

Transition from statistical to hardware-limited scaling in photonic quantum state reconstruction

Attila Baumann,¹ Zsolt Kis²,[✉] János Koltai[✉],¹ and Gábor Vattay[✉],^{1,*}

¹*Institute for Physics and Astronomy, Eötvös Loránd University, Pázmány Péter s 1/A, 1117 Budapest, Hungary*

²*Institute for Solid State Physics and Optics, HUN-REN Wigner Research Centre for Physics, Konkoly-Thege Miklós út 29-33, 1117 Budapest, Hungary*



(Received 28 November 2025; accepted 29 June 2026; published 10 August 2026)

The theoretical efficiency of classical shadow tomography is predicated on a perfect Haar-random unitary ensemble, yet this mathematical ideal remains physically unattainable in near-term hardware. Here, we report the experimental characterization of a hardware-induced systematic error floor on integrated photonic processors: an operational Hardware Horizon, where the reconstruction error undergoes a sharp crossover. While the error initially obeys the predicted statistical scaling $\mathcal{O}(M^{-1/2})$ in a variance-dominated regime, it abruptly saturates at a floor determined by the spectral distortions of the realized unitary group, marking a transition to a bias-dominated regime. By deriving a phenomenological error model, we decouple the competing mechanisms of static coherent spectral distortion and dynamic decoherence, demonstrating that this device-specific noise floor imposes a systematic saturation limit that statistical accumulation cannot overcome. These findings establish that the utility of shadow tomography on NISQ (noisy intermediate-scale quantum) hardware is defined by a specific scaling law involving hardware parameters, necessitating active compensation strategies to bridge the gap between theoretical purity and the noisy reality of integrated photonics.

DOI: [10.1103/3gsc-krrq](https://doi.org/10.1103/3gsc-krrq)

I. INTRODUCTION

Efficiently characterizing quantum states is a fundamental prerequisite for scaling quantum information technologies [1]. Standard quantum tomography, however, faces an insurmountable barrier for large systems, as its resource cost scales exponentially with the number of components [2,3]. This scaling problem has necessitated the development of resource-efficient protocols, including compressed sensing, matrix-product-state tomography, and permutationally invariant tomography [4,5]. Among these, classical shadow tomography has emerged as a versatile alternative [6]. By averaging measurements performed after applying a series of random unitary transformations theoretically drawn from a continuous Haar distribution, this technique allows for the reconstruction of the full density matrix with a measurement count that scales only polynomially with system size [1].

Photonic processors represent a leading platform for implementing such protocols [7]. Devices based on silicon-nitride waveguides offer scalability, high-speed operation, and excellent coherence due to photon isolation at room temperature [8,9]. With appropriate design, these processors can implement complex unitary transformations by propagating photons through a tunable mesh of phase shifters and beam

splitters, effectively realizing the random rotations required for shadow tomography.

However, a critical disparity exists between the mathematical axioms of shadow tomography and the physical reality of noisy intermediate-scale quantum (NISQ) implementations. The protocol's polynomial scaling advantage is predicated on sampling from a perfect Haar-random ensemble—an assumption that is physically unattainable in near-term hardware. Real-world photonic processors do not implement continuous group operations; instead, they realize a spectrally distorted approximation of the unitary group, governed by the inevitable interplay of static coherent spectral distortion and dynamic thermal noise. This creates a structural incompatibility: while statistical errors vanish as $\mathcal{O}(M^{-1/2})$, these hardware-induced spectral distortions impose a systematic, nonvanishing floor on reconstruction fidelity. Consequently, the central challenge shifts from optimizing sample complexity to characterizing this operational Hardware Horizon, which defines the specific resolution limit of the utilized device.

In this paper, we address this gap by reporting the experimental characterization of a hardware-dependent systematic error floor—an operational Hardware Horizon—for shadow tomography on integrated photonics. We implement a classical analog of density matrix reconstruction on a programmable eight-channel QuiX Quantum photonic processor [10]. By contrasting extensive numerical simulations with experimental data, we identify two distinct scaling regimes. In the statistical regime, reconstruction error decreases according to the theoretical prediction $\mathcal{O}(M^{-1/2})$, characteristic of a variance-dominated regime. However, we observe a crossover into a device regime, a bias-dominated regime where the

*Contact author: gabor.vattay@ttk.elte.hu

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

fidelity saturates despite increased measurement statistics (M). This saturation reveals that the reconstruction is no longer limited by sample size but by the intrinsic spectral distortion of the realized unitary group. Through a phenomenological error model, we decouple the mechanisms driving this saturation, demonstrating that it arises from the competition between static coherent spectral distortion and dynamic decoherence. We establish that this device-specific error floor imposes a systematic saturation limit on accuracy that statistical accumulation cannot overcome. These findings suggest that the utility of shadow tomography on NISQ hardware is defined by a specific scaling law involving hardware parameters, necessitating the development of active compensation strategies to bridge the gap between the theoretical purity of the Haar measure and the noisy reality of integrated photonics.

The remainder of this paper is structured as follows: In Sec. II, we describe our methodology, including the experimental setup on the photonic chip and the core principles of the reconstruction algorithm. Section III presents the numerical simulations that validate the analytical model and establish the theoretical baseline. In Sec. IV, we detail the experimental results, deriving the phenomenological error model that defines the Hardware Horizon. Finally, in Sec. V, we discuss the implications of these hardware-imposed limits for future quantum state characterization. The detailed theoretical derivations for the reconstruction formula, Haar-measure generation, and the analytical Frobenius norm are provided in Appendixes A, B, and C.

II. METHODS

The central concept of this work involves reconstructing the quantum density matrix ρ solely from classical measurement data. We consider a d -dimensional quantum system where, prior to detection, a unitary transformation U is applied, effectively rotating the measurement basis relative to the state. The measurement outcome is formalized through the binary indicator variable b_i . Specifically, the value $b_i = 1$ corresponds to the event where the system collapses into the i th state of the computational basis, while $b_j = 0$ for all $j \neq i$. The probability $P(b_i = 1)$ represents the likelihood of this specific event occurring. This probability is determined by the Born rule, which dictates that the likelihood of observing outcome i is equal to the i th diagonal element of the density matrix in the rotated frame:

$$P(b_i = 1) = (U \rho U^\dagger)_{ii} = \sum_{j,k} U_{ij} \rho_{jk} U_{ik}^*. \quad (1)$$

For a fixed U , a set of measurements yields the probabilities $P(b_i = 1)$. These outcomes allow us to define a diagonal matrix ρ^{exp} , where $\rho_{ii}^{\text{exp}} = P(b_i = 1)$. Consequently, an initial estimator $\hat{\rho}^{(1)}$ of the original density matrix elements can be constructed as

$$\hat{\rho}_{ab}^{(1)} = \sum_i U_{ia}^* \rho_{ii}^{\text{exp}} U_{ib} = \sum_{i,j,k} U_{ia}^* U_{ij} \rho_{jk} U_{ik}^* U_{ib}. \quad (2)$$

To generalize this, we repeat the procedure with random matrices drawn from the Haar distribution of the compact group

of unitary matrices $U(d)$. Averaging over this ensemble yields

$$I_{ab} = \mathbb{E}_U[\hat{\rho}_{ab}^{(1)}] = \int_{U(d)} dU \sum_{i,j,k} U_{ia}^* U_{ij} \rho_{jk} U_{ik}^* U_{ib}, \quad (3)$$

where dU denotes the Haar measure. The evaluation of this integral requires computing fourth-order moments of the form $\int dU U_{a'b'} U_{c'd'} U_{e'f'}^* U_{g'h'}^*$, which may be performed using Weingarten calculus [11]. This calculation results in

$$I_{ab} = \frac{1}{d+1} (\rho_{ab} + \delta_{ab}). \quad (4)$$

Therefore, the elements of the original density matrix ρ_{ab} can be expressed directly in terms of the averaged quantity I_{ab} :

$$\rho_{ab} = (d+1)I_{ab} - \delta_{ab}. \quad (5)$$

In a practical experimental setting, I_{ab} is estimated by averaging $\rho_{ab}^{(k)} = \sum_i U_{ia}^* \rho_{ii}^{\text{exp}} U_{ib}$ over a finite number M of random unitary transformations $\{U^{(k)}\}_{k=1}^M$:

$$\overline{I}_{ab}^{(M)} = \frac{1}{M} \sum_{k=1}^M \rho_{ab}^{(k)} = \mathbb{E}_U^{(M)} \left[\sum_i U_{ia}^* \rho_{ii}^{\text{exp}} U_{ib} \right]. \quad (6)$$

