

Extensible universal photonic quantum computing with nonlinearity

In the format provided by the
authors and unedited

CONTENTS

I. Brief Introduction to the Setup	2
II. Experimental setup: Quantum Light Unit	2
A. Squeezed Light Source	2
B. Local Oscillator and Displacements	5
C. Non-Gaussian Photon Number State	5
D. “Multi-Core” Large-Scale Linear Optical Unit	7
E. Time-Polarisation Hybrid Encoding	8
F. “Quantum Cache” and I/O Module	8
G. Delay-Line and Phase Locking	8
III. Experimental setup: Detectors	9
IV. Overview of Units for Extended Functionality	11
A. Inline Squeezing Unit	11
B. Nonlinear Unit: Kerr Gate via Photon Addition and Subtraction	12
V. Selected Applications of Clavina	16
A. Linear Optical Unit: Quantum Simulation and Sampling Tasks	16
B. Linear Optical Unit: Cluster State Generation	18
C. Quantum Optical Neural Network	19
VI. Non-Gaussian State Preparation	20
A. Cat State Breeding	20
B. GKP State Engineering	21
C. GKP Cluster States and Surface Code	26
VII. Bose-Hubbard Hamiltonian simulation	26
References	28

I. BRIEF INTRODUCTION TO THE SETUP

Photonic quantum processors are emerging as versatile platforms for scalable quantum information processing, thanks to their inherent compatibility with low-loss, room-temperature operations and well-established optical technologies. Here, we present Clavina, an extensible photonic quantum computing system consisting of six units: (1) a control unit (CU); (2) a quantum light unit (QLU); (3) a detection unit (DU); (4) a linear operation unit (LOU); (5) an inline squeezing unit (ISU); and (6) a nonlinear unit (NLU, realised by a measurement-induced Kerr module). The QLU generates squeezed, coherent, and heralded photon-number states. Specifically, squeezed states are produced by a low-loss optical parametric oscillator (OPO), while photon-number states are prepared near-deterministically by heralding the squeezed light using a “boosting” protocol. In addition, the squeezing level and amplitude can be fully tuned via a computer-controlled electro-optic modulator (EOM). Fig. S1 shows the detailed layout. Clavina is extensible in two primary ways: first, the LOU can be scaled to perform larger linear transformations by adding computational “cores”; second, diverse functional modules can be connected to the CU in a plug-and-play fashion and enabled or disabled on the fly.

In the current implementation, the LOU contains two computational cores. Each core consists of a pair of EOMs that serve as fundamental processing elements, implementing $SU(2)$ operations between adjacent time-bin modes. These cores can operate independently. For a given circuit depth, the addition of cores reduces the number of loops required within the setup. Time-bin encoding has been utilised in earlier studies [2, 6–9], demonstrating excellent scalability and programmability, particularly when equipped with high-transmittance and high-bandwidth EOMs. This approach also substantially reduces the hardware count. Alongside these cores, the system includes a phase-locked delay line for temporal control, ensuring precise synchronisation. Another crucial component is a 530-meter-long optical fibre system, which incorporates high-efficiency free-space-to-fibre couplers to optimise photon transmission. Finally, an input/output (I/O) module, controlled by two optical switches (OSs), facilitates the dynamic routing of optical signals throughout the setup.

Time-bin encoding naturally provides expandable functionality. By inserting an I/O module at any point along the CU, we can interface additional expansion units such as ISU, NLU, and other useful quantum gates like the CZ gate with the LOU. The control software determines, for each time bin, whether to route the signal through an expansion unit or bypass it, effectively turning the unit on or off. This method yields exceptional hardware resource efficiency, as these units can be reused throughout a single execution of the circuit.

Finally, signals are detected by superconducting nanowire single-photon detectors (SNSPDs), photon-number-resolving detectors (PNRDs in the form of multi-pixel SNSPDs), or homodyne detectors (HDs). Clavina is versatile and can be reconfigured on demand to run linear circuit simulations, implement quantum optical neural networks, and generate error-correcting codes.

II. EXPERIMENTAL SETUP: QUANTUM LIGHT UNIT

A. Squeezed Light Source

Squeezed states are a quantum state of light in which the uncertainty of one quadrature (e.g., amplitude or phase) is reduced below the vacuum noise limit at the expense of increased uncertainty in the conjugate quadrature. This phenomenon adheres to the Heisenberg uncertainty principle, ensuring that the product of the uncertainties remains bounded by the quantum limit. Unlike classical light sources, such as coherent and thermal states, which exhibit symmetrical noise distributions at or above the shot-noise limit, squeezed states allow for sub-shot-noise performance, making them valuable for high-precision applications.

