$, and (4) randomly permute the elements of $\boldsymbol{s}_1$ uniformly. Then, given $\boldsymbol{s}$, $\boldsymbol{s}_0$ and $\boldsymbol{s}_1$, one

can efficiently construct M by M permutation matrix $\boldsymbol{P}$ that satisfies

$$\boldsymbol{s} = \boldsymbol{P}(\boldsymbol{s}_0 \oplus 0_{M/2-M_0} \oplus \boldsymbol{s}_1), \tag{47}$$

where we denote $0_{M/2-M_0}$ as a $(M/2-M_0)$-dimensional vector composed on only zeros. In addition, let $\boldsymbol{t} = \boldsymbol{t}_0 \oplus 0_{M/2-M_0} \oplus \boldsymbol{t}_1$ be an input configuration for the average-case output probability, where $\boldsymbol{t}_1$ is some fixed $M/2$-mode and $(N-N_0)$-photon configuration vector.

Next, given the sampled U, we denote the first-rightmost $M/2$-mode subcircuit of U as U_0, such that U_0 is in $M/2$-mode circuit architecture $\mathcal{A} = \mathcal{BB}^*$ and essentially follows the distribution $\mathcal{H}_{\mathcal{A}}$; we note that, $M/2$-mode $\mathcal{BB}^*$ can be embedded into M-mode $\mathcal{BB}^*$ (see, e.g., Fig. 3). Using the definitions so far, we construct the revised worst-case circuit $C \in (\mathcal{BB}^*)^{q=2}$ as

$$C = \boldsymbol{P}(C_0 \oplus I_{M/2-M_0} \oplus U_0), \tag{48}$$

where we denote $I_{M/2-M_0}$ as a $(M/2-M_0)$-sized identity matrix. We note that the revised worst-case circuit can be embedded in the higher depth circuit architecture $C \in (\mathcal{BB}^*)^{q\geq 2}$, by adding trivial gates in the front. Then, by definition, for input configuration $\boldsymbol{t}$, the output probability $p_{\boldsymbol{s}}(C)$ can be expressed as

$$p_{\boldsymbol{s}}(C) = p_{\boldsymbol{s}_0}(C_0)p_{\boldsymbol{s}_1}(U_0). \tag{49}$$

Therefore, if one can well-estimate the revised worst-case value $p_{\boldsymbol{s}}(C)$ and the average-case value $p_{\boldsymbol{s}_1}(U_0)$, one can also infer the desired worst-case output probability $p_{\boldsymbol{s}_0}(C_0)$, thereby establishing worst-to-average-case reduction.

We now argue how to estimate the worst-case value $p_{\boldsymbol{s}}(C)$ and the average-case value $p_{\boldsymbol{s}_1}(U_0)$. First, estimating the average-case value $p_{\boldsymbol{s}_1}(U_0)$ can be done by similarly accessing the oracle, but now we input the outcome $\boldsymbol{s}_1$ and the circuit U_0 instead. Hence, the success probability for the oracle would increase at most the total variation distance between the distribution of $\boldsymbol{s}_1$ and $\mathcal{Q}_{M/2,N-N_0}$, and if this can be sufficiently bounded, we can obtain the estimated value of $p_{\boldsymbol{s}_1}(U_0)$ with high probability. Also, one can estimate the value $p_{\boldsymbol{s}}(C)$ by the polynomial interpolation, as in the proof of Theorem 4. Specifically, by parameterizing the circuit U as $U(\theta)$, where each gate of U (except those in U_0) is perturbed via the Cayley transform using the corresponding gate from C, one can similarly infer the value $p_{\boldsymbol{s}}(C)$ through oracle access to $p_{\boldsymbol{s}}(U(\theta))$ and applying polynomial interpolation (see Appendix F for details).

