ARTICLE   OPEN

A controlled-NOT gate for frequency-bin qubits

Hsuan-Hao Lu [1], Joseph M. Lukens [2], Brian P. Williams[2], Poolad Imany[1], Nicholas A. Peters [2,3], Andrew M. Weiner [1] and Pavel Lougovski [2]

The realization of strong photon–photon interactions has presented an enduring challenge across photonics, particularly in quantum computing, where two-photon gates form essential components for scalable quantum information processing (QIP). While linear-optic schemes have enabled probabilistic entangling gates in spatio-polarization encoding, solutions for many other useful degrees of freedom remain missing. In particular, no two-photon gate for the important platform of frequency encoding has been experimentally demonstrated, due in large part to the additional challenges imparted by the mismatched wavelengths of the interacting photons. In this article, we design and implement an entangling gate for frequency-bin qubits, a coincidence-basis controlled-NOT (CNOT), using line-by-line pulse shaping and electro-optic modulation. We extract a quantum unitary fidelity of 0.91 ± 0.01 via a parameter inference approach based on Bayesian machine learning, which enables accurate gate reconstruction from measurements in the two-photon computational basis alone. Our CNOT imparts a single-photon frequency shift controlled by the frequency of another photon—an important capability in itself—and should enable new directions in fiber-compatible QIP.

INTRODUCTION

As carriers of quantum information, optical photons feature a host of valuable attributes, such as immunity to environmentally induced decoherence, availability of precise tools for state control, and room temperature operation, enabling quantum information processing (QIP)[1] in a variety of encodings such as space/polarization[2–4] and temporal modes.[5–7] Frequency-bin encoding—which offers additional advantages in terms of compatibility with state-of-the-art fiber-optic networks—has advanced rapidly in recent years, facilitated by the development of integrated frequency-bin photon sources[8–11] and quantum gates based on both nonlinear-optical[12,13] and electro-optical[14,15] mixing approaches. However, two-photon entangling gates for frequency bins have yet to be realized on any platform.

Such entangling gates are required for universal QIP, for an arbitrary quantum operation can be constructed with single-qubit rotations plus a two-qubit entangling gate.[1] While photonics excels for single-qubit gates, the inherent difficulty in realizing photon–photon interactions has made the two-qubit gate a persistent obstacle in photonic QIP. In the absence of a sufficient nonlinearity, such gates can still be achieved via quantum interference, ancilla photons, and single-photon detection. While two-qubit gates succeed only probabilistically in this paradigm, linear-optical quantum computation (LOQC)[2] is in principle scalable with polynomial auxiliary resource requirements and has laid the foundation for many subsequent advances in photonic QIP.[3,16–25] It is this approach which we invoked in proposing *spectral* LOQC—a universal QIP scheme tailored to frequency-bin qubits which makes use of electro-optic phase modulators (EOMs) and Fourier-transform pulse shapers (PSs).[26] Systems implementing designs from spectral LOQC, termed

"quantum frequency processors" (QFPs), have been utilized to demonstrate coherent single-photon operations with near-unity fidelity,[14,15] but a two-photon gate has heretofore proven elusive.

Theoretically, we previously discovered EOM/PS configurations capable of realizing ancilla-based two-qubit gates in spectral LOQC.[26] Yet if one relaxes the gate requirements slightly, by conditioning on the presence of a photon in each pair of qubit modes, it is well-known in standard LOQC that one can engineer a two-qubit gate with no ancillas and success probability $\mathcal{P} = 1/9$.[18,19] Assuming a quantum nondemolition measurement is unavailable, such gates are destructive (succeeding only when both information-carrying photons are detected). Yet they require only two-fold coincidences for characterization, making them excellent choices for experimental studies of basic quantum computing functionalities.

To explore two-qubit coincidence-basis gates with a QFP, we follow the optimization approach in refs. [14,26], numerically finding phase patterns for an EOM/PS sequence which maximize success $\mathcal{P}$ constrained to fidelity $\mathcal{F} \geq 0.9999$. Specifically, with U_{ideal} defined as the desired two-qubit unitary and W the actual Hilbert space transformation,

$$\mathcal{P} = \frac{\text{Tr}(W^\dagger W)}{d} \tag{1}$$

$$\mathcal{F} = \frac{\left|\text{Tr}(U_{ideal}^\dagger W)\right|^2}{d^2 \mathcal{P}}, \tag{2}$$

where $d = 4$ is the dimensionality of the subspace spanned by the coincidence basis.[26] In order to facilitate experimental implementation, we restrict our simulations to sinewave-only electro-optic modulation. We find that a 3EOM/2PS QFP can realize a

frequency-bin CNOT at the optimal success probability of $\mathcal{P} = 1/9$, while a smaller 2EOM/1PS circuit can do so with reduced success: $\mathcal{P} = 0.0445$. (See Methods for the specific EOM/PS modulation patterns.) Due to equipment availability and system complexity, we elect to implement this simpler 2EOM/1PS CNOT in the experiments below.

RESULTS

Figure 1 provides a schematic of the setup. The gate itself comprises the central EOM/PS/EOM sequence (the QFP), and the frequency bins for encoding are defined according to $\omega_n = \omega_0 + n\Delta\omega$, where $\omega_0 = 2\pi \times 193.45$ THz and $\Delta\omega = 2\pi \times 25$ GHz, corresponding to the standard ITU grid and facilitating low-crosstalk, line-by-line shaping by our 10 GHz resolution pulse shapers. The specific bins for encoding follow in Fig. 2a, where $\{C_0, C_1\}$ and $\{T_0, T_1\}$ denote logical $|0\rangle$ and $|1\rangle$ for the control and target, respectively. This particular mode placement makes sense conceptually: mode C_0 is spectrally isolated from the target's logical bins, ensuring a photon in mode C_0 leaves the target unchanged; on the other hand, bin C_1 is close to both target bins, able to be coupled to T_0 and T_1 with equal strength.

Since this gate is based on a linear-optical network, we can estimate its performance using coherent-state-based characterization,[14,27] i.e., probing it with an electro-optic frequency comb and measuring the output spectrum for different input frequency superpositions. This technique allows us to estimate the mode transformation matrix V, which controls how input mode operators $\hat{a}_n$ at each frequency ω_n transform to the output operators $\hat{b}_n$: $\hat{b}_n = \sum_{n'} V_{nn'} \hat{a}_{n'}$. The mode matrix V, averaged over five independent measurements and projected onto the four computational modes, is shown in Fig. 2b. We utilize phasor notation to represent the complex elements $V_{nn'}$; the filled color reflects the amplitude on a logarithmic scale, normalized to the maximum value in the matrix (0.499), and the arrow marks out the phase. (See Methods for values of all matrix elements including uncertainty.) From this matrix V, we can compute the equivalent two-photon state transformation matrix W,[26] plotted for the coincidence basis in Fig. 2c and also normalized to its peak magnitude of 0.222. We note that the implemented mode transformation (Fig. 2b) has the same eight high-amplitude elements as in the original path-encoded CNOT.[18] The phases differ, however, as there are a continuum of combinations which can yield the desired two-photon transformation W. The particular

set was selected numerically as optimal given our experimental constraints.

