

Bayesian tomography of high-dimensional on-chip biphoton frequency combs with randomized measurements

Hsuan-Hao Lu,^{1,*} Karthik V. Myilswamy,^{1,*} Ryan S. Bennink,² Suparna Seshadri,¹ Mohammed S. Alshaykh,³ Junqiu Liu,⁴ Tobias J. Kippenberg,⁴ Daniel E. Leaird,¹ Andrew M. Weiner,¹ and Joseph M. Lukens^{5,†}

¹*School of Electrical and Computer Engineering and Purdue Quantum Science and Engineering Institute, Purdue University, West Lafayette, Indiana 47907, USA*

²*Quantum Computational Science Group, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA*

³*Electrical Engineering Department, King Saud University, Riyadh 11421, Saudi Arabia*

⁴*Institute of Physics, Swiss Federal Institute of Technology Lausanne (EPFL), 1015 Lausanne, Switzerland*

⁵*Quantum Information Science Group, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA*

(Dated: January 25, 2022)

Owing in large part to the advent of integrated biphoton frequency combs (BFCs) [1], recent years have witnessed increased attention to quantum information processing in the frequency domain for its inherent high dimensionality and entanglement compatible with fiber-optic networks. Quantum state tomography (QST) of such states, however, has required complex and precise engineering of active frequency mixing operations [2–4], which are difficult to scale. To address these limitations, we propose a novel solution that employs a pulse shaper and electro-optic phase modulator (EOM) to perform random operations instead of mixing in a prescribed manner. We successfully verify the entanglement and reconstruct the full density matrix of BFCs generated from an on-chip Si_3N_4 microring resonator (MRR) in up to an 8×8 -dimensional two-qudit Hilbert space, the highest dimension to date for frequency bins. More generally, our employed Bayesian statistical model [5, 6] can be tailored to a variety of quantum systems with restricted measurement capabilities, forming an opportunistic tomographic framework that utilizes all available data in an optimal way.

Encoding d levels of quantum information on single photons, known as photonic qudits [7], offers crucial advantages for quantum communication and networking applications [8], such as higher information capacities [9], increased noise tolerance [10, 11], and stronger violations of Bell’s inequalities [12]. Generation and manipulation of photonic qudits have been explored in many degrees of freedom, including path [13, 14], orbital angular momentum [15, 16], frequency bins [2, 3, 17], and time bins [18, 19]. Integrated photonics plays a pivotal role in scaling the complexity of quantum states [20, 21] and quantum operations [22], and the frequency degree of freedom is particularly attractive as on-chip BFCs can produce a large number of spectrally entangled bins in a compact fashion.

Joint spectral intensity (JSI) measurements are commonly used to characterize BFCs, but such measurements are insensitive to phase coherence (and hence entanglement) across frequency-bin pairs. Thus, reconstruction of BFC density matrices has been realized through active mixing of frequency bins [2–4], such that measurements in multiple bases can be realized. In one method, by properly setting the amplitude and phase on a pulse shaper and the modulation voltage on a subsequent EOM, one can filter out overlapping sidebands to perform projective frequency-bin measurements [2, 3]. Alternatively, a quantum frequency processor [23] can be used to synthesize full quantum gates for tomography [4]. Nevertheless, both methods face roadblocks en route to higher dimensions: aggressive amplitude filtering of the input state is inevitable in the first approach, while the number of elements required for arbitrary frequency qudit operations [24] limits the maximum dimensionality possible with current technology.

Accordingly, these existing methods are ill-suited to single-frequency electro-optic modulation: for them, the infinite Fourier series of Bessel functions produced by sinewave electro-optic modulation—a far cry from standard quantum bases—present a challenge to be overcome. In this work, however, we instead leverage the complex mixing behavior of an EOM to our advantage as built-in randomized measurements for BFC characterization. By applying Bayesian tomographic techniques [5, 6], we obtain complete state estimates for any dataset, including uncertainties commensurate with the data gathered, thus bolstering all results obtained from our novel measurement technique with a principled foundation. Importantly, these Bayesian features extend beyond the specific nuances of frequency-bin encoding to any quantum system, offering the promise of meaningful inference irrespective of whatever experimental constraints may have limited the measurements performed.

Figure 1(a) illustrates the experimental setup and concept behind our proposed scheme. The states of interest are BFCs with mode spacing $\Delta\omega/2\pi \sim 40$ GHz and dimension d in both signal and idler photons. The first test source is prepared by pumping a periodically poled lithium niobate (PPLN) waveguide with a continuous-wave laser operating at ~ 780 nm, followed by filtering

* These authors contributed equally to this work.

† lukensjm@ornl.gov

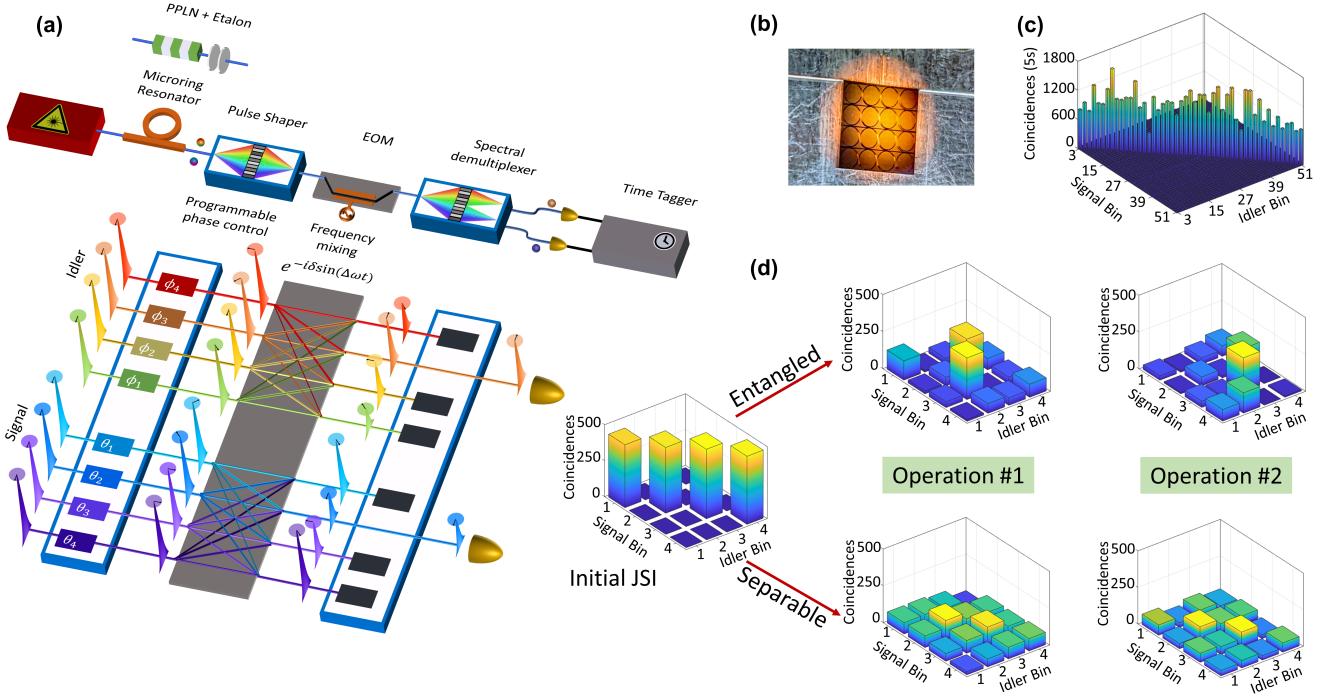


FIG. 1. (a) Experimental setup and conceptual illustration of the proposed method. A random operation is uniquely determined by the d random spectral phases θ_m (ϕ_m) on the signal (idler) bins and the EOM modulation index δ . (b) Microscope image of the MRR chip. The device in the last column of the second row is used experimentally. (c) JSI of the output generated by the MRR. (d) Examples of JSI measurements simulated for the cases of both an entangled state and a classically correlated separable state for two different operations with $\delta = 1.5$ rad (left) and $\delta = 2$ rad (right) and randomly chosen spectral phases.

the broadband spontaneous parametric down-conversion spectrum with a Fabry-Perot etalon [4, 23]. The second source exploits spontaneous four-wave mixing in an on-chip Si_3N_4 MRR [25], where we pump the ring with a tunable continuous-wave laser operating in the optical C-band at one of the ring resonances [2, 3]. Ideal maximally entangled states are of the form $|\Psi_d\rangle = \frac{1}{\sqrt{d}} \sum_{m=1}^d e^{i\alpha_m} |m, m\rangle$, where $|m, m\rangle$ represents the photon pair which is centered at frequency $\omega_0 \pm (m + B)\Delta\omega$ for the signal (idler); α_m is the phase of each pair. The integer B here denotes the number of signal (idler) bins at the center of the biphoton spectrum that are blocked by bandstop filters [omitted in Fig. 1(a)]. Details regarding BFC state preparation can be found in the Methods.

The generated state is then directed to a pulse shaper and an EOM for the implementation of R_{tot} randomly chosen operations. For each operation, we apply a set of d random spectral phases θ_m (ϕ_m) onto the signal (idler) bins with the pulse shaper. The spectral phases (θ_m, ϕ_m) are uniformly sampled between 0 and 2π . The EOM is driven by a sinusoidal voltage with amplitude δ and frequency equal to the mode spacing $\Delta\omega$, imposing a temporal phase $\exp[-i\delta \sin \Delta\omega t]$ onto each photon—equivalently introducing coupling between distinct frequency bins with weights given by Bessel functions of the first kind $J_n(\delta)$. The strength of the imposed phase modulation δ is selected from a set of R_{tot} values equispaced

between $\delta \in [0, \delta_{\text{max}}]$, with δ_{max} set by the maximum radio-frequency (RF) power attainable at $\Delta\omega$. Among these R_{tot} operations, we designate $\delta = 0$ for the first measurement (i.e., the EOM turned off, making it a conventional JSI), while the modulation indices for the remaining $R_{\text{tot}} - 1$ operations are chosen in a random order from the equispaced set, without replacement.