The final reconstructed matrix is thus given by

$$\rho_{ab}^{(\text{rec})} = (d+1)\overline{I}_{ab}^{(M)} - \delta_{ab}. \quad (7)$$

While this formula provides the reconstruction mechanism, the critical limitation lies in the statistical convergence; specifically, how the reconstruction error scales with the dimension d when the sample size M remains finite. The uniqueness of our approach is that, unlike the standard quantum approach, where each measurement yields a single discrete outcome $\hat{b} \in \{1, \dots, d\}$ (a click), in our photonic implementation, we do not observe a single collapse. Instead, we measure the probability of all d outputs simultaneously. This provides the full probability distribution $p_k = \langle k | U \rho U^\dagger | k \rangle$ for a given unitary U in a single experimental run. The estimator for a single snapshot is rank 1:

$$\hat{\rho}_{\text{single}} = (d+1)U^\dagger |\hat{b}\rangle \langle \hat{b}| U - \mathbb{I}. \quad (8)$$

This process introduces significant shot noise because information about the other $d-1$ outcomes is lost in the collapse. The estimator becomes a weighted sum:

$$\hat{\rho}_{\text{int}} = (d+1) \sum_{k=1}^d p_k U^\dagger |k\rangle \langle k| U - \mathbb{I}. \quad (9)$$

By accessing the full distribution, the shot noise associated with outcome selection is eliminated, leaving only the variance from the Haar random sampling of U . As derived in Appendix C, the mean squared error (MSE) scales linearly with d :

$$\mathbb{E}[\|\hat{\rho}^{(M)} - \rho\|_F^2]_{\text{int}} = \frac{d \text{Tr}(\rho^2) - 1}{M}. \quad (10)$$

In the following section, we detail the realization and experimental validation of this scheme on a photonic integrated circuit.




FIG. 1. Schematic of the experimental arrangement. The LASER source acts as the carrier, while the fiber switch (FS) prepares the initial state, denoted by the density matrix ρ . This state propagates through the photonic processor (PP), which implements the specific unitary transformation U . Finally, the photodiode array (PDA) detects the output intensity distribution, generating the input for the reconstruction of the rotated density matrix.

A. Experimental realization

The experiments were conducted using an eight-channel Mach-Zehnder interferometer (MZI) mesh, following the architecture proposed by Clements *et al.* [12] and implemented via a fully programmable Quix Quantum processor [13]. The experimental setup, illustrated in Fig. 1, comprises a 1550 nm laser source (Thorlabs S5FC1005P) and a fiber switch that governs photon injection into specific channels of the eight-channel processor.

The chosen light source is a continuous-wave Fabry-Perot laser diode operating at a central wavelength of 1550 nm with a -3 dB bandwidth of 50 nm. The coherence length of this broadband field is defined by $L_c = \lambda^2/\Delta\lambda$, yielding approximately $48 \mu\text{m}$ in vacuum. Accounting for the refractive index of the silicon nitride host material ($n \approx 2$ at 1550 nm), the effective coherence length within the waveguides is reduced to approximately $24 \mu\text{m}$. While this corresponds to a relatively short coherence window of roughly 31 wavelengths, it does not degrade the implementation of the unitary transformations. The fundamental building block of the integrated photonic network is the balanced Mach-Zehnder interferometer, composed of tunable beam splitters and phase shifters; see Fig. 2.

This specific architecture is engineered such that the structural optical path length remains symmetric regardless of the internal routing. Because the intrinsic path-length differences between the respective interferometer arms are negligible compared to the $24 \mu\text{m}$ coherence length, the network sustains highly stable interference across all physical modes, ensuring the operational fidelity of the setup.

This macroscopic regime allows us to treat the detection voltages at the photodiode array strictly as classical probabilities, effectively isolating the statistical scaling behavior to the variance introduced by the Haar-random sampling of the




FIG. 2. Sketch of the unit cell. Phase shifters ϕ and θ are denoted by red lines and the 50:50 beam splitter by blue lines.

unitary matrices. Output intensities are subsequently recorded by an array of photodiodes (Thorlabs FGA01FC). Both the photonic processor and the photodiode array are directly interfaced to a control computer. For measurement protocols, we employ Python scripts that leverage the proprietary dogwood library from Quix Quantum—this framework not only supports basic chip operations but also enables the precise configuration of arbitrary unitary matrices.

A unit cell of the MZI, shown in Fig. 2, consists of a phase shifter with an external phase ϕ and an internal interferometer phase θ , together with a pair of 50:50 beam splitters.

The unit cell acts on the input channels as the unitary matrix:

$$S(\theta, \phi) = \frac{1}{2} \begin{pmatrix} 1 - e^{-i\theta} & -i(e^{-i\theta} + 1)e^{-i\phi} \\ -i(e^{-i\theta} + 1)e^{-i\phi} & -(1 - e^{-i\theta})e^{-i\phi} \end{pmatrix}.$$

The photonic chip comprises a vertically arranged structure of four layers, each extending to a depth of eight unit cells. This configuration provides the necessary degrees of freedom to implement arbitrary unitary matrices. By modulating the waveguide phases over the full 2π range, any transformation within the $SU(2)$ unitary group can be realized at the unit cell level (for the mathematical construction, see Ref. [13]). These phase shifters are actuated via the thermo-optic effect using resistive heating; consequently, every unitary matrix maps to a distinct array of control voltages. By composing the transfer matrices of the individual unit cells, one can derive the global transfer and scattering matrices of the full system. This implementation achieves a fidelity of 98%, ensuring that the realized matrices exhibit minimal distortion compared to the target transformations. Detailed specifications of the system are provided in Ref. [14].

B. Measurement protocols

A fundamental distinction of our measurement protocol compared to standard classical shadow tomography lies in the nature of the detection phase. Conventional quantum implementations rely on single-shot projective measurements, which yield a single discrete outcome per experimental run. Such approaches inevitably embed projective shot noise into the acquired dataset. Our experimental setup, conversely, acquires the full probability distribution simultaneously through the continuous detection of macroscopic optical intensities across all output channels. By extracting the exact outcome probabilities in a single measurement cycle, the protocol entirely circumvents the shot noise inherent in quantum state collapse. The statistical variance present in our state reconstruction is thus isolated strictly to the finite Haar-random sampling of the unitary matrices. This methodological choice permits a precise evaluation of the scaling limit, free from the obscuring effects of outcome projection. The experimental procedure functions as a classical analog to quantum shadow tomography. It relies on the precise manipulation of spatial modes to encode quantum information, followed by random unitary rotations and intensity detection. The protocol proceeds in three distinct phases:

(1) *State initialization (injection)*: The input quantum state is defined by the spatial distribution of the laser light injected into the processor. Physically, injecting coherent laser light

solely into the a th input waveguide prepares the computational basis state $|a\rangle$. For example, injection into channel $a = 3$ prepares the state $|\psi\rangle = (0, 0, 1, 0, 0, 0, 0, 0)^T$, corresponding to a density matrix ρ where $\rho_{33} = 1$ and all other elements vanish.

(2) *Unitary evolution*: Using the control library, we apply a specific unitary transformation, U , to the photonic processor. This transformation is realized by tuning the mesh of Mach-Zehnder interferometers. For shadow tomography, these unitaries are drawn from the Haar measure (or a specific experimental design as detailed below).

(3) *Detection and probability assignment*: The output optical intensities are recorded by the photodiode array. Since the detectors measure classical intensity rather than single-photon counts in this regime, we derive the quantum probabilities via normalization. Let V_i denote the voltage measured at the i th output photodiode, which is proportional to the light intensity. The probability of observing the system in state $|b_i\rangle$ (the i th output mode) is calculated as

$$\rho_{ii}^{\text{exp}} = P(b_i) = \frac{V_i}{\sum_{j=1}^8 V_j}. \quad (11)$$

This normalized vector represents the diagonal of the density matrix in the rotated basis, serving as the input for the reconstruction algorithm.