Because this estimate predicts all four of the large elements of W to be in-phase, the corresponding inferred fidelity is $\mathcal{F}_{inf} = 0.995 \pm 0.001$; the success probability is $\mathcal{P}_{inf} = 0.0460 \pm 0.0005$. Both values are with respect to the ideal CNOT and in good agreement with theory. We emphasize that, unlike single-qubit gates which act on photons independently, two-qubit entangling gates rely on quantum interference effects that are inherently absent with high-flux laser fields. Thus this inferred fidelity is only an indirect estimate, based on extrapolating measured one-photon interference results to the two-photon case. Nevertheless, it provides strong initial evidence for the phase coherence and proper operation of our gate.

To test our gate with truly quantum states, however, we prepare a biphoton frequency comb (BFC) by pumping a 35 mm-long periodically poled lithium niobate (PPLN) waveguide with a continuous-wave Ti:sapphire laser (at ~4.5 mW) under type-0 phase matching, followed by filtering with a Fabry–Perot etalon with 25 GHz mode spacing and a full-width at half-maximum linewidth of 1.8 GHz (see Fig. 1). The BFC pulse shaper subsequently selects specific modes as input to the gate. By translating the pump frequency to four different values (as shown in Fig. 2a) and filtering out all but the desired modes using the BFC pulse shaper, we can prepare all inputs from the two-qubit computational basis: $|C_0 T_0\rangle = |1_{\omega_0} 1_{\omega_7}\rangle$, $|C_0 T_1\rangle = |1_{\omega_0} 1_{\omega_8}\rangle$, $|C_1 T_0\rangle = |1_{\omega_6} 1_{\omega_7}\rangle$, and $|C_1 T_1\rangle = |1_{\omega_6} 1_{\omega_8}\rangle$. To ensure the photon flux remains constant across the four inputs, we tune the PPLN waveguide temperature to align the peak of the phase-matching spectrum with the pump laser frequency. After the gate, the output photons are frequency-demultiplexed: we send control photon bins to detector A and target photon bins to detector B.

Figure 3a is a conceptual example of the interference underpinning the CNOT, where the rails denote particular frequency bins and the lines trace out probability amplitudes of single photons initially in bins C_1 and T_1; blue follows the control, red the target, and the thickness is proportional to the squared amplitude. Each EOM serves as a multimode interferometer mixing all bins simultaneously; in this particular example, the phases applied by the QFP shaper produce destructive interference of the two amplitudes yielding the output $|C_1 T_1\rangle$, leaving only the possibility $|C_1 T_0\rangle$ in the coincidence basis (the characteristic CNOT bit flip). This picture highlights that, while the general interference phenomena remain the same between path and frequency encoding, the basic manipulations are significantly different: standard beamsplitters interface two input modes with two outputs, while EOMs couple, in principle, infinitely many. Accordingly, this schematic cannot be taken as fully quantitative, but it does, through line weights, give an idea of the coupling magnitudes in this example. Such a lack of direct correspondence between frequency-bin and path primitives is the reason for our use of numerical optimization of the full transformation, rather than constructing and combining individual frequency-bin beamsplitters.

Figure 3b shows the measured coincidences for all 16 input/output mode combinations, integrated over 600 s for each point. As expected, inputs with a photon in control mode 0 retain their quantum state, whereas a photon in control mode 1 leads to a flip in the output target qubit. In Fig. 3c, we plot the accidentals as determined by the product of the singles counts and our timing resolution.[28,29] The nonuniform distribution of accidentals stems from the fact that the singles counts vary significantly across input/output state combinations. Indeed, this is a natural feature of coincidence-basis gates: they are designed to discard cases when one of the qubit spaces is empty or doubly occupied, so that photon detection rates in a specific mode can change without impacting the designed operation.
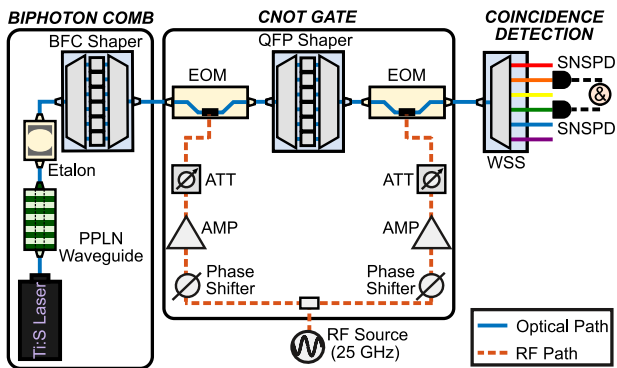

Fig. 1 Experimental setup. PPLN (SRICO Model 2000-1550), Etalon (Optoplex 25 GHz C-Band), BFC Shaper (Finisar WaveShaper 1000A), EOMs (EOSpace 40 Gbps phase modulators), QFP Shaper (Finisar Waveshaper 4000A), WSS wavelength selective switch (Finisar 1 × 9 Flexgrid), SNSPD superconducting nanowire single-photon detector (Quantum Opus model Opus One, >80% detection efficiency), ATT variable radio-frequency (RF) attenuator, AMP RF amplifier
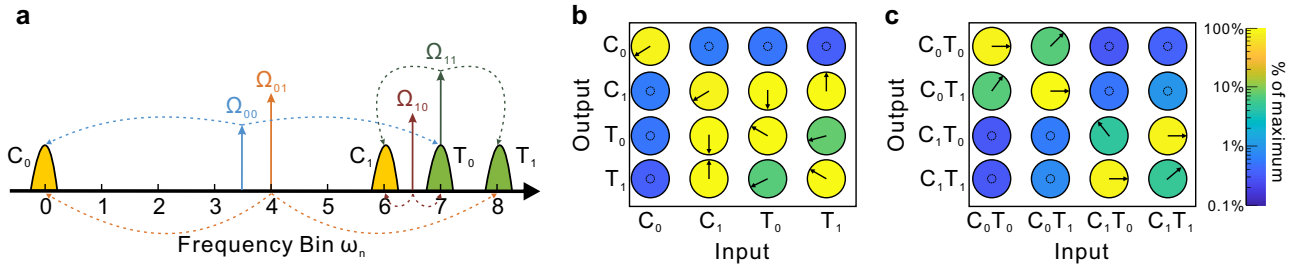
Fig. 2 **a** Mode definitions for frequency-bin control and target qubits. The labels $\{\Omega_{00}, \Omega_{01}, \Omega_{10}, \Omega_{11}\}$ mark the pump frequency values (divided by two) needed to produce each of the computational basis states. **b** Experimentally obtained complex mode transformation V. **c** Inferred two-photon transformation W obtained from permanents of 2×2 submatrices of V. For **b** and **c**, we use phasor notation to represent the complex elements, with filled color signifying the amplitude (normalized by the matrix's maximum value, and shown on a logarithmic scale), and the arrow depicting the phase. Dotted circles denote phases we could not retrieve due to weak amplitudes. (See Methods for details.)