Mathematically, the quadratures of a single-mode electromagnetic field can be defined as $\hat{X} = \frac{1}{\sqrt{2}}(\hat{a} + \hat{a}^\dagger)$ and $\hat{Y} = \frac{1}{\sqrt{2}i}(\hat{a} - \hat{a}^\dagger)$, where $\hat{a}$ and $\hat{a}^\dagger$ are the annihilation and creation operators, respectively. A minimum uncertainty state (such as the vacuum or a coherent state) has equal variances, $\Delta^2 X = \Delta^2 Y$, while a squeezed state satisfies $\Delta^2 X < \Delta^2 X_{\text{vac}}$ and $\Delta^2 Y > \Delta^2 Y_{\text{vac}}$ (or vice versa). The unitary squeezing operator is given by

$$\hat{S}(\zeta) = \exp \left[\frac{1}{2}(\zeta^* \hat{a}^2 - \zeta \hat{a}^{\dagger 2}) \right], \quad (1)$$

where $\zeta = r e^{i\phi}$ is the squeezing parameter, with r controlling the squeezing strength and ϕ determining the orientation of the squeezed quadrature.

The single-mode squeezed vacuum (SMSV) state is obtained by applying the squeezing operator to the vacuum




FIG. S1: **An illustration of the experimental setup for each unit.** **QLU:** A single-mode squeezed vacuum is generated by an OPO. Coherent states derived from the pump serve as the displacement field. Photon-number (Fock) states are prepared by heralding from a two-mode squeezed-vacuum source. See Section II. **LOU:** A fully programmable, dual-core linear interferometer that enables large-scale linear transformations. **Control unit:** The control stack (not fully shown in the figure) contains a phase-locking module, an approximately 520 m long-fibre delay as a cache, a synchronisation module, and host-computer control software. **Input/output (I/O) module:** Provides a common path mode to insert all required components. **Detection unit:** Three detector options are available: superconducting nanowire single-photon detectors (SNSPDs), photon-number-resolving detectors (PNRDs in the form of multi-pixel SNSPDs), and homodyne detectors (HDs). **Inline squeezing unit:** Realised by mixing the input with a highly squeezed ancilla, followed by compensating displacements. The inline squeezing level is set by the EOM. **Nonlinear unit:** A measurement-induced Kerr gate provides optical nonlinearity, and its strength is tunable via the EOM. More details follow in Sections III, IV, and V. This extensible design integrates pre-built modules on a common platform, simplifying scaling, expanding functionality, and tailoring the system to diverse applications. Additional abbreviations not defined in the main text are: periodically poled potassium titanyl phosphate (ppKTP), half-wave plate (HWP), quarter-wave plate (QWP), beam splitter (BS), polarising beam splitter (PBS), right-angled prism mirror (RAP), and piezoelectric transducer mirror (PZT).

state, yielding [19]:

$$|\text{SMSV}\rangle = \frac{1}{\sqrt{\cosh r}} \sum_{n=0}^{\infty} (-e^{i\phi} \tanh r)^n \frac{\sqrt{(2n)!}}{2^n n!} |2n\rangle. \quad (2)$$

This state contains only even photon-number components ($|2n\rangle$), as photons are generated in pairs. Assuming $\phi = 0$, the quadrature variances for the SMSV state are $\Delta^2 X = \frac{1}{2}e^{-2r}$ and $\Delta^2 Y = \frac{1}{2}e^{2r}$. For $r > 0$, one quadrature is thus squeezed below the vacuum level, while the conjugate quadrature is anti-squeezed. **Squeezed State Generated by**

Optical Parametric Oscillator

Squeezed-vacuum states constitute the fundamental resource underlying every continuous-variable (CV) quantum protocol considered in this work. Their faithful generation relies on a low-loss OPO operated in a stabilised configuration. We therefore begin by outlining the OPO cavity design—its geometry, finesse, and phase-matching bandwidth—before detailing the preparation and delivery of both the synchronised pump-pulse sequence and the independent Pound-Drever-Hall (PDH) locking beam.

With this setup, we produce a SMSV pulse train at a 40 MHz repetition rate, achieving approximately 10 dB of squeezing. The squeezed light is routed to the LOU via the I/O module (see Section III.C for details). The pump consists of 40 MHz rectangular pulses of duration $T_p = 3.5$ ns at 775 nm. As illustrated in Fig. S2, these pulses originate from a continuous-wave (CW) laser centred at 1550 nm. A high-speed fibre electro-optic intensity modulator, driven by an arbitrary-waveform generator (AWG), carves the continuous output into pulses. The pulses are subsequently amplified and frequency-doubled to 775 nm by second-harmonic generation (SHG) in a periodically poled lithium niobate (PPLN) waveguide. To counteract slow drifts in pump power over time, an acousto-optic modulator (AOM) actively stabilises the pump intensity [2], using feedback from a photodiode (PD2) that monitors the transmitted pump.