To rigorously benchmark the platform's fidelity, we executed four distinct experimental protocols. These progress from full-dimensional trivial states to complex superpositions embedded in subspaces:

I. *The eight-dimensional trivial state*: This protocol calibrates the system's ground truth. We inject laser light into the first channel, preparing the fundamental state $|1, 0, 0, 0, 0, 0, 0, 0\rangle$. The chip is subjected to a sequence of Haar-random unitaries, $U_{\text{Haar};8 \times 8}$. A successful reconstruction yields the initial projector:

$$\rho_{\text{triv};8 \times 8}^{(\text{rec})} = \rho_{\text{triv};8 \times 8} = \text{diag}(1, 0, 0, 0, 0, 0, 0, 0).$$

II. *The eight-dimensional random state*: To reconstruct an arbitrary coherent superposition, we introduce a fixed randomizing unitary, $U_{\text{rdm};8 \times 8}$. While the input remains channel 1, the effective measurement unitary applied to the chip is the composition

$$U_{\text{meas};8 \times 8} = U_{\text{Haar};8 \times 8} U_{\text{rdm};8 \times 8}.$$

Significantly, the reconstruction algorithm is provided only with the sequence $U_{\text{Haar};8 \times 8}$. If successful, the algorithm reconstructs the state rotated by the hidden randomization matrix:

$$\rho_{\text{rdm};8 \times 8}^{(\text{rec})} = U_{\text{rdm};8 \times 8} \rho_{\text{triv};8 \times 8} U_{\text{rdm};8 \times 8}^\dagger.$$

III. *The four-dimensional trivial state*: We verify control over isolated subspaces by restricting operations to the first four modes. We generate Haar-random unitaries $U_{\text{Haar};4 \times 4}$ and embed them into the full 8×8 processor via the direct sum

$$U_{\text{meas};8 \times 8} = \begin{pmatrix} U_{\text{Haar};4 \times 4} & 0 \\ 0 & \mathbb{I}_{4 \times 4} \end{pmatrix}.$$

Measurements are restricted to the first four output channels, with normalization calculated over that subspace ($\sum_{i=1}^4 V_i$).

This tests the system's ability to decouple from the environment (modes 5–8). The target reconstruction is

$$\rho_{\text{triv};4 \times 4}^{(\text{rec})} = \text{diag}(1, 0, 0, 0).$$

IV. *The four-dimensional random state*: Finally, we combine subspace embedding with arbitrary rotation. We generate a random unitary $U_{\text{rdm};4 \times 4}$ and compose it with the Haar sampling sequence within the upper-left block:

$$U_{\text{meas};8 \times 8} = \begin{pmatrix} U_{\text{Haar};4 \times 4} U_{\text{rdm};4 \times 4} & 0 \\ 0 & \mathbb{I}_{4 \times 4} \end{pmatrix}.$$

By recording the full composition U_{meas} on the device but using only $U_{\text{Haar};4 \times 4}$ for reconstruction data, we recover the rotated state confined strictly to the target subspace:

$$\rho_{\text{rdm};4 \times 4}^{(\text{rec})} = U_{\text{rdm};4 \times 4} \rho_{\text{triv};4 \times 4} U_{\text{rdm};4 \times 4}^\dagger.$$

III. THE STATISTICAL REGIME: IDEAL HAAR SCALING

Prior to deployment on the photonic hardware, we conducted extensive numerical simulations to validate the reconstruction protocols detailed in Sec. II. These simulations emulate the classical analog of shadow tomography by calculating the exact intensity distributions expected from the ideal unitary operations, thereby establishing a theoretical baseline free from hardware imperfections. The primary objective was to substantiate the analytical relationship between the sample size, M , of applied unitary transformations and the reconstruction error, as derived in Appendix C. This step is critical for optimizing the experimental resource budget.

A. Error quantification metrics

Quantifying the deviation between the reconstructed density matrix, $\rho^{(\text{rec})}$, and the target state, ρ , requires a robust metric. While fidelity and trace distance are standard measures in quantum information [15,16], they present specific limitations in this context. Due to finite statistics and numerical precision, the reconstructed matrix $\rho^{(\text{rec})}$ is not guaranteed to be positive semidefinite. Consequently, fidelity calculations can yield unphysical values exceeding unity, rendering them unreliable for precise error scaling. Conversely, the trace distance, while a valid metric, focuses on distinguishability and may not adequately capture the accuracy of individual off-diagonal elements (coherences), which are of central interest in photonic interferometry. Therefore, we employ the Frobenius norm of the difference matrix as our primary figure of merit. It provides a direct, elementwise quantification of the reconstruction accuracy:

$$\|\Delta\rho\|_F \equiv \|\rho^{(\text{rec})} - \rho\|_F = \sqrt{\sum_{i=1}^d \sum_{j=1}^d |\rho_{ij}^{(\text{rec})} - \rho_{ij}|^2}. \quad (12)$$

As derived in Appendix C, the expected value of this error scales inversely with the square root of the measurement count:

$$\mathbb{E}[\|\Delta\rho\|_F^2] = \frac{d \text{Tr}(\rho^2) - 1}{M}. \quad (13)$$

In the context of our measurement protocols, the original density matrix represented a pure state $\rho = |\psi\rangle\langle\psi|$ for which




FIG. 3. Numerical validation of density matrix reconstruction. Left: Reconstruction of an eight-dimensional trivial state. Right: Reconstruction of a four-dimensional random state. The heat maps display the real and imaginary components of the density matrices for the original random state (4×4 case) and after $M = 10$ and $M = 5000$ (both cases) unitary transformations. The numerical values within the cells indicate the reconstructed amplitudes, demonstrating convergence toward the target state.

$\text{Tr}(\rho^2) = 1$, yielding

$$\mathbb{E}[\|\Delta\rho\|_F^2] = \frac{d-1}{M}. \quad (14)$$

B. Simulation results

The simulations mirror the four experimental configurations (I–IV) established previously. By substituting the simulated ideal intensities into the reconstruction formula [Eq. (6)], we verified the protocol for both trivial and random states in four-dimensional and eight-dimensional Hilbert spaces. Figure 3 visualizes the convergence of the reconstruction for an 8×8 trivial state and a 4×4 random state. The heat maps and corresponding matrix elements demonstrate that with $M = 5000$ unitary samples, the reconstructed elements approach the target values with high precision. The quantitative scaling behavior is presented in Fig. 4. Here, we plot the Frobenius norm against the number of unitary transformations M on a logarithmic scale, extending to $M = 10^5$. The simulation data aligns perfectly with the analytical prediction (solid lines), confirming the $\mathcal{O}(M^{-1/2})$ scaling law. These simulations define the statistical regime, where reconstruction fidelity is bounded solely by sample size M . The specific values 0.037 (eight-dimensional) and 0.024 (four-dimensional) represent the asymptotic zero-noise limit. Any deviation from this $\mathcal{O}(M^{-1/2})$ scaling in the physical apparatus identifies the onset of the device regime.

IV. EXPERIMENTAL RESULTS

Figure 5 compares the experimental data with the analytical curves. A clear divergence between the expected and experimental Frobenius norms becomes apparent when averaging 5000 unitary matrices. Specifically, the experimental Frobenius norm saturates to a finite value, whereas the theoretical curve continues to decrease. Consequently, the reconstructed density matrix no longer converges to the input




FIG. 4. Scaling of reconstruction error. The plot shows the dependence of the Frobenius norm on the number of unitary measurements (M) for both four-dimensional and eight-dimensional systems. The simulated data (points) adhere strictly to the analytical prediction $\sqrt{(d-1)/M}$ (lines). The dash-dotted green lines mark the theoretical accuracy limit for the experimentally chosen sample size of $M = 5000$.

density matrix (See Fig. 6). This deviation is not attributable to statistical errors arising from the finite number of measurement samples. Rather, the discrepancy stems from two primary sources. The first source relates to static imperfections in the chip. Implementing unitary matrices on the processor necessitates the fine-tuning of beam splitters and phase shifters. Any imprecision can cause coherent spectral distortion of the realized unitary matrices and their subsequent samples. The second source is processor noise, which induces small, random temporal fluctuations in the unitary matrices.




FIG. 5. Observation of the hardware horizon. The scaling of reconstruction error (Frobenius norm) versus measurement count (M) for four- and eight-dimensional states. The data initially follows the predicted statistical scaling law $\mathcal{O}(M^{-1/2})$ (solid black lines). However, a distinct crossover occurs at a critical sample size (M_{crit}), transitioning from a variance-dominated to a bias-dominated regime where the error diverges from the theoretical prediction and saturates at a hardware-imposed floor. This plateau defines the hardware horizon, governed by the intrinsic spectral distortion of the photonic processor, which prevents convergence to the theoretical accuracy limit (green dash-dotted lines).