Such information-bearing features in the accidentals suggest that incorporating knowledge from single-detector events—as well as the coincidences—can add significant value for quantifying the performance of our gate in the presence of noise. To utilize all of our experimental data in a consistent fashion, we make use of Bayesian machine learning techniques to implement a numerical parameter inference approach built on Bayesian mean estimation (BME).[30] In the context of quantum state retrieval, BME is a powerful method which returns uncertainties on any quantity directly and makes efficient use of all available information, in the sense that the confidence in any estimate naturally reflects the amount of data gathered.[31] BME models for photon pairs including single-detector events have been developed as well, permitting extraction of the quantum pathway efficiencies in conjunction with estimates of the input density matrix.[32] In our BME model here, not only do we account for noise effects, but we can also retrieve meaningful estimates of the full complex matrix V, even though we only prepare and measure states in the computational basis. This represents an entirely new capability in two-photon gate analysis, for previously such truth-table measurements as in Fig. 3b have only been used to establish *magnitudes* in the matrix transformation, with superposition states required to assess the phase.[20]

In our model, the unknown parameters to retrieve include the mode matrix V, the pair generation probability μ, and the total system efficiencies η_A and η_B preceding detection at the control and target photon detectors, respectively. Obtained from independent measurements, and thus taken as fixed and known, are the dark count probabilities d_A and d_B. All probabilities $\{\mu, d_A, d_B\}$ are specified for one resolving time τ (~1.5 ns). For the input photon state $|C_k T_l\rangle$ ($k, l \in \{0, 1\}$) with detectors A and B set to respond to output modes C_r and T_s ($r, s \in \{0, 1\}$), respectively, the probability of a coincidence between detectors A and B is

$$p_{AB} = \mu \eta_A \eta_B |V_{C_r C_k} V_{T_s T_l} + V_{C_r T_l} V_{T_s C_k}|^2 + 2p_A p_B. \quad (3)$$

Here, p_A and p_B are the marginal probabilities for clicks on A or B, irrespective of clicks on the other, during a given time τ. This formula thus contains both a correlated term (from photons of the same pair) and an accidental term. The latter, equal to $2p_A p_B$,[28,29] represents the chance of simultaneous clicks in which at least one detector registers a dark count, or the photons come from different pairs (see Methods for details).

The marginal probabilities p_A and p_B can be found by summing the contributions from each possible number of photons N being present in the monitored mode, sketched formally as, e.g., $p_A = \sum_N P(\text{click}|N \text{ photons})P(N \text{ photons})$. Writing out each term for N = 0, 1, 2, and simplifying, we ultimately arrive at the probabilities for a click on either detector within a time τ (see

Methods):

$$p_A = \mu \eta_A \left(|V_{C_r C_k}|^2 + |V_{C_r T_l}|^2 \right) + d_A$$
$$p_B = \mu \eta_B \left(|V_{T_s C_k}|^2 + |V_{T_s T_l}|^2 \right) + d_B, \quad (4)$$

valid under the assumptions $\mu, \eta_A, \eta_B, d_A, d_B \ll 1$—satisfied in our experiment. In words, a detector can click from either of the following: (i) a photon pair is generated (μ), one of the photons is sent to the monitored frequency bin (through V), and the photon reaches the output and is successfully detected (η_A, η_B); or (ii) the detector fires spontaneously (d_A, d_B). Crucially, the singles probabilities [Eq. (4)] depend only on the moduli of the V-matrix elements, whereas the coincidences also depend on the relative phase [via the permanent term in Eq. (3)]. It is this complementary dependence which underpins our ability to extract the full complex matrix from experimental data.

Specifically, for a single preparation/measurement configuration we possess three numbers as data: clicks on A (N_A), clicks on B (N_B), and coincidences (N_{AB}). This gives us the multinomial likelihood for this specific input/output configuration ($|C_k T_l\rangle \rightarrow |C_r T_s\rangle$):

$$P\left(\mathcal{D}^{C_r T_s}_{C_k T_l} | \beta\right) = (p_A - p_{AB})^{N_A - N_{AB}} (p_B - p_{AB})^{N_B - N_{AB}}$$
$$\times p_{AB}^{N_{AB}} (1 - p_A - p_B + p_{AB})^{M - N_A - N_B + N_{AB}}, \quad (5)$$

where $\mathcal{D}^{C_r T_s}_{C_k T_l} = \{N_A, N_B, N_{AB}\}$ contains all data values for the specific configuration. We have also reexpressed the events to make them mutually exclusive: click on A only, happening $N_A - N_{AB}$ times; click on B only, occurring $N_B - N_{AB}$ times; coincidence between A and B (N_{AB} times); and no clicks (all remaining frames). The symbol β is shorthand for all model parameters ($\beta = \{V, \mu, \eta_A, \eta_B\}$), and M equals the total number of τ frames considered in one counting period (~4×10^{11} in our tests). The complete likelihood comprises 16 factors in the manner of Eq. (5) for all combinations of inputs and outputs. Incidentally, one could retrieve estimates of the system parameters from this likelihood directly, via conventional maximum likelihood estimation (MLE). Computationally simpler than BME, MLE generates only a point estimate, without intrinsic error bars, in contrast to BME which quantifies uncertainty by integrating over the full probability distribution.[31,32] We must emphasize that our model relies on explicit enumeration of system noise sources, and thus is restricted to a parameter space smaller than the set of all two-qubit operations. This stands in contrast to the standard procedure for gate characterization, quantum process tomography (QPT),[1,21,33,34] which is designed to recover a quantum operation treating the system as a black box. Nevertheless, process tomography is complex, requiring a number of measurements that grows exponentially with system size—and in our case these measurements require more components than what we have available. Physically motivated simplifications[35] and alternatives[36] to QPT are thus of significant value in quantum information, and so, in our particular case, the key question is
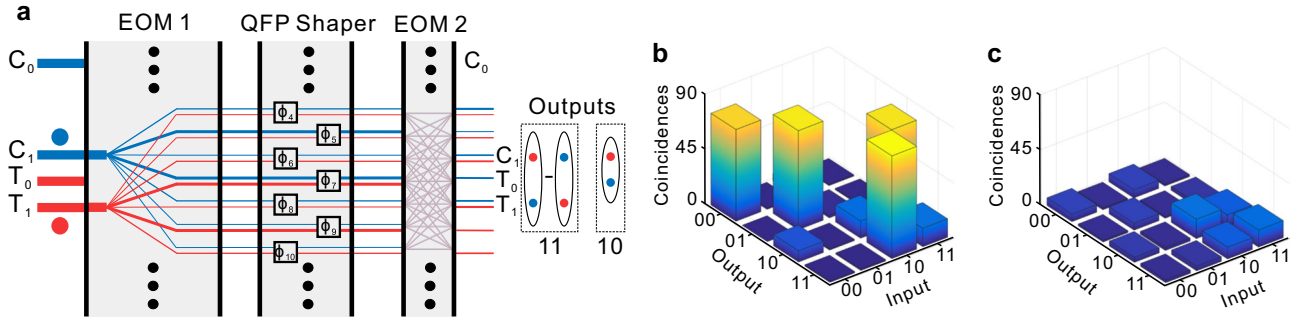
a

b

c

Fig. 3 **a** Conceptual frequency-bin interference in the QFP, for the case of input state 11. Quantum interference suppresses the result 11 at the output, leaving only state 10 in the coincidence basis. **b** Experimentally measured coincidences over 600 s for all input/output state combinations. **c** Estimated number of accidentals computed from the product of single detector counts

whether the model encompasses all relevant physical effects. Theoretically, our understanding of photon generation, frequency-bin operations, and detection suggest no additional sources of decoherence. Empirically as well, we previously explored frequency-bin Hong–Ou–Mandel interference in this QFP (the basic effect behind two-photon LOQC gates[2,3]) measuring ~97% visibility, where the reduction from unity was attributable to accidental coincidences resulting from system inefficiencies and dark counts.[15] Accordingly, our total channel model—linear-optical multiport plus loss, dark counts, and accidentals—is strongly justified, as well as validated ex post facto by the agreement with experiment below. Finally, it is important to note that Bayesian machine learning techniques can be applied toward any model, whether custom-tailored or built on more conventional quantum tomography, indicating additional opportunities for data analysis in a variety of quantum information experiments.[15,31,32]