To sum up, the overall worst-to-average-case reduction process can be schematized as follows.

$$p_{\boldsymbol{s}_0}(C_0) = \frac{p_{\boldsymbol{s}}(C)}{p_{\boldsymbol{s}_1}(U_0)} \quad \xrightarrow[\substack{U\to C \\ \text{Cayley transform} \\ \text{and} \\ \text{polynomial interpolation}}]{} \quad \frac{p_{\boldsymbol{s}}(U)}{p_{\boldsymbol{s}_1}(U_0)} \approx \frac{\mathcal{O}(\boldsymbol{s}, U)}{\mathcal{O}(\boldsymbol{s}_1, U_0)}, \tag{50}$$

where $\mathcal{O}$ represents the oracle for the average-case estimation problem.

We highlight the necessary ingredients that need to be addressed for the successful worst-to-average-case reduction by the above process. First, to make sure that we can also well-estimate $p_{\boldsymbol{s}_1}(U_0)$ with high probability over $\boldsymbol{s}_1$ via the oracle access, one needs to justify that the distribution of $\boldsymbol{s}_1$ is sufficiently close to the distribution $\mathcal{Q}_{M/2,N-N_0}$. To help understand how the distribution of $\boldsymbol{s}_1$ can be close to $\mathcal{Q}_{M/2,N-N_0}$, consider the limiting case that $N_0 = 0$. Here, given that M is much larger than N (but still can be $M = \Theta(N)$), sampling $\boldsymbol{s}_1$ is equivalent to sampling from $\mathcal{Q}_{M,N}$ and post-selecting outcomes that have

all the non-zero elements in the first $M/2$ mode, which yields the uniform distribution $\mathcal{Q}_{M/2,N}$. Therefore, given that M is much larger than N and N_0 is sufficiently small relative to N, we expect that the distribution of $\boldsymbol{s}_1$ would not largely deviate from the distribution $\mathcal{Q}_{M/2,N-N_0}$.

Also, as the final imprecision level increases by the estimated value of $p_{\boldsymbol{s}_1}(U_0)^{-1}$, it must be promised that $p_{\boldsymbol{s}_1}(U_0)$ is lower bounded. More precisely, $p_{\boldsymbol{s}_1}(U_0)$ must be at least as large as the allowed imprecision level of average-case estimation (i.e., imprecision level of the oracle $\mathcal{O}$), commonly referred to as the anti-concentration property [1, 41, 57–60]. Nevertheless, we believe that $p_{\boldsymbol{s}_1}(U_0)$ is at least as large as the allowed imprecision level for most of U_0, based on two reasons: First, the imprecision level required by our current average-case hardness result is sufficiently small. Also, although this has only been demonstrated in finite-sized numerical experiments, random circuits for our circuit architecture closely resemble global Haar-random unitary circuits [47], which are strongly conjectured to exhibit the anti-concentration property [1].

We now present a toy model that could offer an insight that, for a large portion of outcomes of Boson Sampling, a sufficient number of singletons would exist. In the following, we present the classical "balls-into-bins" problem where we throw N balls into M bins uniformly at random, which is indeed the classical analogue of Boson Sampling. We show that, given that $M = \Theta(N^\gamma)$ with $\gamma \geq 1$, it is highly unlikely that the number of bins getting exactly one ball is less than cN^λ for a constant c and $0 < \lambda < 1$.

Let us first denote the random variable X_i as the number of balls in i-th bin (such that $\sum_{i=1}^M X_i = N$), when we throw N balls into M bins uniformly at random. Here, our goal is to verify that over the experiments, it is highly unlikely to observe the outcomes that the number of $X_i = 1$ is smaller than cN^λ. Since all the variables X_i are correlated due to the constraint $\sum_{i=1}^M X_i = N$, instead, we consider i.i.d. random variable Y_i each drawn from Poisson distribution $P(\frac{N}{M})$, such that $\mathbb{E}[\sum_{i=1}^M Y_i] = N$ and $\Pr\left[\sum_{i=1}^M Y_i = N\right] = e^{-N}\frac{N^N}{N!} = \Theta(\frac{1}{\sqrt{N}})$. The motivation for using Poissonian random variables is that the distribution of $(Y_1, \ldots, Y_M)$ conditioned on the sum $\sum_{i=1}^M Y_i = N$ is equivalent to the distribution of the balls-into-bins problem $(X_1, \ldots, X_M)$ (see, e.g., Ref. [61]). Accordingly, if it is exponentially unlikely that the number of $Y_i = 1$ is smaller than cN^λ as we sample $(Y_1, \ldots, Y_M)$, this indicates that it is still exponentially unlikely that the number of $X_i = 1$ is smaller than cN^λ as we sample $(X_1, \ldots, X_M)$, given that the post-selection probability for $\sum_{i=1}^M Y_i = N$ is at most inverse polynomial (i.e., $\Theta(\frac{1}{\sqrt{N}})$).