The photons are then passed to a wavelength-selective switch and we scan the filters to collect coincidences over the original $d \times d$ frequency mode grid for each implemented operation, omitting any photons scattered outside of this computational space. Figure 1(d) shows examples of such measurements, simulated using multinomial statistics, corresponding to two different random operations for the cases of a classically correlated separable state ($\rho_{\text{sep}} = \frac{1}{d} \sum_{m=1}^d |m, m\rangle \langle m, m|$) and a maximally entangled quantum state ($\rho_{\text{ent}} = \frac{1}{d} \sum_{m,n=1}^d |m, m\rangle \langle n, n|$) for $d = 4$. In the absence of modulation ($\delta = 0$), their JSIs are identical, yet when the EOM is turned on, the frequency correlations vary strongly. For example, in the extreme case of complete incoherence between energy-matched bins (ρ_{sep} here), applying random phases on the initial pulse shaper has no impact on the measured output. Such differences imply that a collection of these measurements can be exploited to infer the full density matrix.

Given our knowledge of the quantum operations ap-

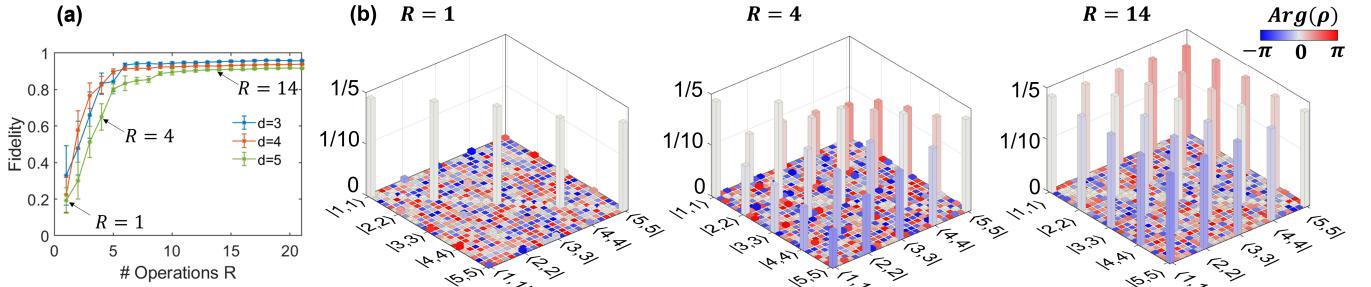


FIG. 2. Evolution of the estimated state for the PPLN BFC. (a) Fidelity with respect to a maximally entangled state as a function of random operations used in Bayesian inference for $d \in \{3, 4, 5\}$. (b) Retrieved density matrices for $d = 5$ when the first $R \in \{1, 4, 14\}$ random operations are considered; their corresponding fidelities are marked by arrows in (a). Shading indicates the phase according to the colorbar scheme shown (here and in the following figures).

plied and a set of $R d \times d$ coincidence results, we then employ Bayesian tomographic techniques [5, 6] to estimate the input quantum state. Conceptually simple—though numerically challenging—Bayesian QST assigns a posterior probability distribution to all unknown parameters \mathbf{x} , given a set of observations \mathcal{D} , according to Bayes' theorem, $\Pr(\mathbf{x}|\mathcal{D}) = \Pr(\mathcal{D}|\mathbf{x})\Pr(\mathbf{x})/\Pr(\mathcal{D})$, which incorporates both a physical model through $\Pr(\mathcal{D}|\mathbf{x})$ and any prior information in $\Pr(\mathbf{x})$. Significantly, the estimator formed by averaging any quantity of interest over the posterior $\Pr(\mathbf{x}|\mathcal{D})$ is guaranteed to offer the lowest squared error on average [26], making the Bayesian mean provably optimal for any number of measurements. Using the model and algorithm described in the Methods, we obtain estimates of the density matrix ρ , fidelity \mathcal{F}_d , and logarithmic negativity E_d ; $E_d > 0$ is a sufficient condition for nonseparability, and E_d upper bounds distillable entanglement [27, 28].

For PPLN experiments, we consider BFCs of qudit dimension $d \in \{3, 4, 5\}$. We implement a total of $R_{\text{tot}} = 21$ randomly chosen operations for each dimension, with a maximum modulation index $\delta_{\text{max}} = 2.5$ rad. We compute the fidelity of the retrieved density matrices at various stages of Bayesian estimation with respect to the ideal state $|\Psi_d\rangle$ with $\alpha_m = \beta_2 L \Delta \omega^2 (m + B)^2$, corresponding to dispersion accumulated over $L = 20$ m of single-mode fiber (approximate length of fiber between the PPLN source and EOM). Figure 2(a) shows the evolution of the Bayesian-estimated fidelity after $R \leq R_{\text{tot}}$ random operations are performed. (For these and all results in this paper, no subtraction of accidentals is performed.) In Fig. 2(b), we plot the mean density matrices retrieved from Bayesian analysis for $d = 5$, at specific numbers of random operations. For the case of $R = 1$, the density matrix resembles a separable state with small off-diagonal elements. As R increases, the off-diagonal elements rise as the phase coherence in the BFC is revealed by operations involving frequency mixing; the fidelity with respect to the ideal state increases accordingly, converging at $R \approx 10$ operations for all three dimensions (see Methods for discussion of possible explanations of

this convergence behavior).

In Fig. 3, we plot both the ideal and the final estimated density matrices ($R = R_{\text{tot}} = 21$). Both absolute values and phases align well with theory—the only discrepancy being the inconsequential phase values of the near-zero off-diagonal elements. From the Bayesian results, we report fidelities \mathcal{F}_d of $\mathcal{F}_3 = (95.8 \pm 0.4)\%$, $\mathcal{F}_4 = (94.0 \pm 0.4)\%$, and $\mathcal{F}_5 = (91.7 \pm 0.4)\%$, which are in the neighborhood of the theoretically predicted $\mathcal{F}_3 = 97.1\%$, $\mathcal{F}_4 = 96.0\%$, and $\mathcal{F}_5 = 94.9\%$ for a white noise model $\rho_\lambda = \lambda |\Psi_d\rangle \langle \Psi_d| + \frac{1-\lambda}{d^2} I_{d^2}$, with λ chosen such that the coincidences-to-accidentals ratio (CAR) matches the experimentally measured value of 90 (see Methods). The Bayesian-estimated log-negativities are $E_3 = 1.523 \pm 0.006$ ebits, $E_4 = 1.911 \pm 0.006$ ebits, and $E_5 = 2.198 \pm 0.007$ ebits, comparable with those of $|\Psi_d\rangle$ ($\log_2 d$) given by 1.58, 2, and 2.32 ebits, respectively.

We then adopt the same methodology for BFCs generated using an on-chip Si_3N_4 MRR [Fig. 1(b)][25]. The measured JSI is shown in Fig. 1(c), with coincidences recorded over 49 signal-idler bin pairs, eight of which (bins 23–30) are selected for testing. The estimated on-chip pair generation rate varies between $\sim 1.3 \times 10^6$ and $\sim 2.2 \times 10^6$ s $^{-1}$ per frequency-bin pair, and the CAR—here defined as the ratio of a given diagonal element to the average of all off-diagonal elements—lies in the interval [17, 30] for the eight bins we test. We perform tomography for qudit dimension $d \in \{2, 3, \dots, 8\}$ with $R_{\text{tot}} = 30$ operations for each d and $\delta_{\text{max}} = 3.4$ rad (increased from 2.5 rad in the PPLN case due to reoptimization of the RF setup for lower loss). We also apply additional spectral phases that compensate for the residual biphoton phase (see Methods), and thus we compute the fidelities of retrieved density matrices with respect to the ideal state $|\Psi_d\rangle$ with $\alpha_m = 0$.

In Fig. 4(a), we plot the final estimated density matrices. The elements indexed by $|m, m\rangle \langle n, n|$ have strong nonzero amplitudes, agreeing well with theory for ideal entangled states. The background corresponding to energy-mismatched bins (gray baseline on the diagonal) is consistent with a white noise model and real-valued ow-

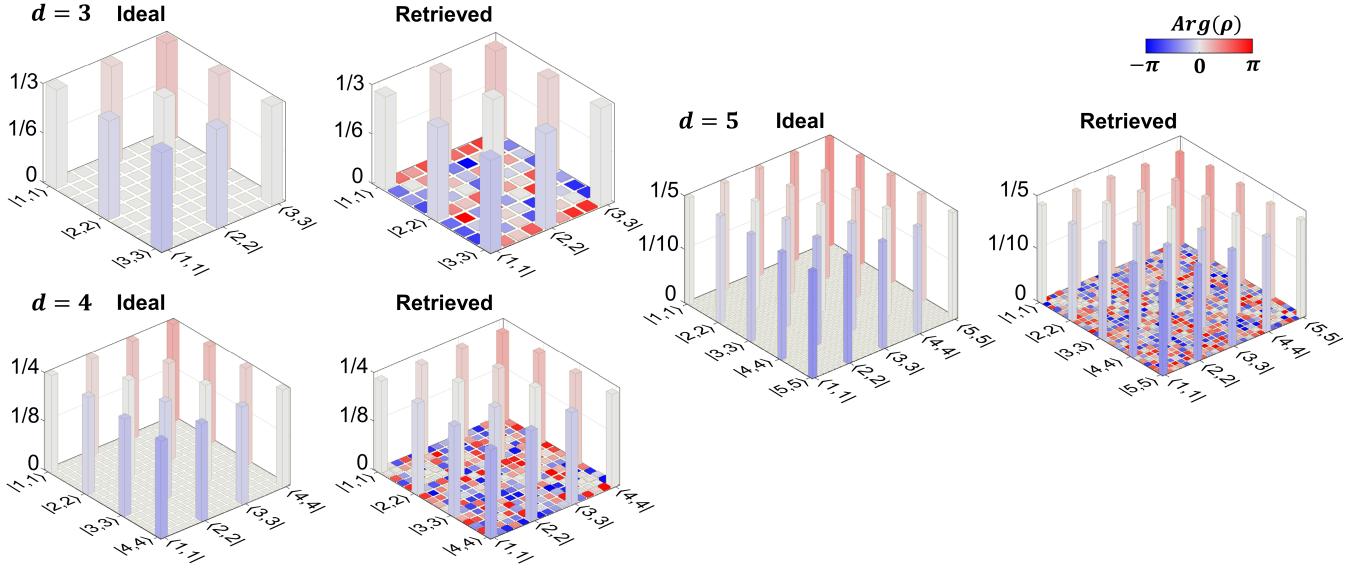


FIG. 3. Final estimated states for the PPLN BFC and $R = 21$ measurements. Ideal and retrieved density matrices for $d \in \{3, 4, 5\}$. Ideal states are d -dimensional Bell states with additional spectral phase equivalent to 20 m of single-mode fiber.