Returning to details of the Bayesian analysis, we next assume uniform prior distributions over the interval (0, 1) for the unknown parameters μ, η_A, and η_B. For the complex matrix V, we express each element in terms of amplitude and phase: $V_{nn'} = r_{nn'}e^{i\phi_{nn'}}$. Since an overall scaling factor on V is indistinguishable from changes to η_A and η_B [Eqs. (3) and (4)], for concreteness we fix the Hilbert–Schmidt norm $\mathrm{Tr}(V^{\dagger}V) = 1.6558$, to match the ideal V matrix obtained from the numerical optimization [see Eq. (7)], thus constraining the sum of the squares of $r_{nn'}$. Otherwise, we let the squared amplitudes vary uniformly over all possible values subject to this condition. Because phase shifts on each of the modes before and after the multiport V are not physically significant, we are free to take some of the $\phi_{nn'}$ as given as well.[27] For convenience, we fix $\{\phi_{C_0C_0}, \phi_{C_1C_1}, \phi_{C_1T_0}, \phi_{C_1T_1}, \phi_{T_0C_1}, \phi_{T_1C_1}\}$ to their theoretical predictions, thus leaving 10 phases to be retrieved via BME.

With the likelihood and prior formally defined, in principle we are done: we have the posterior probability distribution from Bayes' rule, which represents complete knowledge of the parameters given the observed data. However, practically speaking, computing integrals or, equivalently, sampling from this many-parameter multimodal distribution is a formidable challenge. It is here that the techniques of Markov chain Monte Carlo (MCMC) sampling offer a solution, which—with minimal input—enable Bayesian machine learning of complex models. In our case, we employ slice sampling, an MCMC algorithm designed to produce a sequence of samples whose stationary distribution converges to the posterior.[37]

Using the predicted matrix V as an initial guess for the slice sampler, a procedure which we found important to speed up convergence given the large search space of 28 independent variables, we ultimately converge to the Bayesian fidelity estimate $\mathcal{F}_{BME} = 0.91 \pm 0.01$, where $\mathcal{F}$ is defined according to Eq. (2). Our truly quantum measurement does not reach the >0.99 classically
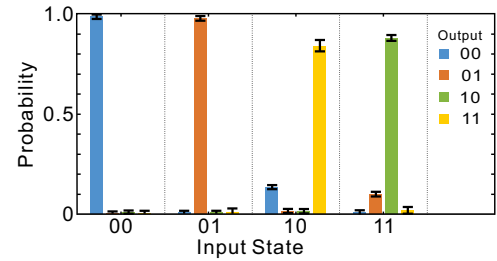

Fig. 4 Output state probabilities retrieved from BME, for each computational-basis input state

inferred $\mathcal{F}_{inf}$, which is a consequence of the relatively few coincidence counts (<100 in all cases) and additional noise from residual light. Nevertheless, the low uncertainty on $\mathcal{F}_{BME}$ indicates high confidence in our BME model, especially in light of its ability to retrieve the full complex fidelity with computational basis measurements. To see how $\mathcal{F}_{BME}$ translates into output state probabilities in the coincidence basis, we plot the Bayesian-estimated pathway probabilities in Fig. 4, where the four outcomes for each input state are normalized to sum to unity. The average probability for obtaining the correct output is 0.92 ± 0.01, computed by taking the mean of the four peaks in Fig. 4. (See Methods for details on all retrieved parameters, including the mode matrix V.)

DISCUSSION

Moving forward, it will be valuable to implement this gate with input states beyond just the computational basis, useful for implementing photonic QIP algorithms such as the variational quantum eigensolver[38] and Shor factoring.[39] Fundamentally, such states would also enable direct demonstration of the gate's ability to entangle two photons, offering independent verification of the quantum phase coherence which here we have estimated through Bayesian machine learning. Probing with such states is readily attainable in the QFP paradigm; for example, one could precede the CNOT operation with Hadamard gates on one or both input photons.[14,15] Yet cascading additional elements at the moment is limited by technical loss; we predict that we could not at present obtain coincidences above the accidental level with the additional equipment required. In order to concentrate on the basic physics in this proof-of-principle experiment, we have constructed our frequency-bin CNOT with off-the-shelf fiber-optic components. While the phase-only EOM/PS operations themselves are unitary, such commercial devices introduce significant additional loss, on the order of 3–4 dB per element. Accordingly, in scaling up to larger QIP systems, improving throughput is a substantial engineering goal.
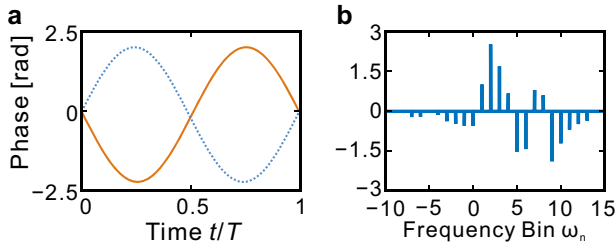
Fig. 5 Numerical solutions for the time-frequency phases required to implement coincidence-basis CNOT gate. **a** Temporal phase modulation applied to the first EOM (solid red) and second EOM (dotted blue), plotted over one period. **b** Phases applied to each frequency mode by the pulse shaper, where modes 0 and 6 denote the control bins $\{C_0, C_1\}$, and modes 7 and 8 represent the target bins $\{T_0, T_1\}$

Fortunately, state-of-the-art fabrication techniques presage significant improvements just over the horizon, via on-chip integration of the fundamental QFP elements. For example, process design kits from photonics foundries[40] suggest that the loss through a pulse shaper channel can be less than 1 dB, while recent experiments have demonstrated lithium-niobate EOMs with loss below 0.5 dB[41] and foundry-compatible EOMs with losses on the order of 1–2 dB.[42] Thus, an on-chip CNOT—identical to the present configuration in terms of functional components, but attaining <3 dB insertion loss—seems feasible with current technology, and we certainly anticipate even better performance as on-chip photonics continues to progress over the coming years.

In conclusion, we have realized an entangling gate on frequency-bin qubits. We confirm high-fidelity operation of the CNOT with two forms of characterization: coherent-state-based matrix retrieval and photon pair measurements in the computational basis. The classically inferred fidelity of $\mathcal{F}_{inf} = 0.995 \pm 0.001$ and Bayesian estimate $\mathcal{F}_{BME} = 0.91 \pm 0.01$ both demonstrate high performance in our system. As the sole realization of a two-photon entangling gate in frequency—and only the second CNOT in the entire field of time-frequency quantum information[5]—our gate significantly expands the potential of single-spatial-mode, fiber-optic-based QIP. More generally, our Bayesian characterization approach provides further evidence of the potential of machine learning in analyzing quantum systems, particularly for extracting information within measurements which traditional methods overlook.