FIG. 6. Experimental reconstruction of density matrices via the photonic processor. Left: The density matrix of a trivial state in an eight-dimensional space. Right: The density matrix of a nontrivial state in a four-dimensional space. Colored squares in the heat maps denote matrix elements; real and imaginary parts are plotted separately with values indicated. Reconstructions utilized up to 5000 unitary transformations.

The measured intensities are, therefore, averages over this ensemble of slightly random unitarities. This averaging mechanism inherently breaks unitarity and constitutes the primary source of decoherence in the device. Next, we develop a phenomenological model to account for these effects. This model will allow us to characterize the level of errors stemming from coherent spectral distortion and decoherence.

A. Reconstruction in the presence of coherent spectral distortion and decoherence

The idealized assumptions of the measurement process must be confronted with the physical realization of the unitary operations. In the experimental hardware, arbitrary unitary matrices are synthesized via a Clements decomposition [12], which maps the target transformation U onto an integrated mesh of Mach-Zehnder interferometers and phase shifters. The global transformation is the ordered product of these local operations, $U = \prod_{k=1}^N T_k(\phi_k)$, where ϕ_k are the programmable phases of the individual optical components.

During physical operation, the exact phase imparted by each component deviates from the ideal setting. We decompose this microscopic phase error into a time-independent systematic calibration offset and a zero-mean dynamic temporal fluctuation: $\delta\phi_k = \delta\phi_k^{\text{static}} + \delta\phi_k^{\text{noise}}$. We model the dynamic temporal noise as independent Gaussian random variables $\delta\phi_k^{\text{noise}} \sim \mathcal{N}(0, \sigma_k^2)$.

As detailed in Appendix D, integrating these microscopic fluctuations over the Haar measure rigorously separates the errors into two distinct macroscopic effects. The intended unitary transformation U is physically realized as a perturbed product UU_c [17], where the static coherent spectral distortion U_c represents the ensemble-averaged unitary deviation from the identity. This distortion modifies the measurement probabilities in the computational basis to

$$P(b_i = 1) = (UU_c\rho U_c^\dagger U^\dagger)_{ii}. \quad (15)$$

Simultaneously, the configuration-dependent temporal noise acting across the interferometric mesh is homogenized by the geometric symmetry of the Haar ensemble. This dynamic degradation converges strictly to a global depolarizing channel [1]:

$$\rho' = (1 - p)\rho + \frac{p}{d}\mathbb{I}, \quad (16)$$

where the parameter p quantifies the severity of decoherence. The combined impact of unitary imperfection and depolarization yields the effective measurement probabilities:

$$P(b_i = 1) = (1 - p)(UU_c\rho U_c^\dagger U^\dagger)_{ii} + \frac{p}{d}. \quad (17)$$

If we keep applying the standard reconstruction procedure to these corrupted statistics, the resulting intermediate term becomes

$$\tilde{\rho}_{ab}^{(1)} = (1 - p) \sum_{i,j,k} U_{ai}^* U_{ij} (U_c \rho U_c^\dagger)_{jk} U_{ik}^* U_{ib} + \frac{p}{d} \delta_{ab}. \quad (18)$$

Averaging over the Haar measure of the unitary group subsequently yields

$$I_{ab} = (1 - p) \frac{1}{d+1} [(U_c \rho U_c^\dagger)_{ab} + \delta_{ab}] + \frac{p}{d} \delta_{ab}. \quad (19)$$

Finally, the application of the reconstruction formula (5) reveals the true nature of the recovered state:

$$\tilde{\rho}_{ab} = (1 - p)(U_c \rho U_c^\dagger)_{ab} + \frac{p}{d} \delta_{ab}. \quad (20)$$

Equation (20) reveals the mechanism of the crossover from a variance-dominated regime to a bias-dominated regime. The mathematical distribution governing the sampled unitary matrices remains the strictly invariant Haar measure dU . Instead, it is the effective measurement operator—and consequently the reconstructed state—that undergoes a physical transformation. The term $(1 - p)(U_c \rho U_c^\dagger)$ represents a persistent rotation away from the target state, coupled with an isotropic contraction of the state vector due to depolarization. This combined effect locks the reconstruction into a permanent, shifted orbit, preventing convergence to the ideal target state even in the limit of infinite statistical accumulation.

B. Spectral analysis of reconstructed density matrices

Our analysis begins with the relationship established in Eq. (20), which links the true quantum state to its reconstructed counterpart. A significant theoretical challenge arises here: the reconstructed state $\tilde{\rho}$ is corrupted simultaneously by coherent spectral distortions (U_c) and the depolarizing channel (p). Distinguishing between these two error sources is nontrivial, as one represents a reversible rotation of the basis while the other represents an irreversible loss of information. To resolve this ambiguity and isolate the error parameters, we exploit the spectral properties of the density matrix. Since eigenvalues are invariant under unitary transformations, they are naturally immune to the coherent spectral distortion inherent in U_c . By performing an eigendecomposition of the original density matrix $\rho = W\Lambda W^\dagger$, where W is the unitary matrix of eigenvectors and Λ is the diagonal matrix of eigenvalues λ_i , we can derive a clear structural representation of the noise. Substituting this decomposition into Eq. (20) and

TABLE I. Eigenvalues of the experimentally reconstructed density matrices for the measurement protocols detailed in Sec. II B, utilizing $M = 5000$ unitary matrices for reconstruction.

Protocol	1	2	3	4	5	6	7	8
I	0.90889	-0.00820	-0.00075	0.03102	0.02638	0.02146	0.01411	0.00707
II	0.90105	0.04071	0.02831	0.01643	0.01461	0.00754	0.00171	-0.01038
III	0.95372	0.00669	0.02906	0.01958				
IV	0.98087	0.02503	0.01116	-0.01150				

applying the identity $U_c W W^\dagger U_c^\dagger = \mathbb{I}$ reveals the simplified form

$$\tilde{\rho} = U_c W D W^\dagger U_c^\dagger, \quad (21)$$

where the diagonal matrix D is given by

$$D = (1 - p)\Lambda + \frac{p}{d}\mathbb{I}. \quad (22)$$

The eigenvalues of the reconstructed state, $D_i = (1 - p)\lambda_i + p/d$, effectively capture the interplay between the ideal state's spectrum and the depolarizing channel. This theoretical framework provides a direct method for extracting p , independent of coherent spectral distortion. It is important to define the scope and limitations of this phenomenological model. While we use the term hardware horizon to describe the observed saturation of the reconstruction error, this threshold represents a device-dependent error floor rather than an immutable law of physics. Modeling hardware imperfections via a single, fixed unitary distortion U_c is an effective, ensemble-averaged approximation over the Haar measure. In physical reality, distortions within the photonic chip inherently depend on the specific dynamic phase settings of the unit cells for each applied unitary operation. Consequently, U_c functions as a statistical abstraction that captures the mean coherent deviation of the processor rather than a strict, universally static hardware constant. However, a practical complication remains. Experimental reconstructions are derived from finite datasets, introducing statistical fluctuations. This noise disproportionately affects the smaller eigenvalues, rendering the lower end of the spectrum unreliable for precise estimation. To mitigate this, we restrict our analysis to the most robust component of the spectrum: the leading eigenvalue. As this value is least susceptible to statistical noise, it offers the most dependable basis for calculating the error parameter. Using the leading eigenvalue, we approximate the depolarizing channel parameter as

$$p \approx \frac{\lambda_1 - D_1}{\lambda_1 - 1/d}. \quad (23)$$

In the context of our measurement protocols, the original density matrix represented a pure state $\rho = |\psi\rangle\langle\psi|$. Consequently, the theoretical spectrum consists of a single dominant eigenvalue $\lambda_1 = 1$, with the remaining $d - 1$ eigenvalues vanishing. Under these conditions, the estimator simplifies to

$$p \approx (1 - D_1) \frac{d - 1}{d}. \quad (24)$$

The eigenvalues obtained from the experimentally reconstructed density matrices are presented in Table I.

Based on the dominant eigenvalues, we estimate the depolarizing channel parameters summarized in Table II. It is notable that the depolarization parameter is significantly higher for the 8×8 matrices compared to the 4×4 cases. This discrepancy arises because the larger matrices necessitate the involvement of more unit cells, thereby compounding the depolarization effect of individual cells within the measurement results.