The resulting train of 775 nm pump pulses is directed into a doubly resonant, semi-monolithic (hemilithic) OPO cavity containing a 10-mm-long, type-0 phase-matched periodically poled KTP (ppKTP) crystal as the nonlinear medium [1, 2]. One crystal end face serves as a dielectric mirror, while the opposite end is terminated by a piezoelectric transducer (PZT)-mounted output coupler. This configuration ensures that the cavity resonates at the 1550 nm wavelength, with the crystal temperature adjusted to maintain simultaneous dual-wavelength resonance. To achieve




FIG. S2: **Squeezed light generated from the OPO.** A continuous-wave laser at 1550 nm is split into two paths. *Lower path:* An intensity modulator carves ~ 3.5 ns pulses that are amplified, frequency-doubled to 775 nm, and used to pump a hemilithic ppKTP OPO. The OPO delivers a 40 MHz train of SMSV pulses. A dichroic mirror (DM) separates the squeezed output from residual pump light, while photodiodes PD1 and PD2 provide feedback for pump-power and cavity-length stabilisation. Additionally, a small portion of the original laser forms the local oscillator (LO) used in homodyne detection and for the displacement operation; further details are provided in the following section. *Upper branch:* Phase-modulated sidebands are filtered by an etalon, up-converted to 775 nm, and serve as the PDH lock beam.

double resonance and near-single-mode operation, carefully selected coatings are applied to both the crystal surfaces and the output coupler, as shown in Table I. Under these conditions, we calculate the lifetime of photons in the cavity to be $\tau_c \simeq 1.44$ ns; thus, $T_p > 2\tau_c$, satisfying the single-temporal-mode condition. By tuning the average power of the pump light from zero to approximately 35 mW, we obtain ~ 10 dB of squeezing. Finally, a mode-cleaner cavity filters out additional non-degenerate two-mode squeezed states and the train of SMSS emerging from the OPO is separated from any residual pump light by a dichroic mirror (DM), coupled into a single-mode fibre, and directed toward the subsequent measurement apparatus. An AR-coated GRIN lens is used for this fibre coupling, achieving an efficiency of approximately 93%.

Simultaneously, we employ an “OPO lock light” beam to maintain the OPO cavity resonances via the PDH technique [2]. This beam differs in wavelength from the squeezed pulses, is orthogonally polarised relative to the pump, and is generated independently of the pump itself. To generate it, we first send a fraction of the main laser through a phase modulator at approximately 7 GHz, then use an etalon to isolate only the +1 sideband. After amplification, we add 80 MHz sidebands using a second phase modulator and up-convert the light from roughly 1550 nm to 775 nm

in a temperature-controlled MgO:PPLN waveguide. Finally, the light transmitted through the OPO is detected by PD1 to produce the PDH error signal, which maintains the cavity lock.

TABLE I: Key OPO cavity parameters for 40 MHz, 3.5 ns pump pulses

Parameters	Values
<i>—Crystal Geometry and coating—</i>	
ppKTP crystal length (L)	10 mm
Radius of curvature of input surface (RoC1)	12 mm
High reflection coating for signal (HRs)	99.9 %
High reflection coating for pump (HRp)	95.0 %
Anti-reflection coating at exit surface (AR)	< 0.2 %
<i>—Piezo-Mounted Mirror Geometry and coating—</i>	
Radius of curvature of piezo-mounted mirror (RoC2)	25 mm
Transmission @ 1550 nm (T)	15 % (85 % R)
Transmission @ 775 nm (T_p)	5 % (95 % R_p)
Air gap (D)	18 mm
<i>—Other parameters—</i>	
Internal loss (L_{int})	$\sim 0.5\%$
Coupling efficiency to single-mode fibre	$\sim 93\%$

B. Local Oscillator and Displacements

In quantum optics, a coherent state plays a dual role: a bright beam acts as the local oscillator (LO) in balanced homodyne detection, while the same field—strongly attenuated—implements the displacement operation that shifts an optical mode in phase space. The latter is one of the three elementary Gaussian unitaries and is written as

$$\hat{D}(\alpha) = \exp(\alpha \hat{a}^\dagger - \alpha^* \hat{a}), \quad (3)$$

where $\alpha \in \mathbb{C}$ is the complex displacement amplitude, and $\hat{a}$ ($\hat{a}^\dagger$) is the annihilation (creation) operator. Experimentally, $\hat{D}(\alpha)$ is implemented by interfering the signal with a weak coherent state $|\alpha\rangle$ on a highly unbalanced beam splitter; only a small fraction of the coherent state is injected into the signal path, whereas the strong reflected portion maintains phase stability and can be reused as the LO.

Fig. S1 outlines the scheme used in this work. A CW laser at 1550 nm is gated by an EOM to generate a 40 MHz train of 3.5 ns rectangular pulses. After the EOM, the temporal width and spectrum of these pulses match those of the SMSV emerging from the OPO. Consequently, the same source, after appropriate attenuation and phase control, serves (i) as the LO for balanced homodyne detection, and (ii) as the displacement beam employed in downstream Gaussian operations. The relative LO phase is tuned with a piezo-driven fibre stretcher, and the pulse energy is actively stabilised by a variable optical attenuator to guarantee reproducible displacement amplitudes throughout the experiment.

C. Non-Gaussian Photon Number State

Photon-number (Fock) states are quantum states of light whose photon number is exactly defined. They are eigenstates of the photon-number operator $\hat{n} = \hat{a}^\dagger \hat{a}$ and satisfy

$$\hat{n} |n\rangle = n |n\rangle, \quad (4)$$

where $n \in \mathbb{N}$ is the photon number. In contrast with coherent or squeezed states, Fock states are *non-Gaussian*: their photon-number distribution is a delta function centred at n , and their Wigner functions exhibit pronounced negativity—an unmistakable signature of non-classicality.