METHODS

Gate design

The optimization approach for designing quantum frequency gates using a series of EOMs/PSs was first proposed in ref. [26], and adopted to experimentally demonstrate a single-photon gate in ref. [14]. In this work, we follow the same procedures, utilizing the MATLAB Optimization Toolbox to search for an optimal set of phases for a particular EOM/PS sequence, constraining fidelity $\mathcal{F} \geq 0.9999$ and maximizing the success probability $\mathcal{P}$ for the two-photon state transformation matrix W. Compared to single-qubit gates, where only one frequency scale appears (the spacing between the two computational bins), two-qubit gates provide a much richer parameter space; namely, the placement of the four computational modes relative to each other can have a profound impact on the EOM/PS complexity needed to realize a specific operation. We have performed a thorough—though non-exhaustive—search over these possible mode placement combinations in each round of optimization. In general, we are guided by the intuition to spectrally isolate control mode 0 (C_0) while packing control mode 1 (C_1) close to both target modes.

For reference, the ideal CNOT matrix is

$$U_{ideal} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \tag{6}$$

against which we compare the numerically obtained two-photon matrix W (a function of V) via Eq. (2). The optimal solution we found for the CNOT gate using a 2EOM/1PS circuit is presented in Fig. 5, with $\mathcal{F} = 0.9999$, $\mathcal{P} = 0.0445$, and modes $\{C_0, C_1, T_0, T_1\}$ at frequency bins $\{0, 6, 7, 8\}$, respectively. The temporal phase modulation on both EOMs are simply π-phase-shifted sinewaves, and combined with the spectral phase modulation imparted by the PS, the corresponding mode transformation matrix V is numerically calculated as:

$$V = [r_{nn'} \measuredangle \phi_{nn'}] = \begin{bmatrix} 0.4407\measuredangle - 2.5976 & 0.0022\measuredangle 0.2103 & 0.0026\measuredangle 1.2938 & 0.0010\measuredangle - 2.0353 \\ 0.0022\measuredangle 0.2104 & 0.4343\measuredangle - 2.6045 & 0.4596\measuredangle - 1.5754 & 0.4549\measuredangle 1.5710 \\ 0.0026\measuredangle 1.2939 & 0.4596\measuredangle - 1.5754 & 0.4830\measuredangle 2.5973 & 0.0030\measuredangle - 2.8778 \\ 0.0010\measuredangle - 2.0352 & 0.4549\measuredangle 1.5710 & 0.0030\measuredangle - 2.8779 & 0.4783\measuredangle 2.5979 \end{bmatrix}, \tag{7}$$

using the phasor shorthand $r_{nn'} \measuredangle \phi_{nn'} \equiv r_{nn'} e^{i\phi_{nn'}}$. This provides a reference to compare the experimental mode transformations below, obtained either by coherent state characterization or BME.

Coherent state measurements

To investigate the performance of a linear-optical multiport, ref. [27] provides an efficient characterization method utilizing only coherent states as sources and power measurements at the output. We follow similar procedures (adopted in ref. [14] for single-qubit frequency gates) by probing our frequency multiport with an electro-optic frequency comb, and measuring the output spectrum for different input frequency super-positions. We first send a continuous-wave laser with center frequency $\Omega_{01} = 2\pi \times 193.550$ THz (see Fig. 2a) into an additional EOM modulated at 25 GHz to create ~10 comb lines, and we utilize a subsequent pulse shaper to prepare specific input states. To obtain the modulus of every matrix element in the four columns of V, each time we send in only one input mode from the set $\{C_0, C_1, T_0, T_1\}$ and measure the spectrum at the output of the gate, collecting all the output modes with power levels within 60 dB of the maximum. This allows us to retrieve the amplitudes $r_{nn'}$. Then by sending in two lines and scanning their relative input phase, we can map out the V-matrix phases $\phi_{nn'}$, where we compute all unknown values relative to phase values we are free by physical considerations to define a priori.[27] We perform five identical measurements of V in order to estimate uncertainty; following are the resulting amplitudes and phases, with each number averaged individually over the five successive, independent matrix acquisitions:

$$[r_{nn'}] = \begin{bmatrix} 0.428 \pm 0.008 & 0.0030 \pm 0.0003 & 0.0027 \pm 0.0001 & 0.0017 \pm 0.0001 \\ 0.0031 \pm 0.0001 & 0.427 \pm 0.001 & 0.451 \pm 0.002 & 0.451 \pm 0.002 \\ 0.0028 \pm 0.0002 & 0.465 \pm 0.005 & 0.478 \pm 0.003 & 0.041 \pm 0.003 \\ 0.0018 \pm 0.0003 & 0.458 \pm 0.002 & 0.036 \pm 0.004 & 0.499 \pm 0.006 \end{bmatrix} \tag{8}$$

$$[\phi_{nn'}] = \begin{bmatrix} -2.5976 \pm 0 & \dots & \dots & \dots \\ \dots & -2.6045 \pm 0 & -1.5754 \pm 0 & 1.5710 \pm 0 \\ \dots & -1.5754 \pm 0 & 2.621 \pm 0.002 & -2.89 \pm 0.05 \\ \dots & 1.5710 \pm 0 & -2.7 \pm 0.1 & 2.631 \pm 0.006 \end{bmatrix}. \tag{9}$$

The phase values with ± 0 uncertainty are those we could fix to the theoretical prediction [Eq. (7)], found by the optimizer to yield high CNOT fidelity. Because the coupling between mode C_0 and $\{C_1, T_0, T_1\}$ is too weak, we could not extract meaningful phase estimates of the elements delineated by "..." in the $\phi_{nn'}$ matrix. However, we have confirmed that setting these phases to any set of random values impacts our calculation of the fidelity at only the fifth decimal place, so that it has no influence on our computed $\mathcal{F}_{inf} = 0.995 \pm 0.001$. From the retrieved amplitudes and phases, we find uncertainties for the eight large elements ($r_{nn'} > 0.4$) at the third significant digit, an indication of the high precision possible with this high-flux characterization method.

Parameter model

In order to make use of the observed data to estimate the key parameters of our quantum gate, we first derive a realistic model connecting the

underlying gate operation to photon counts, encapsulated in a likelihood function $P(\mathcal{D}|\beta)$, for the model parameters β given data $\mathcal{D}$ (proportional to the conditional probability of $\mathcal{D}$ given β). In our case, the set β contains not only the mode transformation matrix V, but also the pair generation probability μ and the system efficiencies η_A and η_B.