C. Frobenius norm and the isolation of coherent spectral distortion

Figure 5 unveils a critical limitation in the reconstruction process: simply increasing the ensemble size M of Haar random unitaries eventually ceases to yield accuracy gains. Instead of converging to zero, the Frobenius norm arrests at a saturation plateau. This persistent residual error compels us to ask what factors prevent perfect reconstruction, even in the limit of infinite sampling. We resolve this by dissecting the MSE. This phenomenological framework allows us to isolate transient statistical fluctuations from the systematic errors induced by decoherence and coherent spectral distortion:

$$\mathbb{E}[\|\tilde{\rho}^{(M)} - \rho\|_F^2] = \mathbb{E}[\|(\tilde{\rho}^{(M)} - \tilde{\rho}) + (\tilde{\rho} - \rho)\|_F^2]. \quad (25)$$

Here, $\tilde{\rho}^{(M)}$ represents the reconstruction from M unitaries, while $\tilde{\rho}$ denotes the physically measured, distorted state [as defined in Eq. (20)], and ρ is the ideal input state. Since $\tilde{\rho}^{(M)}$ acts as an unbiased estimator of $\tilde{\rho}$, the cross terms vanish ($\mathbb{E}[\tilde{\rho}^{(M)} - \tilde{\rho}] = 0$), allowing us to separate the MSE into two distinct terms:

$$\mathbb{E}[\|\tilde{\rho}^{(M)} - \rho\|_F^2] = \mathbb{E}[\|\tilde{\rho}^{(M)} - \tilde{\rho}\|_F^2] + \mathbb{E}[\|\tilde{\rho} - \rho\|_F^2]. \quad (26)$$

The first term represents the statistical error of estimation. As derived in Eq. (13), this error scales inversely with the sample

TABLE II. Depolarizing channel parameters calculated from the largest eigenvalues in Table I via $p \approx (1 - D_1) \frac{d-1}{d}$ (first column) and the average magnitude of the coherent distortion matrix elements $\epsilon = [\frac{\text{slope} - p^2(1-\frac{1}{d})}{2(1-p)(d-1)}]^{1/2}$ calculated from the slopes in Fig. 7 (second column).

Protocol	p	ϵ
I	0.10412	0.01198
II	0.11308	0.01164
III	0.06170	0.01297
IV	0.02550	0.01271

size:

$$\mathbb{E}[\|\tilde{\rho}^{(M)} - \tilde{\rho}\|_F^2] = \frac{d \operatorname{Tr}(\tilde{\rho}^2) - 1}{M}. \quad (27)$$

Substituting our distortion model (20) and considering our experimental protocol where the input $\rho = |\psi\rangle\langle\psi|$ is a pure state, this simplifies to a vanishing contribution:

$$\mathbb{E}[\|\tilde{\rho}^{(M)} - \tilde{\rho}\|_F^2] = (1-p)^2 \frac{d-1}{M}. \quad (28)$$

As M increases, this term disappears. Yet, the saturation observed in Fig. 5 remains. This brings us to the second term in Eq. (26), the dominant factor responsible for the leveling off of the Frobenius norm. Independent of measurement quantity, this term contains the structural imperfections of the system:

$$\begin{aligned} \mathbb{E}[\|\tilde{\rho} - \rho\|_F^2] &= p^2 \left[\operatorname{Tr}(\rho^2) - \frac{1}{d} \right] \\ &\quad + 2(1-p)[\operatorname{Tr}(\rho^2) - \operatorname{Tr}(U_c \rho U_c^\dagger \rho)]. \end{aligned} \quad (29)$$

Here we can isolate the decoherence contribution (proportional to p^2) and the coherent spectral distortion:

$$\mathcal{E}_c = \operatorname{Tr}(\rho^2) - \operatorname{Tr}(U_c \rho U_c^\dagger \rho). \quad (30)$$

For our pure state protocol, the total systematic error crystallizes into

$$\mathbb{E}[\|\tilde{\rho} - \rho\|_F^2] = p^2 \left(1 - \frac{1}{d} \right) + 2(1-p)[1 - |\langle\psi|U_c|\psi\rangle|^2]. \quad (31)$$

This result offers a direct probe into the unitary error matrix U_c . By working in a basis where $|\psi\rangle$ aligns with the first eigenbasis vector $|100\dots\rangle$, the term $\langle\psi|U_c|\psi\rangle$ corresponds to the diagonal element $\{U_c\}_{11}$. Utilizing the unitarity condition $\sum_j^d |\{U_c\}_{1j}|^2 = 1$, we arrive at a precise definition of the coherent spectral distortion in terms of the matrix's off-diagonal elements:

$$\mathcal{E}_c = 1 - |\langle\psi|U_c|\psi\rangle|^2 = \sum_{j \neq 1}^d |\{U_c\}_{1j}|^2. \quad (32)$$

Consequently, the hardware horizon is not merely an accuracy limit; it is a robust spectral fingerprint of the specific processor. Equation (32) demonstrates that the saturation level is invariant under unitary randomization, depending solely on the mean off-diagonal elements of the coherent distortion matrix U_c . Consequently, the saturation of the Frobenius norm is not merely a limit to accuracy; it is a diagnostic tool. It reveals that the coherent spectral distortion is exactly the sum of squares of the off-diagonal ($j \neq 1$) elements of U_c . The mean square of the off-diagonal elements is thus given by

$$\epsilon^2 = \mathbb{E}[|\{U_c\}_{1j}|^2] = \frac{1}{d-1} \sum_{j \neq 1}^d |\{U_c\}_{1j}|^2 = \frac{\mathcal{E}_c}{d-1}, \quad (33)$$

where ϵ is the average magnitude of the coherent distortion matrix elements.

D. Consistency of the scaling law

To isolate the true coherent spectral distortion from the noise, we examine the structure of the MSE curves. By plot-




FIG. 7. Robust spectral fingerprint of the photonic processor. The scaled reconstruction error $M\|\tilde{\rho}^{(M)} - \rho\|_F^2$ as a function of sample size M for the four experimental protocols. The strict linearity confirms the validity of the phenomenological error model [Eq. (31)], separating statistical fluctuations from systematic hardware distortion. While the slopes vary significantly due to state-dependent decoherence (p), the underlying coherent spectral distortion extracted from these slopes remains highly consistent. In all regimes, the static unitary distortion converges to a mean off-diagonal element magnitude of $\epsilon \approx 0.012$ (see Table II), establishing this value as the operational hardware horizon of the device. The panels correspond to protocol I (top left) slope = $1.12873 \times 10^{-2} \pm 1.82596 \times 10^{-5}$, protocol II (top right) slope = $1.18221 \times 10^{-2} \pm 1.53545 \times 10^{-5}$, protocol III (bottom left) slope = $3.80301 \times 10^{-3} \pm 1.06298 \times 10^{-5}$, and protocol IV (bottom right) slope = $1.43278 \times 10^{-3} \pm 9.30978 \times 10^{-6}$.

ting the scaled error $M\|\tilde{\rho}^{(M)} - \rho\|_F^2$ as a function of M , a linear relationship emerges (Fig. 7), characterized by the slope defined in Eq. (31). This linearity allows us to pierce through the measurement uncertainties. Utilizing our established estimates for p (Table II), we determine the magnitude of the coherent distortion, ϵ , via the relation

$$\epsilon = \left[\frac{\text{Slope} - p^2(1-1/d)}{2(1-p)(d-1)} \right]^{1/2}. \quad (34)$$

The visualization of these slopes in Fig. 7 and the calculated average magnitudes in Table II reveal consistency. While the observed slopes and depolarization parameters p fluctuate by orders of magnitude—driven by the differing complexity of the target states—the derived coherent spectral distortion exhibits a rigid invariance. In all scenarios, the static unitary distortion converges to a mean off-diagonal element magnitude of $\epsilon \approx 0.012$. This consistency confirms that the Hardware Horizon is not an artifact of specific measurement protocols but a robust, state-independent spectral fingerprint of the specific photonic architecture. Consequently, ϵ represents the effective resolution limit of the hardware, defining the boundary where device physics inevitably supersedes statistical scaling.