Now, let us newly define $Z = \sum_{i=1}^M Z_i$ as the summation of the independent Bernoulli random variable $Z_i \in \{0, 1\}$, where each $Z_i = 1$ only if $Y_i = 1$, and $Z_i = 0$ otherwise. Then, Z indicates the number of $Y_i = 1$, and thus our problem reduces to verifying that the probability $\Pr\left[Z \leq cN^\lambda\right]$ is exponentially small. To proceed, as $Y_i \sim P(\frac{N}{M})$, note that the probability of obtaining $Z_i = 1$ is given by

$$\Pr[Z_i = 1] = \frac{N}{M}\exp\left(-\frac{N}{M}\right). \tag{51}$$

Then, the expectation value of Z is given by

$$\mathbb{E}[Z] = \sum_{i=1}^M \mathbb{E}[Z_i] = \sum_{i=1}^M \Pr[Z_i = 1] = N\exp\left(-\frac{N}{M}\right). \tag{52}$$

As all the Z_i are independent Bernoulli random variables, we use the Chernoff bound,

which gives the probability bound as

$$\Pr[Z \leq \varepsilon\, \mathbb{E}[Z])] \leq \exp\left(-\frac{(1-\varepsilon)^2}{2}\, \mathbb{E}[Z]\right). \tag{53}$$

Hence, by setting $\varepsilon\, \mathbb{E}[Z] = cN^{\lambda}$, we finally obtain the bound

$$\Pr\left[Z \leq cN^{\lambda}\right] \leq \exp\left(-\frac{1}{2}\left(1 - cN^{\lambda-1}\exp\left(\frac{N}{M}\right)\right)^2 N \exp\left(-\frac{N}{M}\right)\right) = \exp\left(-\Theta(N)\right). \tag{54}$$

Therefore, for the classical balls-into-bins problem, it is exponentially unlikely that the number of bins getting exactly one ball is less than cN^{λ}. Note that this bound even allows $\lambda = 1$ as long as $c\exp\left(\frac{N}{M}\right) < 1$, which implies that a constant portion of bins will get exactly one ball.

We emphasize that the distribution of the classical balls-into-bins problem is not equivalent to the uniform distribution over the possible outcomes of Boson Sampling, because the classical balls are distinguishable while identical bosons are indistinguishable in nature (and indeed, bosons are even more "gregarious" than classical balls [1]). Nevertheless, this example is sufficient to provide an insight that a sufficiently large number of singletons would appear in a large fraction of Boson Sampling outcomes.

We conclude this appendix by introducing an alternative approach, based on the closely related result in [29] that establishes a hardness argument for general Boson Sampling in the saturated regime. Note that our convention on singletons closely parallels the notion of "clicks" introduced in [29], which refers to the number of modes that are occupied by a *non-zero* number of photons. Analogous to our analysis, their key observation is that, even in the saturated regime, most outcomes have a sufficiently large number of clicks. Based on this observation, they show that, given an N-photon outcome $\boldsymbol{s}$ with collisions that has a sufficient number of clicks N_0, there exists a poly-time constructible matrix $A \in \{0,1\}^{N_0 \times N}$ such that computing $\mathrm{Per}(A_{\boldsymbol{s}})$ is worst-case #P-hard, for $A_{\boldsymbol{s}}$ being a submatrix of A by repeating rows according to $\boldsymbol{s}$. Then they establish a worst-to-average-case reduction by perturbing A into their average-case matrix and performing polynomial interpolation (for details, see [29]).