To generate such states near-deterministically, we employ a “boosting” protocol, sketched in Fig. S3. A pulsed OPO first produces a 40 MHz train of SMSV pulses that acts as the resource. These pulses then pass through an




FIG. S3: **Near-deterministic Fock state generation.** A train of SMSV pulses from the OPO is first converted into temporally multiplexed TMSV states, which enable heralded photon-number generation. The signal arm is stored in a fibre delay line while the idler detection triggers electronic feed-forward that drives a high-speed EOM. This EOM gates the selected time bin and routes the heralded pulse—now projected onto the desired Fock state—to the output port, while all unheralded bins are discarded. Because a successful herald can occur at any position in the sequence, a second, loop-based delay line with an additional fast EOM buffers the chosen pulse to a fixed release time, ensuring that the final output states emerge at uniform intervals.

EOM, and their polarisation state changes from horizontal to vertical every other pulse. The pulses subsequently propagate through an interferometric multiplexing module consisting of a pair of polarising beam splitters (PBSs), a half-wave plate (HWP) set to 22.5° , and a delay compensator. This configuration converts the SMSV pulses into a two-mode squeezed vacuum (TMSV) state [18]. Consequently, a PNR measurement of the idler mode can herald a pure photon-number state in the signal mode.

Next, another PBS separates the signal and idler modes, and the idler mode is monitored by a PNRD. A detection heralds the subtraction of a photon from the squeezed pulses and triggers an electronic feed-forward system that labels the corresponding time bin as successful. This feed-forward drives a fast optical switch—implemented with a second EOM and a PBS—which routes the heralded pulse, now projected onto the target Fock state $|n\rangle$ (e.g., $|1\rangle$ for a single-photon event or $|2\rangle$ for a two-photon event), to the output channel.

All unheralded pulses are simply discarded. Each preparation cycle injects ten SMSV pulses into the multiplexing network. By multiplexing these ten time bins, the inherently probabilistic photon-subtraction process is “boosted”




FIG. S4: **Experimental realisation of the “dual-core” LOU.** **a**, Layout of the reconfigurable interferometric mesh for realising arbitrary linear transformations. Each Mach-Zehnder interferometer (MZI) achieves an $SU(2)$ operation. **b**, Odd-numbered (red) and even-numbered (blue) time bins carry horizontal (H) and vertical (V) polarisations, respectively. Adjacent bins (1 & 2, 3 & 4, ...) are driven independently by two high-bandwidth phase modulators, EOMa and EOMB. EOMa sets the relative phase φ , while EOMB—sandwiched between two half-wave plates (HWPs) oriented at 22.5° —tunes the polarisation rotation angle θ . Collectively, this arrangement is equivalent to an MZI. **c**, Laboratory implementation of the dual-core LOU. Core 1 (orange) and Core 2 (purple) each consist of a pair of EOMs, and a free-space Herriott-cell delay provides the required inter-core coupling, which shifts the odd and even time bins in opposite directions. This scheme effectively realises the Clements mesh shown in Fig. S4a. **d**, Equivalent path representation after one round trip. When both cores are activated, a single round-trip evolution can achieve two layers of linear transformation.

into a nearly deterministic source. In principle, increasing the number of multiplexed pulses further raises the success probability. With sufficiently many time bins, the desired Fock state can be selected with a probability arbitrarily close to one.

These non-Gaussian states are indispensable for fault-tolerant quantum computation, long-distance quantum communication, and quantum-enhanced sensing. Their sharply defined photon numbers allow for operations that Gaussian resources cannot provide, such as quantum error correction, entanglement distillation, and, in concert with Gaussian gates, universal quantum computation. In this work, we generate on-demand Fock states using the boosting protocol described above and use them as building blocks for more complex non-classical resources, e.g., Schrödinger cat states prepared by photon addition and Kerr-mediated gates implemented in later sections.

D. “Multi-Core” Large-Scale Linear Optical Unit

The large-scale linear transform circuit is shown in Fig. S4a. This circuit can be decomposed using Clements’ algorithm [11], with a Mach-Zehnder interferometer (MZI) as the primitive for arbitrary $SU(2)$ operations, as shown in the right inset of Fig. S4a. In our experiment, we realise the $SU(2)$ operation with two EOMs; custom anti-reflection coatings give each device a transmittance exceeding 99%. One EOM controls the phase shift φ between the horizontal and vertical polarisation components, while the other adjusts the rotation angle θ , as shown in Fig. S4b. By carefully tuning the voltages applied to the EOMs, we can realise any $SU(2)$ rotation on the Poincaré sphere of the polarisation state.

We also employ a temporal encoding scheme, where two neighbouring time-bins are merged into a single time-bin with orthogonal polarisations (horizontal (H) and vertical (V)). In this way, the polarisation degree of freedom is directly represented by the two neighbouring time bins. The $SU(2)$ operation, implemented by the EOMs, can thus

mix and phase-shift these two modes, effectively coupling adjacent time-bin states. Afterwards, we can separate the modes back into their respective time bins. The two cores are arranged in series along the free-space optical path, with a carefully calibrated delay line (transmittance $\gtrsim 0.985$) inserted between them. This delay line introduces a temporal offset equivalent to one time-bin duration, thereby synchronising the photon arrival times so that different pairs of time bins can interact sequentially. The corresponding experimental setup is shown in Fig. S4c.