V. CONCLUSION

In this work, we have experimentally established the existence of an operational hardware horizon—a hardware-induced systematic error floor that strictly limits the efficacy of classical shadow tomography on integrated photonic processors. By pushing the measurement count M beyond the standard statistical limits, we observed a distinct crossover where the theoretical $\mathcal{O}(M^{-1/2})$ scaling collapses into a saturation regime governed entirely by device physics, marking a transition from a variance-dominated regime to a bias-dominated regime. This discovery reframes the challenge of NISQ-era tomography. We have demonstrated that the primary obstacle to high-fidelity reconstruction is not the acquisition of statistics but the spectral distortion of the unitary implementation itself. Our phenomenological model reveals that this horizon is defined by a specific interplay between static coherent spectral distortion (ϵ) and dynamic decoherence (p), creating a device-specific error floor that no amount of additional sampling can overcome. Consequently, the standard promise of shadow tomography polynomial resource scaling is effectively voided once the system reaches this operational hardware limit. This realization creates an urgent imperative for the field: to move beyond passive error characterization and toward active spectral compensation. The open challenge is no longer just about measuring faster but about whether we can develop algorithmic inversions that learn the distorted unitary measure in real time. To breach this horizon, we propose the development of structure-aware reconstruction kernels that abandon the assumption of an ideal Haar ensemble. Instead, by incorporating the experimentally learned distortion matrix U_c characterized via independent transfer-matrix inversion or efficient process tomography directly into the Weingarten integration, it becomes possible to convert static coherent spectral distortion from a device-dependent saturation limit into a corrigible systematic bias, effectively extending the reach of shadow tomography into the sub-horizon regime.

ACKNOWLEDGMENT

This research was carried out with the support of the Ministry of Culture and Innovation, funded by the National Research Development and Innovation Fund, under Project No. 2022-2.1.1-NL-2022-00004.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not techni-

cally feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: DENSITY MATRIX RECONSTRUCTION

For a d -dimensional quantum system, let the input state be described by the density matrix ρ . Suppose a unitary transformation U is applied to this state. In that case, the probability of observing an outcome b_i (e.g., detecting a particle in the i th output channel of a photonic circuit) is given by the Born rule. For instance, if we are interested in the diagonal elements of the transformed state $U\rho U^\dagger$, which correspond to measurement probabilities in the computational basis, we have

$$P(b_i = 1) = (U\rho U^\dagger)_{ii} = \sum_{j,k} U_{ij} \rho_{jk} U_{ik}^* = \rho_{ii}^{\text{exp}} \quad (\text{A1})$$

where U_{ik}^* is the complex conjugate of U_{ik} . From a single measurement [or a set of measurements yielding probabilities $P(b_i = 1)$ for a fixed U], we can form a diagonal matrix, ρ^{exp} , where $\rho_{ii}^{\text{exp}} = P(b_i = 1)$ as stated in (A1). An estimate of the original density matrix element ρ_{ab} after one such measurement procedure (applying U and measuring probabilities) can be constructed as

$$\hat{\rho}_{ab}^{(1)} = \sum_i U_{ai}^* \rho_{ii}^{\text{exp}} U_{ib} = \sum_i U_{ai}^* \left(\sum_{j,k} U_{ij} \rho_{jk} U_{ik}^* \right) U_{ib}. \quad (\text{A2})$$

Rearranging this expression, we get

$$\hat{\rho}_{ab}^{(1)} = \sum_{i,j,k} U_{ai}^* U_{ij} \rho_{jk} U_{ik}^* U_{ib}. \quad (\text{A3})$$

To obtain a more accurate reconstruction of the original density matrix ρ , we must average the results obtained from many different unitary transformations U , drawn from a suitable distribution, in our case the Haar distribution. Let I_{ab} be the quantity obtained by averaging $\hat{\rho}_{ab}^{(1)}$ over the unitary group with respect to the Haar measure dU :

$$I_{ab} = \mathbb{E}_U[\hat{\rho}_{ab}^{(1)}] = \int dU \sum_{i,j,k} U_{ai}^* U_{ij} \rho_{jk} U_{ik}^* U_{ib}. \quad (\text{A4})$$

The evaluation of this integral involves terms of the form $\int dU U_{a'b'} U_{c'd'} U_{e'f'}^* U_{g'h'}^*$, which are known as Weingarten functions [11]. The general formula for calculating such integrals is given by

$$\int_{U_d} U_{i_1 j_1} \cdots U_{i_q j_q} U_{i'_1 j'_1}^* \cdots U_{i'_q j'_q}^* dU = \sum_{\sigma, \tau \in S_q} \delta_{i'_1 i_{\sigma(1)}} \cdots \delta_{i'_q i_{\sigma(q)}} \delta_{j_1 j_{\tau(1)}} \cdots \delta_{j_q j_{\tau(q)}} Wg(\sigma \tau^{-1}, d), \quad (\text{A5})$$

where $Wg(\sigma, d)$ denotes the Weingarten functions:

$$Wg(\sigma, d) = \frac{1}{q!^2} \sum_{\lambda} \frac{\chi^\lambda(e)^2 \chi^\lambda(\sigma)}{s_{\lambda, d}(1)}, \quad (\text{A6})$$

where d is the dimension of our unitary matrices. In our case for Eq. (A3), the general formula is given by

$$\int dU U_{i_1 j_1} U_{i_2 j_2} U_{i'_1 j'_1}^* U_{i'_2 j'_2}^* = (\delta_{i_1 i'_1} \delta_{i_2 i'_2} \delta_{j_1 j'_1} \delta_{j_2 j'_2} + \delta_{i_1 i'_2} \delta_{i_2 i'_1} \delta_{j_1 j'_2} \delta_{j_2 j'_1}) Wg(1^2, d) \\ - (\delta_{i_1 i'_1} \delta_{i_2 i'_2} \delta_{j_1 j'_2} \delta_{j_2 j'_1} + \delta_{i_1 i'_2} \delta_{i_2 i'_1} \delta_{j_1 j'_1} \delta_{j_2 j'_2}) Wg(2, d), \quad (\text{A7})$$

where

$$Wg(1^2, d) = \frac{1}{d^2 - 1}, \quad Wg(2, d) = \frac{-1}{d(d^2 - 1)}, \quad (\text{A8})$$

and result in

$$\int dU U_{i_1 j_1} U_{i_2 j_2} U_{i'_1 j'_1}^* U_{i'_2 j'_2}^* = \frac{\delta_{i_1 i'_1} \delta_{i_2 i'_2} \delta_{j_1 j'_1} \delta_{j_2 j'_2} + \delta_{i_1 i'_2} \delta_{i_2 i'_1} \delta_{j_1 j'_2} \delta_{j_2 j'_1}}{d^2 - 1} - \frac{\delta_{i_1 i'_1} \delta_{i_2 i'_2} \delta_{j_1 j'_2} \delta_{j_2 j'_1} + \delta_{i_1 i'_2} \delta_{i_2 i'_1} \delta_{j_1 j'_1} \delta_{j_2 j'_2}}{d(d^2 - 1)}. \quad (\text{A9})$$

The symmetries of the problem allow us to rearrange the indices in Eq. (A9) and arrive at the arrangement of (A4), eventually getting the exact solution to the integral in question:

$$\int dU U_{ia}^* U_{ij} U_{ik}^* U_{ib} = \frac{1}{d(d+1)} (\delta_{ja} \delta_{bk} + \delta_{jk} \delta_{ba}). \quad (\text{A10})$$

The result for the specific fourth-order moment in Eq. (A4), after summing over index i (which introduces a factor of d , the dimension of the Hilbert space), the integral simplifies significantly. The relevant integral is

$$\sum_i \int dU U_{ai}^* U_{ij} U_{ik}^* U_{ib} = \frac{1}{d+1} (\delta_{aj} \delta_{kb} + \delta_{ak} \delta_{jb}). \quad (\text{A11})$$

Summing over i in Eq. (A4) using this result (and reindexing summation variables for clarity) yields

$$I_{ab} = \sum_{j,k} \frac{1}{d+1} (\delta_{aj} \delta_{kb} + \delta_{ak} \delta_{jb}) \rho_{jk}. \quad (\text{A12})$$

This simplifies to

$$I_{ab} = \frac{1}{d+1} (\rho_{ab} + \delta_{ab} \text{Tr}(\rho)). \quad (\text{A13})$$

For a physical density matrix, the trace $\text{Tr}(\rho) = 1$. Therefore, we can express the elements of the original density matrix ρ_{ab} in terms of the averaged quantity I_{ab} :

$$\rho_{ab} = (d+1)I_{ab} - \delta_{ab}. \quad (\text{A14})$$

In practice, I_{ab} is estimated by averaging $\rho_{ab}^{(k)} = \sum_i U_{ai}^* \rho_{ii}^{\text{exp}} U_{ib}$ [from Eq. (A3) for each unitary $U^{(k)}$] over a finite number M of unitary transformations randomly chosen from a Haar distribution:

$$\overline{I_{ab}}^{(M)} = \frac{1}{M} \sum_{k=1}^M \rho_{ab}^{(k)} = \mathbb{E}_U^{(M)} \left[\sum_i U_{ai}^* \rho_{ii}^{\text{exp}} U_{ib} \right]. \quad (\text{A15})$$

The reconstructed matrix is then

$$\rho_{ab}^{(\text{rec})} = (d+1) \overline{I_{ab}}^{(M)} - \delta_{ab}. \quad (\text{A16})$$

The accuracy of this reconstruction depends on M and the dimension d , shown later in Appendix C.