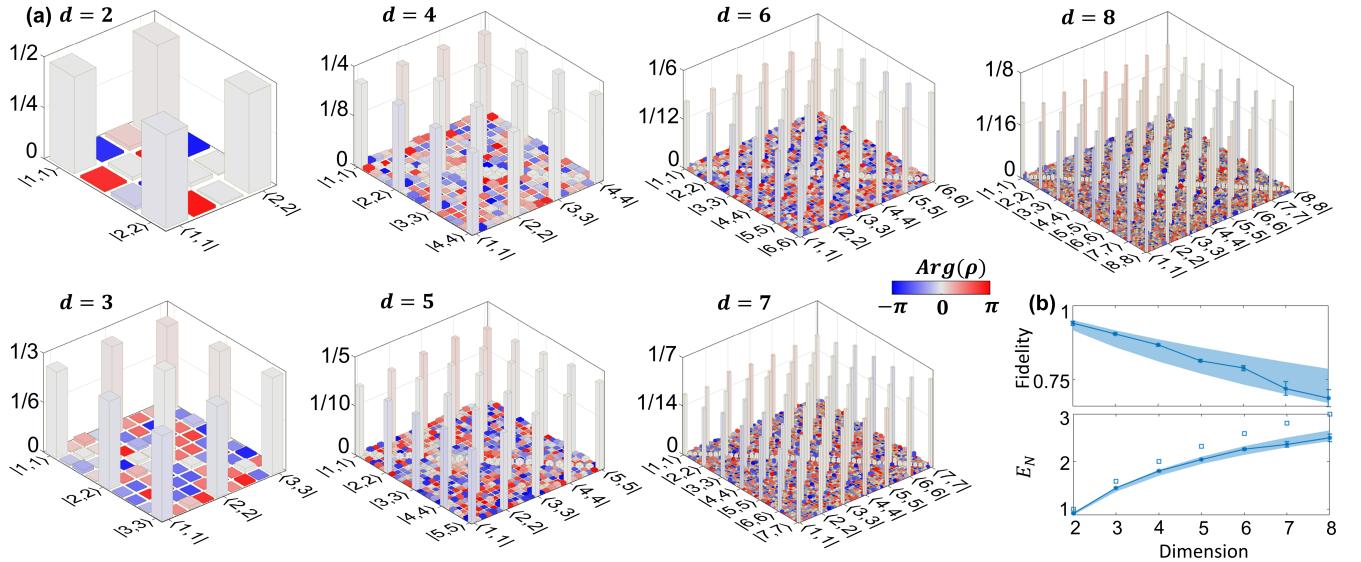


FIG. 4. Tomography results for MRR BFCs and $R = 30$ measurements. (a) Retrieved mean density matrices for dimensions up to $d = 8$. Ideal states $|\Psi_d\rangle$ (not plotted) are d -dimensional maximally entangled states with uniform phase. (b) Fidelities and log-negativities of retrieved density matrices (solid line). Shaded region corresponds to theoretically computed values for ρ_λ (a mixture of a maximally entangled state and white noise) lower- and upper-bounded by CARs of 17 and 30, respectively. Square markers correspond to log-negativities of maximally entangled states, given by $\log_2 d$.

ing to hermiticity. Figure 4(b) plots the fidelities with respect to $|\Psi_d\rangle$ and the corresponding log-negativities E_d , lying comfortably within the range predicted for our noise model using experimentally observed CARs (shaded region). Significantly, our $d = 8$ result of $E_8 = 2.50 \pm 0.08$ ebits can only be achieved by two-qudit states with $d \geq 6$, indicating the genuine high-dimensional nature of the observed entanglement. These results showcase

the highest dimension of a fully reconstructed density matrix—Hilbert space dimension of 64—in experimental frequency-bin encoding.

Technologically speaking, our approach aligns closely with recent ideas presented in EOM-based frequency-bin quantum random walks [29], where here we precede the walk with random spectral phases and consider varying circuit depths. Yet beyond the confines of frequency-

bin encoding, our statistical treatment hints at the much wider value of Bayesian models in quantum information. Neither the choice of measurement settings nor number of datapoints has any bearing on the legitimacy of Bayesian tomography [5, 6]; hence, Bayesian estimation will return a reasonable result for any dataset, with automatic uncertainty quantification indicating the confidence warranted from the data. This feature imparts Bayesian inference with unique flexibility compared to other advanced random measurement approaches, in that it does not assume, e.g., low-rank states [30] or rely on unitary operations drawn from specific distributions [31, 32]. Therefore, any quantum system for which an appropriate physical model can be constructed is ripe to potentially benefit from Bayesian models like the one presented here. Although not required conceptually, well-chosen measurements are *practically* valuable for obtaining final estimates accompanied by low uncertainties. Our experimental results show through example that the datasets from random modulation are more than sufficient to converge from an initially uniform (Bures) prior to density matrices with small error bars and in good agreement with the expected ground truth. Indeed, arguments from a simple theoretical model suggest that random EOM measurements with $\delta_{\max} \sim \mathcal{O}(d)$ efficiently cover the entire Hilbert space of a d -dimensional quantum system (see Methods for details), so that our measurement approach offers a straightforward path for high-dimensional frequency-bin characterization and should open new opportunities for BFCs in quantum information processing.

ACKNOWLEDGMENTS

We thank AdvR for loaning the PPLN ridge waveguide; P. Imany, N. B. Lingaraju, and A. J. Moore for valuable discussions; A. A. N. Ovi for laboratory help; and B. T. Kirby for introducing us to the Bures distribution. This work was performed in part at Oak Ridge National Laboratory, operated by UT-Battelle for the U.S. Department of Energy under contract no. DE-AC05-00OR22725. Funding was provided by the U.S. Department of Energy, Office of Advanced Scientific Computing Research, Early Career Research Program (Field Work Proposal ERKJ353), the National Science Foundation (1839191-ECCS, 2034019-ECCS), the Air Force Office of Scientific Research (Award no. FA9550-19-1-0250), and the Swiss National Science Foundation under Grant no. 176563 (BRIDGE). K.V.M. acknowledges support from the QISE-NET fellowship program of the National Science Foundation (DMR-1747426). M.S.A. acknowledges support from the College of Engineering Research Center at King Saud University. The Si_3N_4 samples were fabricated in the EPFL Center of MicroNanoTechnology (CMi).

AUTHOR CONTRIBUTIONS

H.H.L. and K.V.M. contributed equally to this work. H.H.L. initiated the concepts, led the experiments, and contributed to the theoretical analysis. K.V.M. performed experiments with MRR sources and led writing of the paper. S.S., M.S.A, and D.E.L assisted the experiments. R.S.B. contributed to the theoretical analysis. J.L. designed and fabricated the integrated device, supervised by T.J.K. A.M.W. supervised the work at Purdue and assisted with system design and analysis. J.M.L. managed the theoretical analysis and numerical simulation, and supervised the project. All authors reviewed the results and contributed to the manuscript.

I. METHODS

A. PPLN source setup

Our test source is a 2.1 cm-long fiber-pigtailed PPLN ridge waveguide (AdvR), possessing an internal efficiency of 150 %/W for second-harmonic generation and fiber-coupling efficiency of $\sim 75\%$ per facet. We couple a continuous-wave laser (Toptica) operated at ~ 5 mW and ~ 780 nm into the PPLN waveguide, temperature-controlled to 51.6 °C for SPDC under type-0 phase matching. Broadband spectrally entangled photon pairs spanning >5 THz are generated and subsequently filtered by a tunable fiber-pigtailed Fabry-Perot etalon (Luna Innovations) with 20 GHz mode spacing and a full-width at half-maximum linewidth of 1 GHz. We carefully tune the etalon’s temperature to align its transmission peaks with the generated entangled photons, i.e., to maximize the coincidences between the symmetric, spectrally filtered mode pairs.

To minimize crosstalk in our 20 GHz-resolution demultiplexer (Finisar Waveshaper 4000S/X), we utilize the first pulse shaper (Finisar Waveshaper 1000S) to perform amplitude filtering (on top of the phase masks programmed for the QST) and block every other frequency bin, resulting in a 40 GHz-spaced BFC with a measured CAR ≈ 90 . A 300 GHz bandstop filter (corresponding to the case of $B = 3$) is programmed on the same pulse shaper in the middle of the BFC spectrum, which allows us to apply strong modulation on the EOM without the possibility of the signal photon jumping over into the idler’s modes, and vice versa. For coincidence measurements, we use a window of 128 ps and integration times of 20 s for all dimensions.