In effect, the first core applies an initial $SU(2)$ transformation to a set of time-bin pairs, and after the delay, the second core further processes the newly adjacent time bins. This sequential arrangement allows the system to implement two layers of interferometric evolution within a single operational loop, as shown in Fig. S4d.

By implementing two layers of mode mixing per cycle, our dual-core architecture effectively halves the circuit length required for a given number of modes compared to conventional setups. This reduction in sequential processing elements significantly minimises optical losses, as photons traverse fewer modulators and free-space segments. Consequently, the overall efficiency of the system is improved, making it highly scalable and well-suited for complex quantum operations. This design not only streamlines the optical path but also enhances the fidelity and robustness of quantum state transformations in time-bin-encoded photonic networks.

Another important benefit of this design is that each $SU(2)$ core can be placed in a loop structure, allowing the same hardware to be reused multiple times. As the pulses circulate through the loop, we can modify φ and θ on each pass. This “time-multiplexed” approach significantly reduces the physical circuit footprint while maintaining flexibility and scalability for larger or more complex quantum operations.

E. Time-Polarisation Hybrid Encoding

Supplementary Fig. S5 illustrates the step-by-step temporal evolution of polarisation-encoded optical pulses through the time-multiplexed Gaussian boson sampling (GBS) setup. Each row of the time-series diagram corresponds to a successive stage in the optical path, with red and blue pulses denoting horizontal (H) and vertical (V) polarisations, respectively. At the source, an OPO produces a train of SMSV pulses with a fixed 25 ns spacing; an EOM then alternates the polarisation of each pulse between H (red) and V (blue). Next, a 7.5 m delay line introduces a 25 ns offset (one pulse period), causing each pulse to overlap in time with its immediate successor. These overlapping pulses are superposed at the first coupler stage (Core 1), enabling temporal-mode interference between neighbouring pulses.

Subsequently, a 15 m delay line imparts a 50 ns delay (two pulse periods), aligning pulses that were originally separated by two 25 ns intervals. At the second coupler (Core 2), these previously non-adjacent pulses are now interferometrically combined. A HWP set at 22.5° then applies a Hadamard-like rotation to each pulse’s polarisation, mixing the horizontal and vertical components. After this polarisation rotation, another 15 m delay line ensures that the pulses again arrive simultaneously for interference. Finally, a second HWP at 22.5° completes the polarisation processing, yielding the intended polarisation-encoded output state.

Throughout this scheme, the 25 ns pulse spacing serves as a fundamental time reference that synchronises the successive interference steps. The diagram provides a detailed visualisation of the precise timing and control strategy underlying the time-multiplexed GBS architecture.

F. “Quantum Cache” and I/O Module

Using a long optical fibre spool as a fixed-duration quantum memory or cache is a straightforward and widely adopted technique in photonic experiments. This approach introduces virtually no noise, apart from the negligible attenuation and dispersion of the fibre itself, largely preserving the encoded quantum states. It relies solely on standard telecom fibre and fast optical switches (OSs), making it ideal for short-term storage and synchronisation tasks in photonic quantum computing, as shown in Fig. S6. In our setup, a 520 m fibre acts as the quantum sequential access memory (QuSAM), providing more than $2.5 \mu\text{s}$ of storage time, which is sufficient for the current experiment. The overall coupling efficiency and transmittance of this 520-m-long fibre exceed 95%.

The I/O module is another important element in our experiment. Two OSs are embedded in an interferometer. Each OS can change the polarisation of light between H and V , thus determining whether the input light is sent to the QuSAM or to the detectors.

G. Delay-Line and Phase Locking

We implement a ~ 15 m Herriott cell as a free-space optical delay line, providing a path transmittance of approximately 0.987. This approach differs from our previous work in that a phase-locking scheme is now applied to ensure




FIG. S5: **Time-Series diagram.** Red and blue denote horizontal and vertical polarisations respectively.

stable temporal overlap between pulses originating from adjacent time bins. As shown in Fig. S7, a 1530 nm reference signal is injected to track and correct any phase drift, enabling the realisation of the Clements' architecture for high-fidelity interference of different polarisation components. For long data-collection runs, we employ the same phase-locking method as that used for the OPO source in the QLU.

After passing through the Herriott cell, a 220 m fibre delay line provides additional delay. This ensures that the final time bin of one cycle completes its operations before the first time bin of the subsequent cycle begins, effectively expanding the temporal multiplexing capacity. Free-space light is coupled into the single-mode fibre using a $4f$ beam-shaping system, yielding a measured total efficiency of approximately 95.4%. The remaining optical components within the delay and memory modules, including mirrors, PBSs, lenses, and AOMs, collectively exhibit a transmittance of approximately 97.5%. This combination of a low-loss Herriott cell, robust phase stabilisation, and efficient long-fibre memory underpins the high overall fidelity required for the temporal multiplexing and interference processes central to this experiment.