APPENDIX B: GENERATION OF HAAR MEASURE UNITARY MATRICES

The reconstruction formalism detailed in Sec. II, particularly the use of Weingarten functions for averaging, mandates that the unitary transformations U are drawn from a uniform distribution over the unitary group, according to the Haar measure. To generate such $d \times d$ unitary matrices, we employ the following standard algorithm [18]:

(1) Generate an $d \times d$ complex matrix Z , where each element $Z_{ij} = x_{ij} + iy_{ij}$. The real part x_{ij} and imaginary part y_{ij} are independent random variables drawn from a standard normal distribution $\mathcal{N}(0, 1)$.

(2) Perform a QR decomposition of Z such that $Z = QR$, where Q is a unitary matrix and R is an upper triangular matrix.

(3) Construct a diagonal matrix D where each diagonal element D_{kk} is the phase of the corresponding diagonal element of R : $D_{kk} = R_{kk}/|R_{kk}|$ for $R_{kk} \neq 0$, and $D_{kk} = 1$ if $R_{kk} = 0$.

(4) The resulting unitary matrix $U = QD$ is then distributed according to the Haar measure.

APPENDIX C: VARIANCE OF RECONSTRUCTION

In the photonic implementation of the shadow tomography, we do not obtain a single discrete outcome. Instead, we measure the intensities at all d output ports simultaneously. For a given Haar-random unitary U acting on the state ρ , the probability of detecting a photon in port k is given by the diagonal element of the rotated density matrix:

$$p_k = \langle k | U \rho U^\dagger | k \rangle = (U \rho U^\dagger)_{kk}. \quad (\text{C1})$$

We construct the single-snapshot estimator $\hat{p}$ using the full distribution $\{p_k\}_{k=1}^d$:

$$\hat{p} = (d+1) \sum_{k=1}^d p_k U^\dagger |k\rangle \langle k| U - \mathbb{I} = (d+1)A - \mathbb{I}, \quad (\text{C2})$$

where we have defined the matrix $A = U^\dagger \text{diag}(p)U$.

1. Expectation value (first moment)

We first verify that the estimator is unbiased. The matrix A can be written as

$$A = \sum_{k=1}^d \langle k|U \rho U^\dagger|k\rangle U^\dagger|k\rangle\langle k|U. \quad (\text{C3})$$

Averaging over the Haar measure dU :

$$\mathbb{E}[A] = \int dU \sum_k \langle k|U \rho U^\dagger|k\rangle U^\dagger|k\rangle\langle k|U = \mathcal{D}(\rho), \quad (\text{C4})$$

where $\mathcal{D}(\rho) = \frac{\rho + \mathbb{I}}{d+1}$. Thus,

$$\mathbb{E}[\hat{\rho}] = (d+1) \left(\frac{\rho + \mathbb{I}}{d+1} \right) - \mathbb{I} = \rho. \quad (\text{C5})$$

2. Total reconstruction error (sum of matrix elements)

To find the global error, we calculate the MSE in the Frobenius norm for a single snapshot. Since the estimator is unbiased, the MSE is the trace of the covariance:

$$\text{MSE} = \mathbb{E}[\|\hat{\rho} - \rho\|_F^2] = \mathbb{E}[\text{Tr}(\hat{\rho}^2)] - \text{Tr}(\rho^2). \quad (\text{C6})$$

First, we compute $\text{Tr}(\hat{\rho}^2)$:

$$\hat{\rho}^2 = [(d+1)A - \mathbb{I}]^2 = (d+1)^2 A^2 - 2(d+1)A + \mathbb{I}, \quad (\text{C7})$$

$$\text{Tr}(\hat{\rho}^2) = (d+1)^2 \text{Tr}(A^2) - 2(d+1)\text{Tr}(A) + d. \quad (\text{C8})$$

Note that $\text{Tr}(A) = \sum_k p_k = 1$. Now we analyze A^2 :

$$A^2 = \left(\sum_k p_k U^\dagger|k\rangle\langle k|U \right) \left(\sum_j p_j U^\dagger|j\rangle\langle j|U \right). \quad (\text{C9})$$

Using the orthogonality $U^\dagger|k\rangle\langle k|U U^\dagger|j\rangle\langle j|U = \delta_{kj} U^\dagger|k\rangle\langle k|U$, we simplify

$$A^2 = \sum_k p_k^2 U^\dagger|k\rangle\langle k|U \Rightarrow \text{Tr}(A^2) = \sum_{k=1}^d p_k^2. \quad (\text{C10})$$

Thus, the problem reduces to finding the average purity of the output distribution:

$$\mathbb{E}[\text{Tr}(A^2)] = \mathbb{E} \left[\sum_k \langle \langle k|U \rho U^\dagger|k\rangle \rangle^2 \right]. \quad (\text{C11})$$

Using standard Weingarten calculus for the second moment of Haar random unitary entries, for any state ρ :

$$\mathbb{E} \left[\sum_k p_k^2 \right] = \frac{1 + \text{Tr}(\rho^2)}{d+1}. \quad (\text{C12})$$

Substituting this back into the trace equation:

$$\mathbb{E}[\text{Tr}(\hat{\rho}^2)] = (d+1)\text{Tr}(\rho^2) - 1. \quad (\text{C13})$$

Finally, the MSE for a single snapshot is

$$\text{MSE} = [(d+1)\text{Tr}(\rho^2) - 1] - \text{Tr}(\rho^2) = d \text{Tr}(\rho^2) - 1. \quad (\text{C14})$$

For M snapshots, the variance scales by $1/M$:

$$\mathbb{E}[\|\hat{\rho}^{(M)} - \rho\|_F^2] = \frac{d \text{Tr}(\rho^2) - 1}{M}. \quad (\text{C15})$$

Special case (pure state): If ρ is pure ($\text{Tr}(\rho^2) = 1$), this simplifies to

$$\text{MSE} = \frac{d-1}{M}. \quad (\text{C16})$$

This $1/M$ variance characterizes the statistical regime. The failure of this term to vanish in the experimental data (Fig. 5) constitutes the mathematical definition of the hardware horizon.

APPENDIX D: DERIVATION OF THE EFFECTIVE MEASUREMENT PROBABILITIES FROM HETEROGENEOUS STATISTICAL NOISE

To establish the physical origin of the depolarizing channel and the effective measurement probability expression [Eq. (17)], we analyze the effect of dynamic statistical noise traversing the integrated photonic processor. Let the ideal target unitary matrix be synthesized via the Clements decomposition [12] as the ordered product of local transformations:

$$U = \prod_{k=1}^N T_k(\phi_k). \quad (\text{D1})$$

During physical execution, the exact phase imparted by each component deviates from the ideal setting. We decompose this microscopic phase error into a time-independent systematic calibration offset and a zero-mean dynamic temporal fluctuation:

$$\delta\phi_k = \delta\phi_k^{\text{static}} + \delta\phi_k^{\text{noise}}. \quad (\text{D2})$$

We model the dynamic temporal noise as independent Gaussian random variables with nonuniform variances, $\delta\phi_k^{\text{noise}} \sim \mathcal{N}(0, \sigma_k^2)$, characterizing heterogeneous physical noise sources such as thermal variance or electronic control jitter. The physical transformation realized by the hardware in a single operational snapshot is therefore the perturbed product:

$$\tilde{U} = \prod_{k=1}^N T_k(\phi_k + \delta\phi_k^{\text{static}} + \delta\phi_k^{\text{noise}}). \quad (\text{D3})$$