B. MRR source set up

The Si_3N_4 MRR used in our experiment is fabricated using the optimized photonic Damascene reflow process [33] with a cross-section of $2 \mu\text{m} \times 0.95 \mu\text{m}$. Such a process has enabled ultralow loss waveguides that have paved

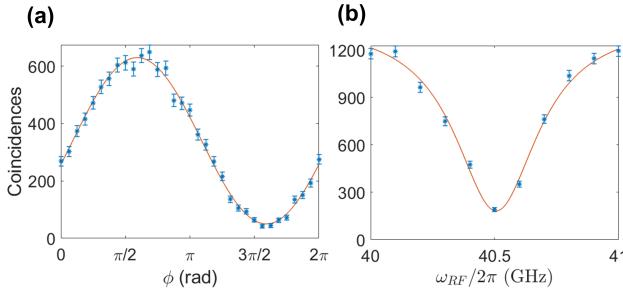


FIG. 5. Characterization of residual BFC spectral phase and FSR of MRR. Two adjacent signal-idler bin pairs are selected and equally mixed using an EOM. Biphoton time-correlation functions of one energy-matched pair are then measured. (a) Coincidences at $\tau = 0$ as joint spectral phase ϕ is swept. (b) Coincidences integrated over τ as modulation frequency ω_{RF} is swept at a spectral phase that minimizes coincidences at $\tau = 0$ (details in text).

the way for dissipative Kerr solitons with free spectral ranges (FSRs) as low as 10 GHz [25]. In the current work, the radius of the ring is 561 μm , corresponding to an FSR of 40.5 GHz—within the range of commercial EOMs—allowing us to drive the EOM at a frequency equal to the FSR for random operations. The gap between the ring and the bus waveguide is 0.3 μm , resulting in strong overcoupling with an intrinsic Q -factor of $\sim 10^7$ and a loaded Q -factor of $\sim 10^6$. We pump the ring using a tunable continuous-wave laser (New Focus) operating in the optical C-band, with an on-chip power of ~ 10 mW—well below the classical comb threshold of ~ 80 mW. The pump is amplified using an erbium doped fiber amplifier (EDFA), which is subsequently filtered using a set of two 100 GHz-wide dense wavelength division multiplexing (DWDM) filters to suppress amplified spontaneous emission from the EDFA. Lensed fibers are used to couple the pump into the MRR, which is positioned on a temperature-controlled stage maintained at ~ 23 $^{\circ}\text{C}$. The fiber-to-fiber coupling loss of the ring is around ~ 4 dB. Such a low loss was realized using engineered inverse waveguide tapers [34]. When the pump is tuned into the ring resonance (~ 1550.5 nm) and operated below threshold for classical comb generation, it gives rise to BFC states. Since the pump and newly generated biphotons are in the same wavelength band, it is essential to suppress the residual pump after the ring to reduce accidentals in coincidence measurements. We use a set of three 200 GHz-wide DWDM filters, which when combined with bandstop filters in the pulse shaper and demultiplexer gives a net pump suppression of ~ 100 dB. We also tap a portion of the pump power to track its wavelength using a wavelength meter, setting up a computer-based feedback loop to ensure that the pump is operated at the intended resonance. For coincidence measurements, we use a window of 2048 ps (roughly equivalent to the inverse resonance linewidth) and integration time of 5 s for $d \in \{2, 3, \dots, 6\}$. For $d \in \{7, 8\}$, we reduce the integra-

tion time to 3 s in the interest of minimizing the total experimental duration.

It is necessary to experimentally characterize the FSR of the MRR to determine the precise modulation frequency needed. In addition to this, we also characterize the residual spectral phase accumulated by the biphotons, likely due to fiber dispersion, to precompensate for it. For characterization of these quantities, the experimental setup remains the same as shown in Fig. 1(a). We select two adjacent signal-idler bin pairs in the first pulse shaper and drive the EOM at a frequency ω_{RF} with amplitude $\delta = 1.43$ such that $|J_0(\delta)| = |J_1(\delta)|$ for equal mixing. We then pass one energy-matched pair of signal-idler frequency bins through the demultiplexer, now consisting of equal contributions from the adjacent bin due to phase modulation, and measure the biphoton time-correlation function. The theoretical expression for the coincidence rate, assuming both bins have identical Lorentzian lineshape and equal probability amplitude, is given by [35]

$$R(\tau) \propto e^{-\gamma|\tau|} \{1 - \cos [\phi + \phi_0 - (\omega_{FSR} - \omega_{RF})\tau]\} \quad (1)$$

where γ represents the Lorentzian linewidth, τ the signal-idler delay, ϕ the joint spectral phase between the bins applied by the pulse shaper, ϕ_0 the residual biphoton phase, and ω_{FSR} the FSR. In Fig. 5(a), we plot the coincidences at $\tau = 0$ as a function of ϕ for a set of two adjoining signal-idler bin pairs, where $\tau = 0$ is defined as the peak of the histogram in the unmodulated case. The offset of the experimentally obtained cosine function from the origin can be used to deduce the residual spectral phase between the bins. This process is repeated for all sets of adjoining signal-idler bin pairs used in our experiment to deduce the residual spectral phase of the entire biphoton state. For finding ω_{FSR} , we again work with a set of two adjoining signal-idler bin pairs and operate at a spectral phase that minimizes the coincidences at $\tau = 0$. In Fig. 5(b) we plot the coincidences integrated over τ as we sweep the applied RF frequency ω_{RF} . When $\omega_{RF} = \omega_{FSR}$, the coincidences are minimized, so that Fig. 5(b) implies $\omega_{FSR}/2\pi = 40.5$ GHz. The solid lines in Fig. 5 are theoretical estimates, scaled and vertically offset to match the data points via least squares. The vertical offset and the scaling account for both the nonzero accidentals and unequal bin probability amplitudes. The width of the trace in Fig. 5(b) can be expressed as 2γ from Eq. (1), from which $\gamma/2\pi \approx 200$ MHz is inferred. These obtained resonator linewidth and FSR values are consistent with linear spectroscopy measurements.

C. Bayesian inference model

The ultimate goal of the inference process is to estimate the full BFC state, which can be represented as a $d^2 \times d^2$ density matrix ρ . The frequency bins of the signal (idler) qudit possess annihilation operators $\hat{a}_k^{(S)} (\hat{a}_l^{(I)})$

where $k, l \in \{1, \dots, d\}$ correspond to center frequencies of $\omega_k = \omega_0 + (k + B)\Delta\omega$ and $\omega_l = \omega_0 - (l + B)\Delta\omega$. The density matrix of interest can be formally written as

$$\rho = \sum_{k,l,k',l'=1}^d \rho_{(kl)(k'l')} [\hat{a}_k^{(S)}]^\dagger [\hat{a}_l^{(I)}]^\dagger |\text{vac}\rangle \langle \text{vac}| \hat{a}_{k'}^{(S)} \hat{a}_{l'}^{(I)}, \quad (2)$$

where $|\text{vac}\rangle$ is the vacuum state.

We use the index s to denote a specific measurement, which consists of a particular EOM and pulse shaper setting $r(s) \in \{1, \dots, R_{\text{tot}}\}$, the signal frequency bin measured $m(s) \in \{1, \dots, d\}$, and the idler bin measured $n(s) \in \{1, \dots, d\}$. For notational convenience, the explicit s -dependence is suppressed in many of the formulas below, but remains implied for (r, m, n) . For a given r , the frequency bins undergo unitary transformations into

the output modes via

$$\begin{aligned} \hat{b}_m^{(S)} &= \sum_{k=1}^d V_{mk}^{(r)} \hat{a}_k^{(S)} \\ \hat{b}_n^{(I)} &= \sum_{l=1}^d W_{nl}^{(r)} \hat{a}_l^{(I)}. \end{aligned} \quad (3)$$

The unitary operations $V^{(r)}$ and $W^{(r)}$ consist of line-by-line phase shifts on each of the input modes, followed by sinewave electro-optic modulation. The index r defines the phase shifts and modulation index for a specific setting: $\theta_n^{(r)}, \phi_n^{(r)}, \delta^{(r)}$. We can therefore write

$$\begin{aligned} V_{mk}^{(r)} &= J_{m-k}(\delta^{(r)}) e^{i\theta_k^{(r)}} \\ W_{nl}^{(r)} &= J_{l-n}(\delta^{(r)}) e^{i\phi_l^{(r)}}, \end{aligned} \quad (4)$$

where $J_n(\cdot)$ is the Bessel function of the first kind, and for definiteness we have assumed a modulation function of the form $\exp[-i\delta^{(r)} \sin \Delta\omega t]$. The only major difference between the signal and idler equations is the input/output index reversal in the Bessel function order, which results from our definition of signal frequencies that increase with index and idler frequencies that decrease with index.

After these operations, we look for coincidences between bins m and n , the probability of which can be computed for the input density matrix as

$$\begin{aligned} p_s &= \langle m, n | \rho | m, n \rangle = \langle \text{vac} | \hat{b}_m^{(S)} \hat{b}_n^{(I)} \rho [\hat{b}_m^{(S)}]^\dagger [\hat{b}_n^{(I)}]^\dagger | \text{vac} \rangle \\ &= \sum_{k,l,k',l'=1}^d \rho_{(kl)(k'l')} V_{mk}^{(r)} W_{nl}^{(r)} [V_{mk'}^{(r)}]^\dagger [W_{nl'}^{(r)}]^\dagger \\ &= \sum_{k,l,k',l'=1}^d \rho_{(kl)(k'l')} J_{m-k}(\delta^{(r)}) J_{l-n}(\delta^{(r)}) J_{m-k'}(\delta^{(r)}) J_{l'-n}(\delta^{(r)}) e^{i[\theta_k^{(r)} + \phi_l^{(r)} - \theta_{k'}^{(r)} - \phi_{l'}^{(r)}]}, \end{aligned} \quad (5)$$

using Eqs. (2–4) to simplify. This expression provides a linear mapping from the density matrix elements $\rho_{(kl)(k'l')}$ to each output probability p_s .

We then need to relate these probabilities to the observed coincidence counts N_s through an appropriate likelihood function. In typical quantum tomographic contexts where the full Hilbert space is detected at each setting, a multinomial model is the most conceptually straightforward [5]; in our case this would consist of products of the factors $p_s^{N_s}$. However, such a model does not readily apply to situations in which some of the outcomes are unmonitored. In our particular experiment, many bins outside of the original d^2 -dimensional computational space can be populated, owing to the nonzero Bessel function weights in Eq. (5). As a rule of thumb, the success probability (i.e., the possibility of a single photon

staying in the original d -mode computational space) is roughly $1/2$ when the modulation depth $\delta \sim \mathcal{O}(d)$ [36]. Rather than attempting to measure all output mode combinations—which in principle involves an infinite-dimensional Hilbert space and in practice means the addition of many measurements with few counts—we focus only on the central $d \times d$ space here.

On the model side, we can account for unobserved outcomes by introducing an additional flux parameter K , defined as the average number of total coincidences that *would* be measured if all (m, n) combinations were tested. Since the EOM and pulse shaper operations are unitary when considered over all modes—apart from an overall insertion loss that does not vary with setting r —this scale factor is fixed for all measurement settings. It also automatically accounts for efficiency; explicitly, it

can be written as $K = \eta_S \eta_I \Phi \Delta T$, where η_S (η_I) is the total system efficiency from generation through detection for the signal (idler), Φ the photon pair generation flux, and ΔT the integration time. Although these quantities themselves could be inferred by considering both singles counts as well as coincidences [4, 37, 38], they are not of interest in the present investigation, and so through K we are able to reduce to their combined effect only.