III. EXPERIMENTAL SETUP: DETECTORS

In this work, three varieties of detectors are used to detect quantum states of light, SNSPDs PNRDs and HDs. For the graph-based computational tasks in this work, we only require SNSPDs, as they demand only threshold detection, and the system losses reduce the average photon number per time bin to well below one, rendering photon-number discrimination superfluous. The SNSPDs used here achieve a detection efficiency of approximately 0.89, a level that ensures the reliable recording of single-photon events. Moreover, we separate orthogonal polarisation components within each time bin and direct them to two distinct SNSPDs, maintaining independent detection channels for each polarisation state. This arrangement effectively captures all relevant photonic information while preserving the advantages of SNSPDs—namely, high efficiency, fast response, and low noise—and obviates the need for more complex, photon-number-resolving instrumentation.

The HD efficiency is modelled as four independent losses whose efficiencies multiply to give η_{HD} : the optical




FIG. S6: Experimental setup of “QuSAM” and I/O modules.




FIG. S7: Details of the phase-locking setup. Phase-locking procedure for the delay-line module.

throughput η_{opt} , the spatial mode overlap (mode matching) η_{MM} , the photodiode quantum efficiency η_{PD} , and the electronic signal-to-noise factor η_{el} . Thus, $\eta_{\text{HD}} = \eta_{\text{opt}}\eta_{\text{MM}}\eta_{\text{PD}}\eta_{\text{el}}$.

All four factors were characterised independently. The optical path contains a single 99.9:0.1 beam splitter; the loss coming from all other components—such as lenses, collimation packages, and mirrors—is very small and can be neglected, giving $\eta_{\text{opt}} \approx 0.99$. The high Hong-Ou-Mandel visibility between the signal light and the LO provides $\eta_{\text{MM}} \approx 0.975$. The photodiodes provide $\eta_{\text{PD}} > 0.970$. With an ~ 12 mW local oscillator, the electronic signal-to-noise ratio exceeds 20 dB, which provides $\eta_{\text{el}} \approx 0.990$. Combining these values yields $\eta_{\text{HD}} = 0.990 \times 0.975 \times 0.970 \times 0.990 \simeq 0.927$, representing an overall homodyne efficiency of approximately 93%. This figure is used in the error budget and numerical simulations presented in the main text.




FIG. S8: **Phase-locking data.** Signal light is injected into the delay line from both directions, and each path must be phase-locked. **a**, Phase-locking data for signal 1. Photon counts recorded after the delay line in the QLU: the blue trace shows counts without phase locking, and the red trace shows counts with phase locking engaged. **b**, Phase-locking data for signal 2.

IV. OVERVIEW OF UNITS FOR EXTENDED FUNCTIONALITY

A. Inline Squeezing Unit

Our ISU employs a configuration in which the squeezing operation is applied directly along the optical path, thereby minimising optical losses and spurious mode coupling. As shown in Fig. S9, a PBS first combines the input state with an ancillary squeezed vacuum, prepared off-path by a dedicated optical parametric amplifier. The two beams co-propagate through an EOM, which imparts phase-dependent modulation on the combined field. An adjustable delay line is inserted after the EOM to match optical path lengths and enable precise synchronisation of the measurement and feed-forward stages.

To realise the squeezing operation, we split the beam again with a second PBS, directing part of the field towards a balanced homodyne detection module. This module consists of a beam splitter (BS) that interferes the signal with an LO of tunable phase, followed by high-efficiency SNSPDs or, alternatively, conventional balanced photodiodes, depending on the target measurement bandwidth. By measuring both quadratures of the field, we obtain real-time information used to condition the subsequent inline feed-forward electronics.

Following detection, electronic signals are processed and passed through amplitude and phase modulators placed in the main optical path. These modulators, driven by the measured quadratures, perform the required displacement and squeezing transformations in a single pass, thereby reducing complexity compared to iterative or cavity-based methods. This single-pass approach is particularly advantageous when high-speed feedback is required, since the timing constraints are relaxed relative to typical cavity lifetimes. By refining the squeezing interaction strength with the EOM, as well as the gains and phases in the modulators, we can tune both the magnitude and orientation of the squeezing ellipse.

The key benefit of the inline design is that the beam undergoes minimal redirection, resulting in lower coupling losses and a cleaner spatial mode. This not only improves the overall fidelity of the quantum state but also simplifies




FIG. S9: **Measurement-based single-mode ISU.** The input pulse is mixed on the first polarising beam splitter with an orthogonally polarised ancillary squeezed state. An EOM couples the two polarisations, after which a second PBS separates them. The signal passes through a short fibre delay line, while the ancilla is measured by HD. The homodyne outcome is processed in real time and fed forward to a pair of high-bandwidth EOMs that apply the required phase-space displacement to the delayed signal, yielding the squeezed output.

scaling to larger, multi-mode systems. Indeed, by cascading or parallelising this platform across multiple spatial or temporal modes, one can implement more complex Gaussian or even non-Gaussian transformations directly on the optical field. Such capability is instrumental for advanced protocols in CV quantum information processing, including cluster-state generation, error correction, and high-fidelity non-Gaussian state engineering.