Initially, we focus on how the input quantum state propagates through the multiport—for the moment neglecting loss, which will be incorporated later. The total optical network (defined over countably infinite frequency bins) maps inputs $\hat{a}_n$ to outputs $\hat{b}_n$ according to

$$\hat{b}_n = \sum_{n'=-\infty}^{\infty} V_{nn'}\hat{a}_{n'}, \qquad (10)$$

with V unitary when considered over all modes. For a particular counting experiment, we take the prepared input state as

$$|\Psi\rangle = |1_u 1_v\rangle = \hat{a}_u^\dagger \hat{a}_v^\dagger |\text{vac}\rangle, \qquad (11)$$

where $u \neq v$. Specifying such a state relies on several assumptions. For one, it neglects contributions from other frequency-bin pairs, justified experimentally by the BFC shaper's ability to suppress adjacent frequency bins by >40 dB. Additionally, this state expression—and the multiport model in general—treats each frequency bin as a pure single mode. Experimentally, as a consequence of the pump laser's ~kHz linewidth (much narrower than our 1.8 GHz-thick bins), a given photon pair is highly frequency-entangled, containing substructure absent in the separable state of Eq. (11). While such hidden entanglement would markedly reduce, e.g., the purity of *heralded* frequency-bin photons, it does not degrade the correlations in the two-photon experiments we conduct here. The counts registered for a particular pair of bins do result from a continuum of photon pairs with slightly different frequency offsets, implying that the net result is the incoherent sum of partially distinguishable probability amplitudes. However, as all such frequency pair combinations under the same bin lineshapes undergo matching frequency operations, the net measurement result is identical to the case in which all bins are purely single mode, apart from an overall scaling constant (see discussion of frequency filtering below). Finally, Eq. (11) does not include higher-order pair generation (e.g., four, six, eight, etc., photon terms) explicitly. Incidentally, the ansatz we incorporate for accidental coincidences [see Eq. (17) below] ends up capturing the main effects of multiple photon pairs on our data in a simpler fashion.

We define $p_\mu(1_m 1_n)$ as the probability for one photon to be found in mode m and the other in mode n at the output (again assuming no loss). This is given by

$$p_\mu(1_m 1_n) = |V_{mu}V_{nv} + V_{mv}V_{nu}|^2 \qquad (n \neq m). \qquad (12)$$

When $n = m$ (two photons in the same mode), the probability is

$$p_\mu(2_m) = 2|V_{mu}V_{mv}|^2, \qquad (13)$$

with the factor of two a consequence of boson statistics. From these results, we can also compute the marginal probability for one-photon occupancy in a particular mode,

$$
\begin{aligned}
p_\mu(1_m) &= \sum_{\substack{n=-\infty \\ n\neq m}}^{\infty} |V_{mu}V_{nv} + V_{mv}V_{nu}|^2 \\
&= \sum_{n=-\infty}^{\infty} \left(|V_{mu}V_{nv} + V_{mv}V_{nu}|^2\right) - 4|V_{mu}V_{mv}|^2 \\
&= |V_{mu}|^2 + |V_{mv}|^2 - 4|V_{mu}V_{mv}|^2,
\end{aligned}
\qquad (14)
$$

with the last line following from the unitarity of V and the fact that $u \neq v$ in our input state.

We then map these fundamental "per-pair" probabilities to expected detection rates. For accounting purposes, we define all detection probabilities within a specific temporal frame τ, the time within which clicks on detector A (t_A) and B (t_B) are deemed coincident: $|t_A - t_B| < \tau$. Our stationary (continuous-wave pumped) source ensures that all such probabilities are equal in every length-τ time bin. With μ defined as the pair generation probability within such a frame, the marginal probabilities for single-detector clicks are

$$
\begin{aligned}
p_A &= \mu[\eta_A + (1-\eta_A)\eta_A]p_\mu(2_m) + \mu\eta_A p_\mu(1_m) + d_A \\
p_B &= \mu[\eta_B + (1-\eta_B)\eta_B]p_\mu(2_n) + \mu\eta_B p_\mu(1_n) + d_B
\end{aligned}
\qquad (15)
$$

for detector A monitoring frequency bin m and B frequency bin n. The probabilities d_A and d_B represent the dark (or more generally, background)

count probabilities; we measure these independently and take them as fixed at $d_A = 9.60 \times 10^{-7}$ and $d_B = 7.77 \times 10^{-7}$, corresponding to dark count rates of 640 Hz and 518 Hz, respectively. The efficiencies η_A and η_B include all loss effects through the system, from generation in the crystal to photon detection; we assume them to be mode-independent— validated by the relatively small bandwidth comprising all modes of interest (~500 GHz)—yet they can vary by the different relative efficiencies of our superconducting nanowire detectors. And while spectral filtering per se does not modify these general considerations, the multimode frequency substructure (mentioned above), coupled with the Lorentzian linewidth profile of the etalon, introduces an effective transmission given by the average over all frequency offsets—we believe this contributes to lower overall η_A and η_B retrieved in BME. Next, we make use of the fact that the system efficiencies $\eta_A, \eta_B \ll 1$. Plugging in Eqs. (13) and (14), we obtain

$$
\begin{aligned}
p_A &= \mu\eta_A\left(|V_{mu}|^2 + |V_{mv}|^2\right) + d_A \\
p_B &= \mu\eta_B\left(|V_{nu}|^2 + |V_{nv}|^2\right) + d_B.
\end{aligned}
\qquad (16)
$$

The simple addition of pair and dark-count contributions is justified in our case by their small values (~10^{-6}), so that there is no concern for p_A or p_B approaching or exceeding 1 in the numerical analysis below.

To establish the probability for a coincidence between detectors A and B in our model, we make a sharp distinction between two types of events: (i) correlated coincidences, deriving from two photons of the same pair; and (ii) accidental coincidences, in which two random clicks (from at least one dark count, or photons from two different pairs) overlap within the resolving time τ. We note that, in principle, such a distinction is not necessary: it should be possible to derive a completely ab initio model for coincidences, with an input density matrix including higher-order pair generation effects, and positive-operator valued measures (POVMs) incorporating dark count noise. However, our approach proves much simpler, requiring fewer parameters while still satisfying conceptual demands.

For event (i), the click probability follows from multiplying the per-pair probability $p_\mu(1_m 1_n)$ by $\mu\eta_A\eta_B$, so that $p_{AB}^{(i)} = \mu\eta_A\eta_B|V_{mu}V_{nv} + V_{mv}V_{nu}|^2$, which assumes that τ is sufficiently large to integrate over the full two-photon correlation time. Regarding event (ii), in general the rate of accidental coincidences between two independent detectors is given by a product of the rates of the two detectors individually: $R_{AB}^{(ii)} = 2\tau R_A R_B$,[28,29] where the factor of two follows from the fact that—under our definition of τ—all events such that $(t_A - t_B) \in (-\tau, \tau)$ register as coincidences. Making the connection $p_j = \tau R_j$ then allows us to write $p_{AB}^{(ii)} = 2p_A p_B$, so that the total coincidence probability becomes

$$
\begin{aligned}
p_{AB} &= p_{AB}^{(i)} + p_{AB}^{(ii)} \\
&= \mu\eta_A\eta_B|V_{mu}V_{nv} + V_{mv}V_{nu}|^2 + 2p_A p_B,
\end{aligned}
\qquad (17)
$$

with p_A and p_B defined as in Eq. (16). Expanding $2p_A p_B$, the expected noise sources appear naturally: a μ^2 term reflects clicks from two different pairs, while μd_A and μd_B terms give coincidences from a photon and dark count. In this way, we can recover noise effects otherwise absent in the physical model, via what can be called an "accidentals correction" term $2p_A p_B$. Finally, we emphasize that the accuracy of Eq. (17) relies again on the relative order of magnitudes of the probabilities involved: $p_{AB}^{(i)} \sim 10^{-10}$, so that the differences between alternative forms one could conceivably argue for—such as $p_B \rightarrow p_B - p_{AB}^{(i)}$, to help ensure that singles counts from correlated coincidences do not also count toward accidental probabilities—become numerically inconsequential.