Thus, the mean number of coincidences for a specific setting s becomes Kp_s , which we can model with a Poissonian distribution, the product of which produces the full likelihood $L_{\mathcal{D}}(\rho, K) \propto \Pr(\mathcal{D}|\rho, K)$

$$L_{\mathcal{D}}(\rho, K) = \prod_{s=1}^{Rd^2} e^{-Kp_s} (Kp_s)^{N_s}, \quad (6)$$

where $\mathcal{D} = \{N_1, \dots, N_{Rd^2}\}$ denotes the set of measured coincidences for all Rd^2 settings. This likelihood matches the form we adopted previously in the construction of JSIs (not the full quantum state) from random measurements [39]. And not only does it readily handle probabilities that do not encompass the full Hilbert space, but it also more accurately reflects the physical situation of high-dimensional measurements with single-outcome detectors. For example, our use of raster scanning with two single-photon detectors means that, for any given configuration, we measure events only for a specific (m, n) frequency-bin pair. Although one can pool the results for all pairs (m, n) and view them synthetically as resulting from a single ΔT integration time of true d^2 -outcome measurements—the situation assumed by a multinomial distribution—this ultimately does not align with the actual measurement procedure. Accordingly, the Poissonian likelihood provides both a practical and conceptually satisfying model for our tomographic scenario.

With the likelihood defining the relationship between a given density matrix and the observed data, Bayesian inference next requires specification of a suitable prior distribution for ρ and K . In the case of QST, uniform priors are generally preferred, as these apply appreciable weights to all possible states and thereby minimally bias the final results. Although a variety of reasonable uninformative priors can be posited for density matrices, we select the Bures distribution, which enjoys a unique position as the single monotone metric which reduces to both the Fisher and Fubini–Study metrics in the classical and pure state limits, respectively [40], and in this sense can claim preference as a “definitive” uniform prior for Bayesian inference [41].

To work with the Bures ensemble, it is convenient to express any $d^2 \times d^2$ density matrix ρ in the computational

basis as

$$\rho = \frac{(I_{d^2} + U)GG^\dagger(I_{d^2} + U^\dagger)}{\text{Tr}[(I_{d^2} + U)GG^\dagger(I_{d^2} + U^\dagger)]}, \quad (7)$$

where the $d^2 \times d^2$ matrices I_{d^2} , U , and G are the identity, a unitary matrix, and a general complex matrix, respectively. This expression automatically satisfies all physicality conditions (unit-trace, hermiticity, and positive semidefiniteness); by sampling U from the Haar distribution and G from the Ginibre ensemble, the ρ thus formed represents a single draw from the Bures distribution [41]. We can in turn parameterize ρ by the complex vector $\mathbf{y} = (y_1, \dots, y_{2d^4})$, with each y_k observing a complex standard normal distribution $y_k \stackrel{\text{i.i.d.}}{\sim} \mathcal{CN}(0, 1)$. d^4 of the components comprise the $d^2 \times d^2$ elements of the Ginibre matrix G directly, while the remaining d^4 parameters form a second Ginibre matrix which is converted to the unitary U through the Mezzadri algorithm [42], thereby ensuring Haar randomness.

In addition to the parameters forming ρ , the scale factor K must also be suitably parameterized. Following [39], we find it convenient to write $K = K_0(1 + \sigma z)$, where K_0 and σ are hyperparameters defined separate of the inference process, and z is taken to follow a standard normal distribution $\mathcal{N}(0, 1)$, leading to a normal prior on K of mean K_0 and standard deviation $K_0\sigma$. We take $\sigma = 0.1$ and K_0 equal to the sum of the counts in all d^2 bins for the first JSI measurement ($r = 1$), where the absence of modulation ensures that all initial photon flux remains in measured bins, i.e., $K_0 = \sum_{s=1}^{d^2} N_s$. This provides an effectively uniform prior, since a fractional deviation of 0.1 is much larger than the maximum amount of fractional uncertainty $1/\sqrt{K_0} \approx 0.02$ expected from statistical noise at our total count numbers; the use of a normal distribution simplifies the sampling process.

The total parameter set can therefore be expressed as the vector $\mathbf{x} = (\mathbf{y}, z) = (y_1, \dots, y_{2d^4}, z)$, with the prior distribution

$$\pi_0(\mathbf{x}) \propto \left(\prod_{k=1}^{2d^4} e^{-\frac{1}{2}|y_k|^2} \right) e^{-\frac{1}{2}z^2}. \quad (8)$$

We note that this parameterization entails a total of $4d^4 + 1$ independent real numbers ($2d^4$ complex parameters for ρ , one real parameter for K)—noticeably higher than the minimum of $d^4 - 1$ required to uniquely describe a density matrix. Nevertheless, this $\rho(\mathbf{y})$ parameterization is to our knowledge the only existing constructive method to produce Bures-distributed states, and is straightforward to implement given its reliance on independent normal parameters only.

Following Bayes’ rule, the posterior distribution becomes

$$\pi(\mathbf{x}) = \frac{1}{\mathcal{Z}} L_{\mathcal{D}}(\mathbf{x}) \pi_0(\mathbf{x}) \propto \left(\prod_{s=1}^{Rd^2} e^{-K(z)p_s(\mathbf{y})} [K(z)p_s(\mathbf{y})]^{N_s} \right) \left(\prod_{k=1}^{2d^4} e^{-\frac{1}{2}|y_k|^2} \right) e^{-\frac{1}{2}z^2}, \quad (9)$$

where \mathcal{Z} is a constant such that $\int d\mathbf{x} \pi(\mathbf{x}) = 1$. We have adopted this notation for Bayes' theorem—rather than the more traditional $\text{Pr}(\mathbf{x}|\mathcal{D}) = \text{Pr}(\mathcal{D}|\mathbf{x})\text{Pr}(\mathbf{x})/\text{Pr}(\mathcal{D})$ —to emphasize the functional dependencies on \mathbf{x} , which are all that must be accounted for in the sampling algorithm below. From $\pi(\mathbf{x})$, the Bayesian mean estimator f_B of any quantity (scalar, vector, or matrix) expressible as a function of \mathbf{x} can be estimated as

$$f_B = \int d\mathbf{x} \pi(\mathbf{x}) f(\mathbf{x}) \approx \frac{1}{S} \sum_{j=1}^S f(\mathbf{x}^{(j)}), \quad (10)$$

where, in lieu of direct integration, S samples $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(S)}\}$ are obtained from the distribution $\pi(\mathbf{x})$ through Markov chain Monte Carlo (MCMC) techniques, as described below.

D. MCMC Sampling

Acquiring the samples necessary for computation of high-dimensional integrals of the form in Eq. (10) forms the primary bottleneck in Bayesian inference. The most common family of solutions to address this challenge fall under the general umbrella of MCMC, in which samples from a well chosen Markov chain are designed to approach the statistics of $\pi(\mathbf{x})$ asymptotically [26, 43]. Recently, we applied a particularly efficient MCMC algorithm—known as preconditioned Crank–Nicolson (pCN) [44]—to the problem of Bayesian QST, finding significant computational improvements over previous implementations [6]. We utilize the same pCN approach here; the only difference from the algorithm in [6] is a simpler acceptance probability depending exclusively on the likelihood ratio, a consequence of the fact that all parameters here have normally distributed priors [44].

For each collection of measurement results considered in the main text, we retain $S = 2^{10}$ samples $\{(\mathbf{y}^{(j)}, z^{(j)})\}$ from an MCMC chain of total length ST , where T is a thinning factor that is successively doubled until convergence is obtained. From these samples, we then compute the Bayesian mean estimate of the density matrix

$$\rho_B = \frac{1}{S} \sum_{j=1}^S \rho(\mathbf{y}^{(j)}), \quad (11)$$

examples of which are plotted in Figs. 2–4. The mean and standard deviation of fidelity with respect to some ideal state $|\Psi_d\rangle$ can also be computed as

$$\overline{\mathcal{F}_d} = \frac{1}{S} \sum_{j=1}^S \langle \Psi_d | \rho(\mathbf{y}^{(j)}) | \Psi_d \rangle \quad (12)$$

and

$$\Delta \mathcal{F}_d = \left(\frac{1}{S} \sum_{j=1}^S [\langle \Psi_d | \rho(\mathbf{y}^{(j)}) | \Psi_d \rangle]^2 - \overline{\mathcal{F}_d}^2 \right)^{\frac{1}{2}}, \quad (13)$$

respectively. For the PPLN states specifically, which are not dispersion-compensated, we take the state for comparison as

$$|\Psi_d\rangle = \frac{1}{\sqrt{d}} \sum_{m=1}^d e^{i\beta_2 L \Delta \omega^2 m^2} |m, m\rangle, \quad (14)$$

with $\beta_2 = 2.06 \times 10^{-2} \text{ ps}^2 \text{ m}^{-1}$ for standard single-mode fiber and $L = 20 \text{ m}$. The phases of the MRR states are precompensated, so for computing fidelities for them we use the uniform-phase version

$$|\Psi_d\rangle = \frac{1}{\sqrt{d}} \sum_{m=1}^d |m, m\rangle. \quad (15)$$

In order to monitor convergence with T , we consider the sequential fidelity, defined for a given $T = 2^n$ as

$$\mathcal{F}_{n,n-1} = \left(\text{Tr} \sqrt{\sqrt{\rho_B^{(n-1)}} \rho_B^{(n)} \sqrt{\rho_B^{(n-1)}}} \right)^2, \quad (16)$$

where we use the notation $\rho_B^{(n)}$ to denote the Bayesian mean estimate [Eq. (10)] for a chain of length ST^n . For a sufficiently large thinning value $\mathcal{F}_{n,n-1}$ will converge to unity and remain for all subsequent n . We note that this metric contains no reference to the ideal state $|\Psi_d\rangle$, but checks for consistency between subsequent Bayesian estimates only.