B. Nonlinear Unit: Kerr Gate via Photon Addition and Subtraction

Implementing strong optical nonlinearity is a central challenge in photonic quantum computing. The annihilation ($\hat{a}$) or creation ($\hat{a}^\dagger$) operators shift the field up or down one Fock level. By coherently interfering the “add-then-subtract” and “subtract-then-add” paths, $A\hat{a}\hat{a}^\dagger + B\hat{a}^\dagger\hat{a} = (A+B)\hat{n} + A$, we can approximately realise an adjustable Kerr unitary evolution, $e^{ixt\hat{n}^2}$, with its strength determined by the ratio A/B [4, 5]. In our work, photon subtraction is realised by a low-reflectivity beam splitter followed by photon-number-resolving detection—the detection of a single photon heralds a successful photon subtraction. Photon addition, on the other hand, is realised by interference with a Fock state on a similar low-reflectivity beam splitter, where a successful operation is now heralded by a null detection. This measurement-induced route provides a probabilistic but practical method for embedding strong nonlinearities into linear optical platforms.

The probabilistic nature can be mitigated through offline preparation and boosting. Specifically, we propose preparing the measurement-induced Kerr gate offline in a repeat-until-success fashion and subsequently teleporting it onto the target state. This approach transfers the non-Gaussian operation onto the input mode and, given high-quality resource states and efficient detection, can achieve high-fidelity operation on arbitrary input states.

Beyond this boosting method, a two-level emitter such as a trapped atom or a semiconductor quantum dot (QD) may offer a promising route to deterministic photon addition [34, 35] and subtraction [36, 37]. Recent advances have delivered near-unity emitter-mode coupling [38] and picosecond-scale QD-cavity dynamics, potentially providing the fast, low-loss interfaces (excluding collection loss) required for effectively deterministic non-Gaussian operations. These capabilities are important for realising universal photonic quantum at the physical level, where Gaussian resources alone are insufficient. Notably, Clavina enables the interfacing of these platforms without a complete redesign of the architecture.

In this section, we present a detailed derivation of how a Kerr nonlinearity can be realised through photon addition and subtraction.

We start by considering a Kerr-type nonlinearity described by the Hamiltonian:

$$\hat{H}_{\text{Kerr}} = \hbar\kappa(\hat{a}^\dagger)^2\hat{a}^2 = \hbar\kappa\hat{n}(\hat{n}-1), \quad (5)$$

where $\hat{a}^\dagger\hat{a} = \hat{n}$ and $\hat{a}\hat{a}^\dagger = \hat{n} + 1$.

Under unitary time evolution, the Kerr operator transforms photon-number states as:

$$|n\rangle \xrightarrow{U_{\text{Kerr}}} e^{-i\Phi n(n-1)} |n\rangle, \quad (6)$$

where $|n\rangle$ denotes the photon-number state, and the effective Kerr strength is defined as $\Phi = \kappa t$.

For an arbitrary state $|\Psi_{\text{in}}\rangle = c_0|0\rangle + c_1|1\rangle + c_2|2\rangle$ within the $\{|0\rangle, |1\rangle, |2\rangle\}$ subspace, the Kerr gate transforms the state into $c_0|0\rangle + c_1|1\rangle + c_2e^{-2i\Phi}|2\rangle$. Clearly, the Kerr effect introduces a photon-number-dependent phase proportional to the square of the photon number associated with each basis state. However, realising this nonlinear interaction is extremely challenging in weak light fields. In Ref. [4], Costanzo *et al.* proposed a practical method to achieve a strong Kerr effect at the single-photon level by combining photon addition and subtraction operations.

The combination of photon addition and subtraction can achieve a photon-dependent phase transformation (accompanied by a certain gain) on the Fock basis, exhibiting a strong single-photon-level nonlinear effect in weak light fields; it is precisely this property that enables the effective simulation of the Kerr effect for finite photon-number states through an appropriate combination of these operations.

Following Ref. [4], we first define an operator $V(\hat{n}) = A\hat{a}\hat{a}^\dagger + B\hat{a}^\dagger\hat{a} = (A+B)\hat{n} + A$. For an arbitrary input state within the two-photon subspace,

$$|\Psi_{\text{in}}\rangle = c_0|0\rangle + c_1|1\rangle + c_2|2\rangle, \quad (7)$$

applying the operation $V(\hat{n})$ transforms $|\Psi_{\text{in}}\rangle$ to:

$$|\Psi_{\text{out}}\rangle = c_0V(0)|0\rangle + c_1V(1)|1\rangle + c_2V(2)|2\rangle. \quad (8)$$

The target output state after the Kerr gate is:

$$|\Psi_{\text{target}}\rangle = c_0|0\rangle + c_1|1\rangle + e^{-2i\Phi}c_2|2\rangle. \quad (9)$$

It can be seen that in this subspace, there is only a corresponding phase change applied to the two-photon component.

By comparing the resulting state with the target state, we can obtain the specific form of V by solving the system of equations:

$$\begin{cases} c_0V_0 = \lambda c_0 & \Rightarrow V_0 = \lambda, \\ c_1V_1 = \lambda c_1 & \Rightarrow V_1 = \lambda, \\ c_2V_2 = \lambda e^{-2i\Phi}c_2 & \Rightarrow V_2 = \lambda e^{-2i\Phi}. \end{cases} \quad (10)$$

This yields:

$$\begin{cases} V_0 = V_1, \\ V_0 = e^{-2i\Phi}V_2. \end{cases} \quad (11)$$

In this case, we cannot obtain an effective solution unless both A and B are set to zero.