Accordingly, since we employ a conceptually similar strategy in the main text, establishing worst-case hardness for shallow-depth circuits with a large enough number of clicks is sufficient to obtain an analogous hardness result for shallow-depth Boson Sampling. This requires constructing a shallow-depth circuit (as done in the proof of Theorem 3) for which estimating an output probability with a large enough number of clicks is worst-case #P-hard, in direct analogy to the construction of the worst-case matrix A above. Since this lies beyond the scope of the presented work, we leave this as future work.

I Numerical simulation for the collision patterns of Boson Sampling for the local random circuit over $(\mathcal{BB}^*)^q$

In this appendix, we present numerical simulation results of collision patterns of Boson Sampling, for solely the local random circuit ensemble in Definition 4 over our shallow-depth circuit architecture $(\mathcal{BB}^*)^q$ in Definition 3.

To show the bosonic birthday paradox property for the local random circuit ensemble over $(\mathcal{BB}^*)^q$, we show numerically that the local random circuit ensemble over $(\mathcal{BB}^*)^q$

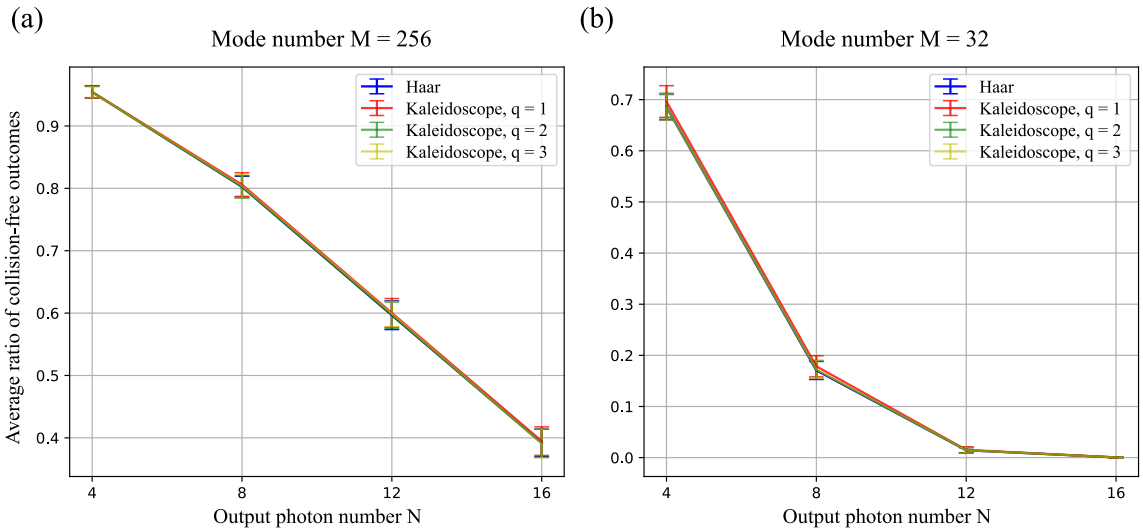

Figure 6: The average ratio of collision-free outcomes from (a) 256-mode and (b) 32-mode unitary circuits drawn from local random circuit ensemble $\mathcal{H}_{\mathcal{A}}$ over the q-Kaleidoscope circuit architecture $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^q$ with the repetition number $q = 1, 2, 3$, and global Haar random circuit ensemble. We sample 500 unitary circuits from each of the circuit ensembles, and we sample 500 outcomes on input $N = 4, 8, 12, 16$ photons for each of the randomly chosen circuits. Then we count the collision-free outcomes over the 500 sampled outcomes and compute the average ratio of collision-free outcomes over 500 randomly chosen circuits for each of the circuit ensembles. We present in the figure the average ratio of collision-free outcomes for local random circuits over $(\mathcal{B}\mathcal{B}^*)^q$ with $q = 1, 2, 3$ and global Haar random circuit ensemble, for the photon number given by $N = 4, 8, 12, 16$ (the error bar represents standard deviation). For all of the repetition numbers $q = 1, 2, 3$ and the photon numbers $N = 4, 8, 12, 16$, local random circuits for $(\mathcal{B}\mathcal{B}^*)^q$ exhibits almost equivalent average collision-free ratio to the global Haar random unitary circuits.