Figure 6 plots these sequential fidelities as a function of thinning ($n = \log_2 T$) for all BFCs characterized in the main text. In all cases, we consider the full measurement sets, $R = 21$ for the PPLN results and $R = 30$ for the MRR findings, as these cases generally possess the slowest MCMC convergence. As expected, the sequential fidelity converges more rapidly with T for lower d ; nonetheless, all examples ultimately reach $\mathcal{F}_{n,n-1} > 0.99$ at their respective maximum of n , indicating high MCMC convergence for all reported ρ_B matrices. Continuing to even larger values of T would certainly be desirable, particularly for $d = 7$ and $d = 8$, yet we are currently limited by computational power. For example, the $d = 8$ MCMC chain with $T = 2^{14}$ required almost a week to complete; with a total of 16,385 parameters to infer $(4d^4 + 1)$, it is no surprise that we are pushing the limits of our desktop computer. Indeed, to our knowledge the present results

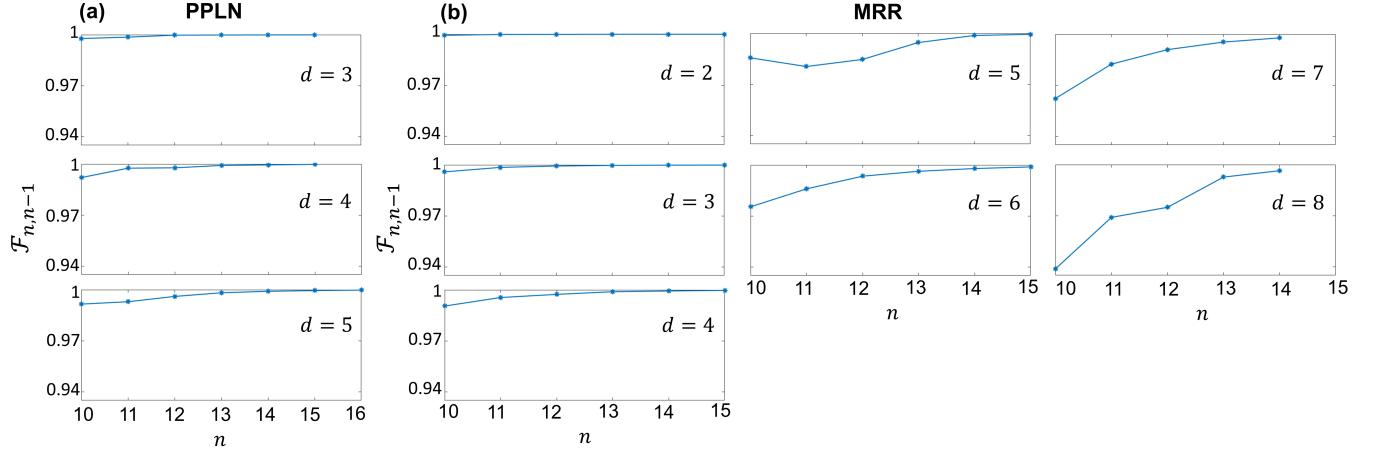


FIG. 6. Sequential fidelity vs. thinning factor $n = \log_2 T$ for all dimensions for BFC states generated from (a) PPLN and (b) MRR.

with $d^2 = 64$ represent a record-high Hilbert-space dimension for complete Bayesian QST with a fully general (mixed state) prior, for any quantum system, simultaneously implying the efficiency of our existing MCMC algorithm as well as the importance of pursuing highly parallelizable MCMC methods [45] to reach even larger dimensions in the future.

E. Theoretical fidelities and ebits

In our quantum system, several sources of noise, such as multipair emission and dark counts, are expected to be uniform across the BFC bins. Accordingly, as a simple model, we theoretically anticipate a ground truth quantum state of the form

$$\rho_\lambda = \lambda |\Psi_d\rangle \langle \Psi_d| + \frac{1-\lambda}{d^2} I_{d^2}, \quad (17)$$

where $|\Psi_d\rangle$ is the ideal maximally entangled state [Eq. (14) or (15)], I_{d^2} is the $d^2 \times d^2$ identity operator, and $\lambda \in [0, 1]$ determines the noise level. In this section, we provide quantities of interest (CAR, fidelity \mathcal{F}_d , and log-negativity E_d) for ρ_λ in order to compare against those found in Bayesian estimation:

$$\begin{aligned} \text{CAR} &= \frac{\max_m \langle m, m | \rho_\lambda | m, m \rangle}{\min_{m,n} \langle m, n | \rho_\lambda | m, n \rangle} = 1 + \frac{d\lambda}{(1-\lambda)} \\ \mathcal{F}_d &= \langle \Psi_d | \rho_\lambda | \Psi_d \rangle = \frac{(d^2-1)\lambda+1}{d^2} \\ E_d &= \begin{cases} 0 & ; \mathcal{F}_d < 1/d \\ \log_2 d\mathcal{F}_d & ; \mathcal{F}_d > 1/d \end{cases} \end{aligned} \quad (18)$$

where the last formula is adapted from [28]. The fidelities and log-negativities reported for ρ_λ in the main text are calculated using these equations. In general, both fidelity and E_d deviate more strongly from their respective ideals as d increases for a fixed CAR. The good agreement in

the main text between the Bayesian-estimated states and this simple white noise model suggests that our understanding of noise processes in the system is well justified.

F. Theoretical analysis of measurement efficiency

In this section we provide a basic theoretical analysis of the effectiveness of our proposed tomographic method. We first consider a highly simplified model that captures the key aspects of how phase modulation and mode mixing enable tomographic reconstruction, concentrating in this case on a single frequency-bin qudit occupying d modes $x \in \{1, \dots, d\}$. While the experiments in the main text examine the joint state of *two* qudits instead, this simpler single-qudit case reveals the basic principles of the tomography method with minimal distractions. The only type of measurement considered is projection onto the photon's output mode, reflecting the condition of frequency-resolved detection; the probability of obtaining bin x is $p_x^{(0)} \equiv \rho_{xx}$, and the qudit assumption ensures that $\rho_{xx} = 0$ for $x < 1$ or $x > d$.

Consider the mode mixing operator S_k which weakly mixes each mode x with modes $x+k$ and $x-k$:

$$S_k |x\rangle = -\epsilon |x-k\rangle + |x\rangle + \epsilon |x+k\rangle. \quad (19)$$

For $k = 1$, this approximates the action of an EOM with small modulation depth $\delta = 2\epsilon$ and modulation frequency equal to the fundamental mode spacing. More generally, for small ϵ , S_k is approximately unitary and describes a weak, translation-invariant mode mixing operation [Fig. 7(a)].

Suppose we apply S_k to the photon and measure its bin. The probability of observing bin x is

$$p_x^{(k)} \equiv \langle x | S_k \rho S_k^\dagger | x \rangle. \quad (20)$$

Using $S_k^\dagger |x\rangle = \epsilon |x-k\rangle + |x\rangle - \epsilon |x+k\rangle$ and the fact that $\rho_{yx} = \rho_{xy}^*$ we obtain

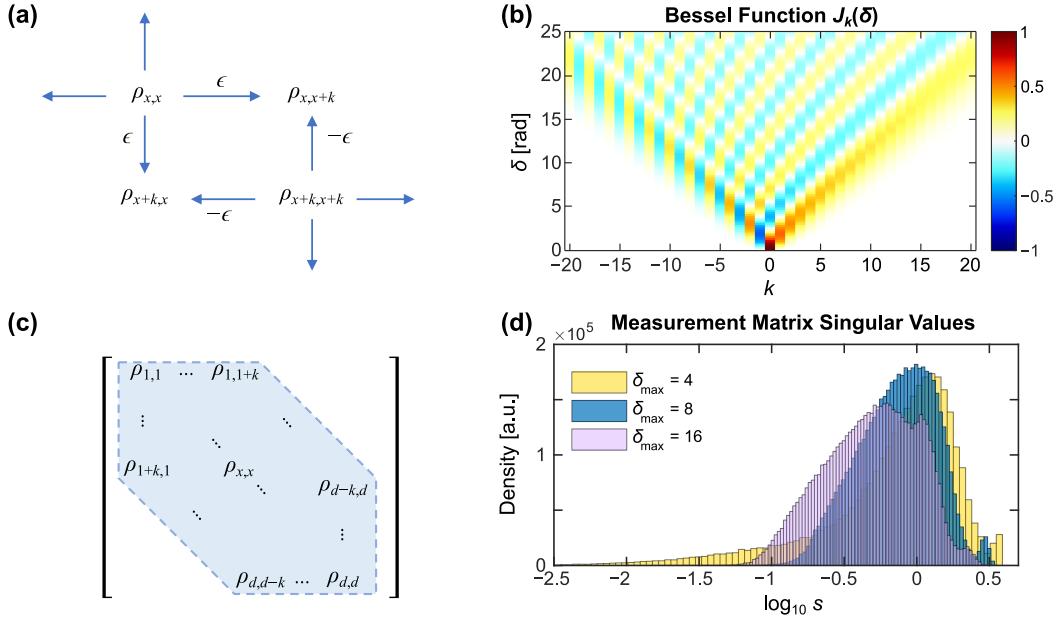


FIG. 7. Theoretical analysis of measurement approach. (a) Effect of the mode-mixing operator S_k [Eq. (19)] on the density operator ρ . (b) Behavior of the Bessel functions, which describe the modulation operators T_δ [Eq. (27)]. (c) The modulation operator T_δ with $\delta \approx k$ probes a k -band of ρ , i.e., the elements at most k above or below the diagonal. (d) Distribution of singular values for random measurement matrices O for $d = 8$, $\delta \in [0, \delta_{\max}]$, and $R = 2d$: (gold) $\delta_{\max} = 4$; (blue) $\delta_{\max} = 8$; (violet) $\delta_{\max} = 16$.

$$\begin{aligned}
 p_x^{(k)} &= (\epsilon \langle x - k | + \langle x | - \epsilon \langle x + k |) \rho (\epsilon |x - k \rangle + |x \rangle - \epsilon |x + k \rangle) \\
 &= \rho_{xx} + \epsilon (\rho_{x-k,x} + \rho_{x,x-k} - \rho_{x+k,x} - \rho_{x,x+k}) + \mathcal{O}(\epsilon^2) \\
 &= \rho_{xx} + 2\epsilon (\Re \rho_{x-k,x} - \Re \rho_{x,x+k}) + \mathcal{O}(\epsilon^2).
 \end{aligned} \tag{21}$$