Finally, with these probabilities established, we can write the likelihood using a multinomial distribution for all event types. Over the course of a single measurement of duration T, we experience $M = T/\tau$ total frames, in which we can register one of the four mutually exclusive outcomes: click on A only, click on B only, coincidence, or no clicks. The likelihood for the specific input/output mode configuration (defined by the mode numbers $uv \rightarrow mn$) is

$$
\begin{aligned}
P(\mathcal{D}_{uv}^{mn}|\beta) &= (p_A - p_{AB})^{N_A - N_{AB}}(p_B - p_{AB})^{N_B - N_{AB}} \\
&\quad \times p_{AB}^{N_{AB}}(1 - p_A - p_B + p_{AB})^{M - N_A - N_B + N_{AB}},
\end{aligned}
\qquad (18)
$$

where we emphasize that both the dataset $\mathcal{D}_{uv}^{mn} = \{N_A, N_B, N_{AB}\}$ and probabilities $\{p_A, p_B, p_{AB}\}$ themselves depend on the mode configuration $uvmn$. The total likelihood follows by multiplying out all 16 individual

combinations

$$P(\mathcal{D}|\beta) = \prod_{\substack{u, m \in \{C_0, C_1\} \\ v, n \in \{T_0, T_1\}}} P(\mathcal{D}_{uv}^{mn}|\beta), \tag{19}$$

where the modes $\{C_0, C_1, T_0, T_1\}$ are as defined in the main text. (We also neglect unimportant scaling factors which do not depend on the parameters β.) This likelihood forms the basis for estimating the parameters $\beta = \{V, \mu, \eta_A, \eta_B\}$ from the dataset $\mathcal{D} = \cup \mathcal{D}_{uv}^{mn}$.

Bayesian machine learning

To estimate these values along with their uncertainties, we make use of Bayes' rule for the posterior probability distribution

$$P(\beta|\mathcal{D}) = \frac{1}{\mathcal{Z}} P(\mathcal{D}|\beta) P(\beta), \tag{20}$$

with $\mathcal{Z} = \int d\beta \, P(\mathcal{D}|\beta) P(\beta)$ the (undetermined) normalizing factor. $P(\beta)$ represents the prior probability distribution for the parameters. We take $P(\beta)$ as uniform over (0, 1) for each of μ, η_A, and η_B; uniform over $(0, 2\pi)$ for all phases $\phi_{nn'} = \arg V_{nn'}$ which are not taken as fixed $\{\phi_{C_0C_0}, \phi_{C_1C_1}, \phi_{C_1T_0}, \phi_{C_1T_1}, \phi_{T_0C_1}, \phi_{T_1C_1}\}$; and uniform for all squared moduli $r_{nn'}^2$, subject to the constraint $\sum_{nn'} r_{nn'}^2 = 1.6558$ from Eq. (7). This uninformative prior allows the estimates to be fully determined by the counting data itself.

Due to the complexity of integrating Eq. (20) over our parameter space, we employ slice sampling[37] and retrieve 4096 samples of all 28 parameters from the unnormalized $P(\mathcal{D}|\beta) P(\beta)$. We use best guesses of all parameters as the starting point to enable convergence, invoking a burn-in period and thinning until stationarity is achieved. At each sample of β, we can compute any quantity of interest, and use the statistics over all samples to produce the mean and standard deviation. Specifically, we find

$$\mu = 0.024 \pm 0.002 \tag{21}$$

$$\eta_A = (3.5 \pm 0.3) \times 10^{-4} \tag{22}$$

$$\eta_B = (4.7 \pm 0.3) \times 10^{-4} \tag{23}$$

$$\mathcal{F}_{BME} = 0.91 \pm 0.01. \tag{24}$$

The retrieved pathway efficiencies are smaller by ~9 dB compared to our insertion loss alone, which we estimate to be ~25 dB from generation to detection. While we have fully characterized the insertion loss of the gate components themselves (12.9 dB in total: each EOM contributes ~2.8 dB; the pulse shaper, ~4.7 dB; and the remainder comes from polarization controllers and fiber patch cords), uncertainties remain in the state preparation and measurement components, such as the breakdown of loss inside the fiber-pigtailed photon source, as well as questions of how strongly the spectrally varying transmission of the etalon reduces its effective transmission from its peak value. Otherwise, the retrieved μ and fidelity match predictions. Even though $\mathcal{F}_{BME}$ is smaller and has higher uncertainty than the classically inferred $\mathcal{F}_{inf}$, the fact it still exceeds 90% with fairly sparse measurements is strong confirmation of excellent performance, particularly in light of the uninformative prior, which permits high fidelity only based on the strength of the observed data.

We also compute the mean and standard deviation for all elements of the retrieved transformation V, for both the magnitude and phase:

$$[r_{nn'}] = \begin{bmatrix} 0.452 \pm 0.005 & 0.124 \pm 0.009 & 0.06 \pm 0.01 & 0.02 \pm 0.02 \\ 0.06 \pm 0.03 & 0.465 \pm 0.008 & 0.475 \pm 0.006 & 0.411 \pm 0.006 \\ 0.04 \pm 0.01 & 0.463 \pm 0.005 & 0.470 \pm 0.005 & 0.03 \pm 0.01 \\ 0.028 \pm 0.009 & 0.455 \pm 0.005 & 0.02 \pm 0.01 & 0.413 \pm 0.005 \end{bmatrix} \tag{25}$$

$$[\phi_{nn'}] = \begin{bmatrix} -2.5976 \pm 0 & -2.8 \pm 0.2 & 1.3 \pm 0.1 & -2.01 \pm 0.09 \\ 0.30 \pm 0.09 & -2.6045 \pm 0 & -1.5754 \pm 0 & 1.5710 \pm 0 \\ 1.35 \pm 0.09 & -1.5754 \pm 0 & 2.6 \pm 0.1 & 0.7 \pm 0.2 \\ -2.0 \pm 0.1 & 1.5710 \pm 0 & 0.3 \pm 0.1 & 2.5 \pm 0.1 \end{bmatrix}. \tag{26}$$

As before, the phases with uncertainties ± 0 are those fixed prior to parameter retrieval. Comparing this result to the design [Eq. (7)] and coherent-state-retrieved matrix [Eqs. (8) and (9)], the most significant mismatch occurs for the element in row 1, column 2 (the coupling from mode C_1 to C_0). At 0.124, this value is significantly larger than designed, and contributes to the higher error for the cases $|C_1T_0\rangle \to |C_0T_0\rangle$ and $|C_1T_1\rangle \to |C_0T_1\rangle$ in Fig. 4. While the source of this error is still uncertain, experimentally we did observe extraneous counts on detector A during these integration times, beyond the theoretical prediction. Bayesian retrieval succeeds in finding matrix elements to account for this observation, as intended.

DATA AVAILABILITY

Data and analysis code used in this study are available from the corresponding authors on request.