samples sufficiently many collision-free outcomes over the experiments. Specifically, we prepare random unitary circuits drawn from the local random circuit ensemble over $(\mathcal{B}\mathcal{B}^*)^q$, sample the outcomes by classical Boson Sampling algorithm for each randomly chosen circuit, and compute the ratio of the collision-free outcomes over all the sampled outcomes for each randomly chosen circuit. To evaluate whether these local random circuits output sufficiently many collision-free outcomes, as a comparison, we also prepare the global Haar random circuits, which have the bosonic birthday paradox property as depicted in [1]. In other words, if the ratio of collision-free outcomes for local random circuit ensemble over $(\mathcal{B}\mathcal{B}^*)^q$ is as large as the ratio of collision-free outcomes for global Haar random circuit, this clearly gives evidence of the bosonic birthday paradox for our circuit architecture.

More specifically, we consider 256-mode random unitary circuits drawn from local random circuit ensemble $\mathcal{H}_{\mathcal{A}}$ over the q-Kaleidoscope circuit architecture $\mathcal{A} = (\mathcal{B}\mathcal{B}^*)^q$ with the repetition number $q = 1, 2, 3$, and global Haar random circuit ensemble. First, we sample 500 unitary circuits from each of the circuit ensembles described above. For each of the randomly chosen circuits, on input $N = 4, 8, 12, 16$ photons in the first N modes over 256 modes, we sample 500 outcomes using the Clifford-Clifford algorithm for Boson Sampling [62]. Then, we count the number of collision-free outcomes over the 500 sampled outcomes and compute the average ratio of collision-free outcomes for each N and circuit ensemble, where the average is taken over the 500 sampled unitary circuits for each circuit ensemble.

In Fig. 6, we present the average ratio of collision-free outcomes and standard deviation (error bar) for the photon number given by $N = 4, 8, 12, 16$, from the local random cir-

cuits over $(\mathcal{BB}^*)^q$ with the repetition number $q = 1, 2, 3$ and global Haar random circuits. Remarkably, local random circuits for $(\mathcal{BB}^*)^q$ shows great resemblance to the global Haar random circuits: For all of the photon numbers $N = 4, 8, 12, 16$ and the repetition numbers $q = 1, 2, 3$, local random circuits for $(\mathcal{BB}^*)^q$ exhibits almost identical average collision-free ratio (and also the standard deviation) to the global Haar random circuits. Given that the bosonic birthday paradox is guaranteed for the global Haar random circuit ensemble, also given that the mode number we used covers all recent large-sized Boson Sampling experiments [4–7], our numerical result shows that the local random circuit ensemble for $(\mathcal{BB}^*)^q$ guarantees the bosonic birthday paradox at least for the near-term finite-sized experiments. Even though this does not guarantee the bosonic birthday paradox of the local random circuit ensemble for $(\mathcal{BB}^*)^q$ in the asymptotic limit as the system size scales, still, we believe that constant repetition number q would be sufficient for the bosonic birthday paradox, due to the close resemblance to the global Haar random circuit ensemble as shown in the figure.

Additionally, we similarly present a numerical simulation result in Fig. 6 for the mode number $M = 32$ and photon number $N = 4, 8, 12, 16$, which indicates the saturated regime where the bosonic birthday paradox does not hold. Even in this regime, the collision patterns of local random circuits over $(\mathcal{BB}^*)^q$ also show a great resemblance to those of global Haar random circuits. This result gives hope that one can possibly extend our result to the saturated regime, by conducting a similar analysis for our circuit ensemble to the analysis for global Haar random circuits [29].