Now, from the fact that $|\rho_{xy}|^2 \leq \rho_{xx}\rho_{yy}$ and the fact that the photon was initially restricted to modes $\{1, \dots, d\}$ we have that the support of $\rho_{x,x+k}$ is $x \in \{1, \dots, d-k\}$ and the support of $\rho_{x-k,x}$ is $x \in \{k+1, \dots, d\}$. Considering just the outcomes in bins $\{1, \dots, d\}$ yields the approximate system of equations

$$\begin{aligned}
 -\Re \rho_{1,1+k} &= \frac{1}{2\epsilon} (p_1^{(k)} - p_1^{(0)}) \\
 &\vdots \\
 -\Re \rho_{k,2k} &= \frac{1}{2\epsilon} (p_k^{(k)} - p_k^{(0)}) \\
 \Re \rho_{1,k+1} - \Re \rho_{k+1,2k+1} &= \frac{1}{2\epsilon} (p_{k+1}^{(k)} - p_{k+1}^{(0)}) \\
 &\vdots \\
 \Re \rho_{d-2k,d-k} - \Re \rho_{d-k,d} &= \frac{1}{2\epsilon} (p_{d-k}^{(k)} - p_{d-k}^{(0)}) \\
 \Re \rho_{d-2k+1,d-k+1} &= \frac{1}{2\epsilon} (p_{d-k+1}^{(k)} - p_{d-k+1}^{(0)}) \\
 &\vdots \\
 \Re \rho_{d-k,d} &= \frac{1}{2\epsilon} (p_d^{(k)} - p_d^{(0)})
 \end{aligned} \tag{22}$$

valid for $k \in \{1, \dots, d-1\}$ and small ϵ . In principle, this (generally overdetermined) system of d equations enables one to determine the real parts of $\rho_{1,k+1}, \dots, \rho_{d-k,d}$, the nonzero part of the k th diagonal of ρ . (Note that the $-k$ th diagonal of ρ is just the conjugate of the k th diagonal.) The elements on diagonals $\{-k, \dots, 0, \dots, k\}$ will be collectively referred to as the k -band.

To determine the imaginary parts of the k th diagonal, each pair of components $\rho_{x,x}$ and $\rho_{x+k,x}$ must be mixed with a relative phase of $\pi/2$. Let Φ_k be the operation

$$\Phi_k |x\rangle = \varpi_k^x |x\rangle \quad (23)$$

where $\varpi_k = e^{i\pi/2k}$ is the k th principle root of i . Suppose we apply Φ_k and then S_k to the photon and measure its mode. The probability of observing bin x in this case is

$$\begin{aligned} p_x^{(k)'} &\equiv \langle x | S_k \Phi_k \rho \Phi_k^\dagger S_k^\dagger | x \rangle \\ &= \langle x | S_k \rho' S_k^\dagger | x \rangle \end{aligned} \quad (24)$$

where $\rho' = \Phi_k \rho \Phi_k^\dagger$. Using $\rho'_{xy} = \varpi_k^{x-y} \rho_{xy}$ we have

$$\begin{aligned} p_x^{(k)'} &= (\epsilon |x - k\rangle + \langle x| - \epsilon \langle x + k|) \rho' (\epsilon |x - k\rangle + |x\rangle - \epsilon |x + k\rangle) \\ &= \rho'_{xx} + \epsilon (\rho'_{x-k,x} + \rho'_{x,x-k} - \rho_{x+k,k} - \rho_{x,x+k}) + \mathcal{O}(\epsilon^2) \\ &= \rho_{xx} + \epsilon (\varpi_k^{-k} \rho_{x-k,x} + \varpi_k^k \rho_{x,x-k} - \varpi_k^k \rho_{x+k,k} - \varpi_k^{-k} \rho_{x,x+k}) + \mathcal{O}(\epsilon^2). \end{aligned} \quad (25)$$

Since $\varpi_k^{\pm k} = \pm i$ we have

$$\begin{aligned} p_x^{(k)'} &= \rho'_{xx} + \epsilon (-i \rho_{x-k,x} + i \rho_{x,x-k} - i \rho_{x+k,k} + i \rho_{x,x+k}) + \mathcal{O}(\epsilon^2) \\ &= \rho_{xx} + 2\epsilon (\Im \rho_{x-k,x} - \Im \rho_{x,x+k}) + \mathcal{O}(\epsilon^2). \end{aligned} \quad (26)$$

Using the probabilities of bins 1 through d we obtain a system of d equations which, together with $p^{(0)}$, in principle enables determination of the imaginary parts of the k th diagonal.

To summarize: $p^{(0)}$ determines the diagonal elements of ρ , then $p^{(k)}$ and $p^{(k)'} (obtained by applying S_k and $S_k \Phi_k$, respectively) determine the off-diagonal elements of ρ that lie k positions above and below the diagonal. Thus, all d^2 elements of the single-qudit density matrix can be probed by $1 + 2(d - 1) = 2d - 1$ experimental settings with d outcomes each.$

To relate this to the experimental approach, we note that the EOM effects a mode mixing operation T_δ of the form [cf. Eq. (4)]

$$T_\delta |x\rangle = \sum_{k=-\infty}^{\infty} J_k(\delta) |x+k\rangle \quad (27)$$

As shown in Fig. 7(b), $J_k(\delta)$ is oscillatory in k with approximate support $k \in \{-\lceil \delta \rceil, \dots, \lceil \delta \rceil\}$. That is, T_δ is approximately a linear combination of $\{S_1, \dots, S_{\lceil \delta \rceil}\}$, mixing each x with multiple modes ranging from roughly $x - \lceil \delta \rceil$ to $x + \lceil \delta \rceil$ and yielding information about the elements in the $\lceil \delta \rceil$ -band of ρ [Fig. 7(c)]. If the elements in the $(\lceil \delta \rceil - 1)$ -band have already been determined, T_δ provides new information primarily concerning the $\lceil \delta \rceil$ -band. In our specific experiments, we applied a set of random phases with the pulse shaper in addition to the modulation, allowing us to probe linear combinations of both the real and imaginary parts of the k -band similar to the S_k and Φ_k operators in the discussion above. In this way, a collection of measurements with modulation

indices uniformly distributed $\delta \in [0, k]$ along with random phases is sufficient to fully probe the k -band. And since the complete $d \times d$ density matrix ρ is encompassed within the $(d - 1)$ -band, this model suggests $\delta_{\max} \sim \mathcal{O}(d)$ as an ideal design choice for efficient state tomography with our method.

This intuitive picture reveals how an EOM-based mode mixer can be designed to respond to all elements of the density matrix, yet it does not quantify how well such measurements span the space of hermitian operators; if the mixing weights are small or produce excessive scattering into modes outside of the computational space, the number of observations required to reach a desired accuracy will be high, even if the measurements are tomographically complete. To address this question explicitly, we next explore the specifics of the measurement operations through singular value decomposition. Suppose we have a set of R different measurements $\mathbf{M}_1, \dots, \mathbf{M}_R$, each of which is described by a set of positive operator-valued measures (POVM): $\mathbf{M}_i = \{M_{i1}, \dots, M_{im_i}\}$ where m_i is the number of possible outcomes of \mathbf{M}_i (d in our case) and each M_{ij} is a positive semidefinite operator. To account for photon scattering outside of the measured computational space, the operators need not conserve probability. When \mathbf{M}_i is performed on state ρ , the probability of outcome j is

$$p_{ij} = \text{Tr}(M_{ij} \rho). \quad (28)$$

This may be summarized as the linear system

$$p = O^\dagger \vec{\rho} \quad (29)$$

where $\vec{\rho}$ is the vectorization of ρ and

$$O = [\vec{M}_{11} \ \cdots \ \vec{M}_{1m_1} \ \cdots \ \vec{M}_{R1} \ \cdots \ \vec{M}_{Rm_R}] . \quad (30)$$

is a $d^2 \times m$ matrix where $m = m_1 + \cdots + m_R$. The measurement set $\mathcal{M} = \{\mathbf{M}_1, \dots, \mathbf{M}_R\}$ is called *informationally complete* if ρ is uniquely determined by measured probabilities p . This occurs iff O is full rank (i.e., rank d^2), which requires at least d different measurements ($R \geq d$) with d outcomes each. Practically speaking, however, more important than the attainment of informational completeness is the actual distribution of singular values $s_1 \geq \cdots \geq s_{d^2}$ of O . If the singular values are all of order 1, then \mathcal{M} determines all components of ρ with comparable sensitivity. But any singular values of O that are much smaller than 1 correspond to components of ρ to which \mathcal{M} is only weakly sensitive. The estimates of such components will be susceptible to statistical fluctuations in the experimental data.

As examples, we compute singular values for a variety of settings at $d = 8$ —the maximum single-qudit dimension characterized experimentally. We consider $R = 2d = 16$ settings, which we found sufficient to obtain robust and repeatable distributions. Each measurement configuration involves a modulation index δ chosen uniformly at random in the interval $[0, \delta_{\max}]$, preceded by a random phase vector $\vec{\phi} \in [0, 2\pi]^d$ applied by the pulse shaper. Figure 7(d) plots histograms of singular values from 2,000 measurement matrices O each for $\delta_{\max} \in \{4, 8, 16\}$, corresponding to $d/2$, d , and $2d$ at $d = 8$. A narrow peak around $\log_{10} s \approx 0$ indicates high and comparable sensitivity for all elements in the density matrix; indeed, a complete set of mutually unbiased bases—an ideal choice for tomography [46]—possesses d^2 equal singular values at $s = 1$. By this criterion, $\delta_{\max} = d = 8$ is seen to provide the most efficient measurement distribution of the three examples. For the smaller index of $\delta_{\max} = 4$, the main peak is accompanied by a strong tail indicating small sensitivity to an appreciable percentage of the matrix elements. And for the larger index $\delta_{\max} = 16$, the main peak shifts to lower values, which reflects “over-modulating” of the quantum state; taking δ_{\max} beyond d increases the probability of scattering the input outside of the computational space [see Fig. 7(c)] without any improvement to mixing within the space. Our numerical findings therefore join the simple theory above in suggesting a maximum modulation index of $\delta_{\max} \approx d$. Nevertheless, we note that all cases in Fig. 7(d) sample the Hilbert space comprehensively, so that any $\delta_{\max} \sim \mathcal{O}(d)$ is likely to prove sufficient in a tomographic context. Ultimately, we emphasize that the theory developed here is meant to provide heuristic guidelines for implementing EOM-based frequency-bin tomography, not to imply the optimality of our spe-

cific measurements. For example, it is quite possible that alternative distributions for δ —e.g., other than the uniform draw $\delta \in [0, \delta_{\max}]$ —may show more favorable properties with further research. Nevertheless, our reliance on Bayesian estimation ensures that these questions need not be answered for useful inference.