By isolating the target phase from the errors, we can express each local transformation using its Hermitian phase generator G_k . To factor out the ideal target unitary U , we commute the error generators to the right side of the ordered product. This commutation dresses the local generators with the subsequent ideal phase transformations, yielding effective configuration-dependent error generators:

$$\tilde{G}_k(\phi) = \left(\prod_{j=k}^N T_j(\phi_j) \right)^\dagger G_k \left(\prod_{j=k}^N T_j(\phi_j) \right). \quad (\text{D4})$$

The complete physical unitary matrix separates into the ideal transformation and an exponential containing both static and dynamic error terms. Ignoring higher-order commutators

between the small error generators, we decouple the static distortion from the dynamic noise:

$$\tilde{U} \approx U \exp \left(i \sum_{k=1}^N \delta\phi_k^{\text{static}} \tilde{G}_k(\phi) \right) \exp \left(i \sum_{k=1}^N \delta\phi_k^{\text{noise}} \tilde{G}_k(\phi) \right). \quad (\text{D5})$$

We define the configuration-dependent static distortion as $U_c(\phi) = \exp(i \sum_k \delta\phi_k^{\text{static}} \tilde{G}_k(\phi))$ and the dynamic noise operator as $V(\phi) = \exp(i \sum_k \delta\phi_k^{\text{noise}} \tilde{G}_k(\phi))$. For a given target phase configuration ϕ , the exact probability of detecting a photon in the i th output mode is dictated by the Born rule:

$$P(b_i = 1|\phi) = (UU_c(\phi)V(\phi)\rho V^\dagger(\phi)U_c^\dagger(\phi)U^\dagger)_{ii}. \quad (\text{D6})$$

The first statistical step requires evaluating the expectation value over the dynamic temporal noise occurring during the finite measurement window. Expanding the noise operator $V(\phi)$ to second order and applying the expectation operator $\mathbb{E}_{\text{noise}}$ over the Gaussian distributions isolates the surviving terms. The linear terms vanish identically due to the zero-mean nature of the fluctuations. The independence of the spatial noise sources removes all cross terms, satisfying $\mathbb{E}[\delta\phi_k^{\text{noise}} \delta\phi_j^{\text{noise}}] = \sigma_k^2 \delta_{kj}$. This reduces the noise interaction to a configuration-dependent dephasing map $\mathcal{N}_\phi(\rho)$ acting on the state

$$\mathcal{N}_\phi(\rho) \approx \rho - \frac{1}{2} \sum_{k=1}^N \sigma_k^2 [\tilde{G}_k(\phi), [\tilde{G}_k(\phi), \rho]]. \quad (\text{D7})$$

The second statistical step encompasses the classical shadow protocol, which necessitates averaging over an extensive ensemble of random unitary operations drawn from the Haar measure. Uniformly sampling the target matrix U from the $\mathcal{U}(d)$ group mandates that the required phase settings ϕ continuously scan the entire parameter space. Averaging the density matrix over this ensemble defines the global effective quantum channel:

$$\Phi(\rho) = \int d\phi P(\phi) \mathcal{N}_\phi(\rho). \quad (\text{D8})$$

Because the Haar measure is left and right invariant, the ensemble of intermediate matrices within $\tilde{G}_k(\phi)$ rotates each local hardware generator isotropically across the entire Hilbert space. Integrating this multimode dephasing double commutator over a unitarily invariant distribution constitutes a fundamental twirling operation. According to Schur's lemma, twirling projects any quantum channel onto the commutant of the unitary group, inevitably contracting the state toward the maximally mixed state. Consequently, the integration of each individual k th term yields a discrete depolarizing channel with a severity proportional to its specific physical variance σ_k^2 . Because the sum of depolarizing channels is itself a depolarizing channel, the spatially heterogeneous noise collapses entirely into a single, homogeneous global expression:

$$\Phi(\rho) = (1 - p)\rho + \frac{p}{d} \mathbb{I}. \quad (\text{D9})$$

The macroscopic depolarization parameter p cleanly absorbs the microscopic Gaussian variances σ_k^2 , totally removing the error model's dependence on any specific spatial index k . Simultaneously, the Haar average projects the configuration-dependent static distortion $U_c(\phi)$ onto its dominant, nonvanishing mean component $U_c = \mathbb{E}_{\text{Haar}}[U_c(\phi)]$. This extracted matrix represents the persistent, configuration-independent coherent rotational bias of the hardware.

Substituting the macroscopic depolarizing channel and the ensemble-averaged static distortion into the probability expression evaluates the combined measurement statistics. Because the maximally mixed term is invariant under any subsequent unitary rotation, the application of UU_c leaves it unaffected:

$$P(b_i = 1) = \left(UU_c \left[(1 - p)\rho + \frac{p}{d} \mathbb{I} \right] U_c^\dagger U^\dagger \right)_{ii}. \quad (\text{D10})$$

Expanding the linear terms within the diagonal element immediately yields the final expected distribution utilized in our phenomenological reconstruction framework:

$$P(b_i = 1) = (1 - p)(UU_c \rho U_c^\dagger U^\dagger)_{ii} + \frac{p}{d}. \quad (\text{D11})$$

- [1] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [2] Z. Hradil, Quantum-state estimation, *Phys. Rev. A* **55**, R1561 (1997).
- [3] R. T. Thew, K. Nemoto, A. G. White, and W. J. Munro, Qudit quantum-state tomography, *Phys. Rev. A* **66**, 012303 (2002).
- [4] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Quantum state tomography via compressed sensing, *Phys. Rev. Lett.* **105**, 150401 (2010).
- [5] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, Efficient quantum state tomography, *Nat. Commun.* **1**, 149 (2010).
- [6] S. Aaronson, Shadow tomography of quantum states, in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2018), pp. 325–338.
- [7] S. Slussarenko and G. J. Pryde, Photonic quantum information processing: A concise review, *Appl. Phys. Rev.* **6**, 041303 (2019).
- [8] S. Shekhar, W. Bogaerts, L. Chrostowski, J. E. Bowers, M. Hochberg, R. Soref, and B. J. Shastri, Roadmapping the next generation of silicon photonics, *Nat. Commun.* **15**, 751 (2024).
- [9] C. G. Roeloffzen, M. Hoekman, E. J. Klein, L. S. Wevers, R. B. Timens, D. Marchenko, D. Geskus, R. Dekker, A. Alippi, and R. Grootjans, Low-loss Si3N4 TriPleX optical waveguides: Technology and applications overview, *IEEE J. Sel. Top. Quantum Electron.* **24**, 1 (2018).
- [10] Quix quantum, <https://www.quixquantum.com>, accessed November 28, 2025.
- [11] B. Collins, S. Matsumoto, and J. Novak, The Weingarten calculus, *Not. Am. Math. Soc.* **69**, 734 (2022).
- [12] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walmsley, Optimal design for universal multiport interferometers, *Optica* **3**, 1460 (2016).

- [13] C. Taballione, R. van der Meer, H. J. Snijders, P. Hooijschuur, J. P. Epping, M. de Goede, B. Kassenberg, P. Venderbosch, C. Toebes, and H. van den Vlekert, A universal fully reconfigurable 12-mode quantum photonic processor, *Mater. Quantum Technol.* **1**, 035002 (2021).
- [14] C. Taballione, T. A. W. Wolterink, J. Lugani, A. Eckstein, B. A. Bell, R. Grootjans, I. Visscher, D. Geskus, C. G. H. Roeloffzen, J. J. Renema, I. A. Walmsley, P. W. H. Pinkse, and K.-J. Boller, 8×8 reconfigurable quantum photonic processor based on silicon nitride waveguides, *Opt. Express* **27**, 26842 (2019).
- [15] R. Jozsa, Fidelity for mixed quantum states, *J. Mod. Opt.* **41**, 2315 (1994).
- [16] M. A. Nielsen and I. L. Chuang, *Quantum Information and Quantum Computation* (Cambridge University Press, Cambridge, 2000).
- [17] A. Strikis, S. C. Benjamin, and B. J. Brown, Quantum computing is scalable on a planar array of qubits with fabrication defects, *Phys. Rev. Appl.* **19**, 064081 (2023).
- [18] F. Mezzadri, How to generate random matrices from the classical compact groups, *Not. Am. Math. Soc.* **54**, 592 (2007).