AUTHOR CONTRIBUTIONS

H.-H.L. led the experiments and data integration, contributed to the theoretical analysis, and assisted in the writing. J.M.L. discovered the CNOT solution, developed the parameter estimation model, and led writing of the paper. B.P.W. set up the time-tagging electronics and guided the estimation and sampling algorithms. P.I. initiated CNOT gate development and design. N.A.P. performed detector setup and assisted with data analysis. A.M.W. supervised the work at Purdue and assisted with system design and analysis. P.L. initiated the concepts, managed the theoretical analysis, and supervised the project. All authors reviewed the results and contributed to the manuscript.

ADDITIONAL INFORMATION

Competing interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

REFERENCES

1. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2000).
2. Knill, E., Laflamme, R. & Milburn, G. J. A scheme for efficient quantum computation with linear optics. *Nature* **409**, 46–52 (2001).
3. Kok, P. et al. Linear optical quantum computing with photonic qubits. *Rev. Mod. Phys.* **79**, 135–174 (2007).
4. Babazadeh, A. et al. High-dimensional single-photon quantum gates: concepts and experiments. *Phys. Rev. Lett.* **119**, 180510 (2017).
5. Humphreys, P. C. et al. Linear optical quantum computing in a single spatial mode. *Phys. Rev. Lett.* **111**, 150501 (2013).
6. Reddy, D. V. & Raymer, M. G. High-selectivity quantum pulse gating of photonic temporal modes using all-optical Ramsey interferometry. *Optica* **5**, 423–428 (2018).
7. Ansari, V., Donohue, J. M., Brecht, B. & Silberhorn, C. Tailoring nonlinear processes for quantum optics with pulsed temporal-mode encodings. *Optica* **5**, 534–550 (2018).
8. Reimer, C. et al. Generation of multiphoton entangled quantum states by means of integrated frequency combs. *Science* **351**, 1176–1180 (2016).
9. Jaramillo-Villegas, J. A. et al. Persistent energy-time entanglement covering multiple resonances of an on-chip biphoton frequency comb. *Optica* **4**, 655–658 (2017).
10. Kues, M. et al. On-chip generation of high-dimensional entangled quantum states and their coherent control. *Nature* **546**, 622–626 (2017).
11. Imany, P. et al. 50-GHz-spaced comb of high-dimensional frequency-bin entangled photons from an on-chip silicon nitride microresonator. *Opt. Express* **26**, 1825–1840 (2018).
12. Kobayashi, T. et al. Frequency-domain Hong–Ou–Mandel interference. *Nat. Photonics* **10**, 441–444 (2016).
13. Clemmen, S., Farsi, A., Ramelow, S. & Gaeta, A. L. Ramsey interference with single photons. *Phys. Rev. Lett.* **117**, 223601 (2016).

14. Lu, H.-H. et al. Electro-optic frequency beam splitters and tritters for high-fidelity photonic quantum information processing. *Phys. Rev. Lett.* **120**, 030502 (2018).

15. Lu, H.-H. et al. Quantum interference and correlation control of frequency-bin qubits. *Optica* **5**, 1455–1460 (2018).

16. Pittman, T. B., Jacobs, B. C. & Frason, J. D. Probabilistic quantum logic operations using polarizing beam splitters. *Phys. Rev. A* **64**, 062311 (2001).

17. Pittman, T. B., Jacobs, B. C. & Franson, J. D. Demonstration of nondeterministic quantum logic operations using linear optical elements. *Phys. Rev. Lett.* **88**, 257902 (2002).

18. Ralph, T. C., Langford, N. K., Bell, T. B. & White, A. G. Linear optical controlled-NOT gate in the coincidence basis. *Phys. Rev. A* **65**, 062324 (2002).

19. Hofmann, H. F. & Takeuchi, S. Quantum phase gate for photonic qubits using only beam splitters and postselection. *Phys. Rev. A* **66**, 024308 (2002).

20. O'Brien, J. L., Pryde, G. J., White, A. G., Ralph, T. C. & Branning, D. Demonstration of an all-optical quantum controlled-NOT gate. *Nature* **426**, 264–267 (2003).

21. O'Brien, J. L. et al. Quantum process tomography of a controlled-NOT gate. *Phys. Rev. Lett.* **93**, 080502 (2004).

22. Crespi, A. et al. Integrated photonic quantum gates for polarization qubits. *Nat. Commun.* **2**, 566 (2011).

23. Carolan, J. et al. Universal linear optics. *Science* **349**, 711–716 (2015).

24. Lahini, Y., Steinbrecher, G. R., Bookatz, A. D. & Englund, D. Quantum logic using correlated one-dimensional quantum walks. *npj Quantum Inf.* **4**, 2 (2018).

25. Zeuner, J. et al. Integrated-optics heralded controlled-NOT gate for polarization-encoded qubits. *npj Quantum Inf.* **4**, 13 (2018).

26. Lukens, J. M. & Lougovski, P. Frequency-encoded photonic qubits for scalable quantum information processing. *Optica* **4**, 8–16 (2017).

27. Rahimi-Keshari, S. et al. Direct characterization of linear-optical networks. *Opt. Express* **21**, 13450–13458 (2013).

28. Eckart, C. & Shonka, F. R. Accidental coincidences in counter circuits. *Phys. Rev.* **53**, 752–756 (1938).

29. Pearson, B. J. & Jackson, D. P. A hands-on introduction to single photons and quantum mechanics for undergraduates. *Am. J. Phys.* **78**, 471–484 (2010).

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

31. Blume-Kohout, R. Optimal, reliable estimation of quantum states. *New J. Phys.* **12**, 043034 (2010).

32. Williams, B. P. & Lougovski, P. Quantum state estimation when qubits are lost: a no-data-left-behind approach. *New J. Phys.* **19**, 043003 (2017).

33. Poyatos, J. F., Cirac, J. I. & Zoller, P. Complete characterization of a quantum process: the two-bit quantum gate. *Phys. Rev. Lett.* **78**, 390–393 (1997).

34. Chuang, I. L. & Nielsen, M. A. Prescription for experimental determination of the dynamics of a quantum black box. *J. Mod. Opt.* **44**, 2455–2467 (1997).

35. Branderhorst, M. P. A., Nunn, J., Walmsley, I. A. & Kosut, R. L. Simplified quantum process tomography. *New J. Phys.* **11**, 115010 (2009).

36. Knill, E. et al. Randomized benchmarking of quantum gates. *Phys. Rev. A* **77**, 012307 (2008).

37. Neal, R. M. Slice sampling. *Ann. Stat.* **31**, 705–767 (2003).

38. Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213 (2014).

39. Politi, A., Matthews, J. C. & O'Brien, J. L. Shor's quantum factoring algorithm on a photonic chip. *Science* **325**, 1221 (2009).

40. AIM Photonics. *Process Design Kit.* http://www.aimphotonics.com/pdk/ (2018).

41. Wang, C. et al. Integrated lithium niobate electro-optic modulators operating at CMOS-compatible voltages. *Nature* **562**, 101–104 (2018).

42. Kieninger, C. et al. Ultra-high electro-optic activity demonstrated in a silicon-organic hybrid modulator. *Optica* **5**, 739–748 (2018).