Extending to the experimental task of two-qudit characterization, our model thus indicates one should consider joint measurements of the form $T_{\delta_1} \otimes T_{\delta_2}$ with $\delta_1, \delta_2 \in [0, d]$. Considering our experimental values of $\delta_{\max} = 2.5$ for the PPLN tests and 3.4 for the MRR, δ_{\max} falls in the range of $0.425d$ and $1.7d$ for all qudit dimensions examined, aligning well with the $\mathcal{O}(d)$ desideratum. Since the entire density matrix of a d^2 -dimensional two-qudit state requires specification of $d^4 - 1$ real parameters and each experimental pulse shaper/EOM setting r provides d^2 outcomes, $R \sim \mathcal{O}(d^2)$ would be expected to be required to fully probe the two-qudit Hilbert space. Empirically, however, we were able to attain low-error state reconstruction with fewer measurements: e.g., $R = 10$ instead of 25 for the $d = 5$ PPLN BFC [Fig. 2(a)] and $R = 30$ instead of 64 for the $d = 8$ MRR BFC [Fig. 4(b)]. Several aspects are likely responsible for this reduction. From a broad perspective, the positive semidefiniteness of the density matrix imposes additional constraints that are not reflected in a linear system analysis but automatically accounted for in the Bayesian inference procedure. Thus, simply comparing the number of measurements with the number of parameters generally yields an overly pessimistic assessment of the information required for tomography.

Moreover, we suspect that specific features of our quantum state also contribute to more efficient reconstruction. The highly correlated nature revealed in the first JSI measurement implies that only off-diagonal elements of the form $\rho_{(xx)(x+k,x+k)}$ —i.e., those satisfying biphoton energy conservation—can be significantly different than zero. This reduces the number of appreciable elements in our density matrix from $\mathcal{O}(d^4)$ to $\mathcal{O}(d^2)$. Of course, the other off-diagonal density matrix elements are not strictly zero in practice, and the Bayesian model neither requires nor assumes such a simplification. However, the strong frequency correlations do eliminate a large portion of the Hilbert space from consideration. Consequently, future experiments with more general states may require more experimental settings for low-uncertainty estimation than we have currently used.

Data Availability

The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Code Availability

Analysis code used in this study is available from the corresponding author on request.

[1] M. Kues, C. Reimer, J. M. Lukens, W. J. Munro, A. M. Weiner, D. J. Moss, and R. Morandotti, Nat. Photonics

13, 170 (2019).

[2] M. Kues, C. Reimer, P. Roztocki, L. R. Cortés, S. Sciara, B. Wetzel, Y. Zhang, A. Cino, S. T. Chu, B. E. Little, D. J. Moss, L. Caspani, J. Azaña, and R. Morandotti, *Nature* **546**, 622 (2017).

[3] P. Imany, J. A. Jaramillo-Villegas, O. D. Odele, K. Han, D. E. Leaird, J. M. Lukens, P. Lougovski, M. Qi, and A. M. Weiner, *Opt. Express* **26**, 1825 (2018).

[4] H.-H. Lu, J. M. Lukens, N. A. Peters, B. P. Williams, A. M. Weiner, and P. Lougovski, *Optica* **5**, 1455 (2018).

[5] R. Blume-Kohout, *New J. Phys.* **12**, 043034 (2010).

[6] J. M. Lukens, K. J. H. Law, A. Jasra, and P. Lougovski, *New J. Phys.* **22**, 063038 (2020).

[7] M. Erhard, M. Krenn, and A. Zeilinger, *Nat. Rev. Phys.* **2**, 365 (2020).

[8] D. Cozzolino, B. Da Lio, D. Bacco, and L. K. Oxenløwe, *Adv. Quantum Technol.* **2**, 1900038 (2019).

[9] J. T. Barreiro, T.-C. Wei, and P. G. Kwiat, *Nat. Phys.* **4**, 282 (2008).

[10] N. J. Cerf, M. Bourennane, A. Karlsson, and N. Gisin, *Phys. Rev. Lett.* **88**, 127902 (2002).

[11] S. Ecker, F. Bouchard, L. Bulla, F. Brandt, O. Kohout, F. Steinlechner, R. Fickler, M. Malik, Y. Guryanova, R. Ursin, and M. Huber, *Phys. Rev. X* **9**, 041042 (2019).

[12] T. Vértesi, S. Pironio, and N. Brunner, *Phys. Rev. Lett.* **104**, 060401 (2010).

[13] J. Wang, S. Paesani, Y. Ding, R. Santagati, P. Skrzypczyk, A. Salavrakos, J. Tura, R. Augusiak, L. Mančinska, D. Bacco, D. Bonneau, J. W. Silverstone, Q. Gong, A. Acín, K. Rottwitt, L. K. Oxenløwe, J. L. O'Brien, A. Laing, and M. G. Thompson, *Science* **360**, 285 (2018).

[14] X. Qiang, X. Zhou, J. Wang, C. M. Wilkes, T. Loke, S. O'Gara, L. Kling, G. D. Marshall, R. Santagati, T. C. Ralph, J. B. Wang, J. L. O'Brien, M. G. Thompson, and J. C. F. Matthews, *Nat. Photonics* **12**, 534 (2018).

[15] J. Bavaresco, N. H. Valencia, C. Klöckl, M. Pivoluska, P. Erker, N. Friis, M. Malik, and M. Huber, *Nat. Phys.* **14**, 1032 (2018).

[16] F. Brandt, M. Hiekkämäki, F. Bouchard, M. Huber, and R. Fickler, *Optica* **7**, 98 (2020).

[17] H.-H. Lu, E. M. Simmernan, P. Lougovski, A. M. Weiner, and J. M. Lukens, *Phys. Rev. Lett.* **125**, 120503 (2020).

[18] A. Martin, T. Guerreiro, A. Tiranov, S. Designolle, F. Fröwis, N. Brunner, M. Huber, and N. Gisin, *Phys. Rev. Lett.* **118**, 110501 (2017).

[19] T. Ikuta and H. Takesue, *New J. Phys.* **19**, 013039 (2017).

[20] G. Moody, L. Chang, T. J. Steiner, and J. E. Bowers, *AVS Quantum Sci.* **2**, 041702 (2020).

[21] Z. Yang, M. Jahanbozorgi, D. Jeong, S. Sun, O. Pfister, H. Lee, and X. Yi, *Nature Communications* **12** (2021), 10.1038/s41467-021-25054-z.

[22] A. W. Elshaari, W. Pernice, K. Srinivasan, O. Benson, and V. Zwiller, *Nat. Photonics* **14**, 285 (2020).

[23] H.-H. Lu, A. M. Weiner, P. Lougovski, and J. M. Lukens, *IEEE Photon. Technol. Lett.* **31**, 1858 (2019).

[24] J. M. Lukens, H.-H. Lu, B. Qi, P. Lougovski, A. M. Weiner, and B. P. Williams, *J. Light. Technol.* **38**, 1678 (2020).

[25] J. Liu, E. Lucas, A. S. Raja, J. He, J. Riemensberger, R. N. Wang, M. Karpov, H. Guo, R. Bouchand, and T. J. Kippenberg, *Nat. Photonics* **14**, 486 (2020).

[26] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods* (Springer, New York, 1999).

[27] A. Peres, *Phys. Rev. Lett.* **77**, 1413 (1996).

[28] G. Vidal and R. F. Werner, *Phys. Rev. A* **65**, 032314 (2002).

[29] P. Imany, N. B. Lingaraju, M. S. Alshaykh, D. E. Leaird, and A. M. Weiner, *Sci. Adv.* **6**, eaba8066 (2020).

[30] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, *Phys. Rev. Lett.* **105**, 150401 (2010).

[31] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, *Science* **364**, 260 (2019).

[32] H.-Y. Huang, R. Kueng, and J. Preskill, *Nat. Phys.* **16**, 1050 (2020).

[33] J. Liu, G. Huang, R. N. Wang, J. He, A. S. Raja, T. Liu, N. J. Engelsen, and T. J. Kippenberg, *Nat. Commun.* **12**, 1 (2021).

[34] J. Liu, A. S. Raja, M. H. Pfeiffer, C. Herkommer, H. Guo, M. Zervas, M. Geiselmann, and T. J. Kippenberg, *Opt. Lett.* **43**, 3200 (2018).

[35] K. V. Myilswamy, M. S. Alshaykh, H.-H. Lu, J. Liu, D. E. Leaird, T. J. Kippenberg, and A. M. Weiner, in *CLEO: 2021* (Optical Society of America, 2021) p. JM3F.5.

[36] H.-H. Lu, J. M. Lukens, N. A. Peters, O. D. Odele, D. E. Leaird, A. M. Weiner, and P. Lougovski, *Phys. Rev. Lett.* **120**, 030502 (2018).

[37] B. P. Williams and P. Lougovski, *New J. Phys.* **19**, 043003 (2017).

[38] H.-H. Lu, J. M. Lukens, B. P. Williams, P. Imany, N. A. Peters, A. M. Weiner, and P. Lougovski, *npj Quantum Inf.* **5**, 24 (2019).

[39] E. M. Simmernan, H.-H. Lu, A. M. Weiner, and J. M. Lukens, *Opt. Lett.* **45**, 2886 (2020).

[40] H.-J. Sommers and K. Życzkowski, *J. Phys. A: Math. Gen.* **36**, 10083 (2003).

[41] V. A. Osipov, H.-J. Sommers, and K. Życzkowski, *J. Phys. A: Math. Theor.* **43**, 055302 (2010).

[42] F. Mezzadri, *Not. Am. Math. Soc.* **54**, 592 (2007).

[43] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, UK, 2003).

[44] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White, *Statist. Sci.* **28**, 424 (2013).

[45] P. E. Jacob, J. O'Leary, and Y. F. Atchadé, *J. R. Statist. Soc. B* **82**, 543 (2020).

[46] W. K. Wootters and B. D. Fields, *Ann. Phys.* **191**, 363 (1989).