If an ideal, noise-free linear amplifier with gain g (i.e., $\hat{a} \mapsto g\hat{a}$ without introducing excess noise) is available, the properly normalised target state can be expressed as:

$$|\Psi_{\text{target}}\rangle = \mathcal{N}(g) \left(c_0|0\rangle + g c_1|1\rangle + g^2 e^{-2i\Phi} c_2|2\rangle \right), \quad \mathcal{N}(g) = \left(|c_0|^2 + g^2 |c_1|^2 + g^4 |c_2|^2 \right)^{-1/2}. \quad (12)$$

By comparing the target state with the state after interaction with $V(\hat{n})$, defined as $V(\hat{n})|\Psi_{\text{in}}\rangle = c_0V_0|0\rangle + c_1V_1|1\rangle + c_2V_2|2\rangle$, we can solve the following system of equations:

$$\begin{cases} c_0V_0 = \lambda c_0, \\ c_1V_1 = \lambda c_1 g, \\ c_2V_2 = \lambda c_2 g^2 e^{-2i\Phi}, \end{cases} \quad \Rightarrow \quad \begin{cases} V_0 = \lambda, \\ V_1 = \lambda g, \\ V_2 = \lambda g^2 e^{-2i\Phi}. \end{cases} \quad (13)$$




FIG. S11: **Kerr effect realised by two-photon subtraction.** The half-wave plate (HWP) and quarter-wave plate (QWP) before the polarising beam splitter (PBS) are used to adjust B and C in operator V .

We can then find the relation between coefficients A and B :

$$\begin{cases} gA = 2A + B, \\ 3A + 2B = A g^2 e^{-2i\Phi}, \end{cases} \implies A = \frac{1}{g-2} B. \quad (15)$$

At this point, we can determine the gain factor g by the required Kerr intensity Φ , thus determining the relationship between A and B . According to the previous equations, we can find the relation between g and Φ :

$$2g - 1 = g^2 e^{-2i\Phi}. \quad (16)$$

Solving for g yields:

$$g = e^{2i\Phi} \left(1 \pm \sqrt{1 - e^{-2i\Phi}} \right) = |g(\Phi)| e^{i\theta(\Phi)}. \quad (17)$$

If we set $\Phi = \frac{\pi}{2}$, then $g = -\left(1 \pm \sqrt{1 - (-1)}\right) = -(1 \pm \sqrt{2})$. Therefore, in this case, the gain factor is $g = 1 + \sqrt{2}$ with a phase rotation $R_\pi = e^{-i\pi\hat{n}}$.

If we want to achieve this Kerr effect without noise-free amplification, we need the ability to add and subtract two photons. **Note that:**

$$\begin{aligned} (\hat{a}^\dagger)^2 \hat{a}^2 &= \hat{n}(\hat{n} - 1), \\ \hat{a}^2 (\hat{a}^\dagger)^2 &= (\hat{n} + 1)(\hat{n} + 2). \end{aligned} \quad (18)$$

For the state $|\Psi_{\text{in}}\rangle = c_0|0\rangle + c_1|1\rangle + c_2|2\rangle$, the two-photon addition and subtraction operations yield:

$$\begin{aligned} \hat{a}^2 (\hat{a}^\dagger)^2 |\Psi_{\text{in}}\rangle &= 2c_0|0\rangle + 6c_1|1\rangle + 12c_2|2\rangle, \\ (\hat{a}^\dagger)^2 \hat{a}^2 |\Psi_{\text{in}}\rangle &= 0 \cdot c_0|0\rangle + 0 \cdot c_1|1\rangle + 2c_2|2\rangle. \end{aligned} \quad (19)$$

Defining the V operator as:

$$V = A \hat{a}^2 (\hat{a}^\dagger)^2 + B (\hat{a}^\dagger)^2 \hat{a}^2 + C \hat{I}, \quad (20)$$

we find:

$$V|\Psi_{\text{in}}\rangle = (2A + C) c_0|0\rangle + (6A + C) c_1|1\rangle + (12A + 2B + C) c_2|2\rangle. \quad (21)$$

Comparing this to the target state $c_0|0\rangle + c_1|1\rangle + e^{-2i\Phi}c_2|2\rangle$, we obtain the system of equations:

$$\begin{cases} 2A + C = 1, \\ 6A + C = 1, \\ 12A + 2B + C = e^{-2i\Phi}, \end{cases} \quad (22)$$

and finally, we can solve for the coefficients:

$$A = 0, \quad C = 1, \quad B = \frac{e^{-2i\Phi} - 1}{2}. \quad (23)$$

In this case, we find that as long as a two-photon subtraction operation is available, the Kerr effect can be induced without requiring noise-free amplification, as shown in Fig. S11.

Deterministic single-photon subtraction has already reached a technologically credible regime on III-V QD platforms. Du *et al.* [20] showed that a charged InAs QD embedded in a high- Q , birefringent photonic-crystal cavity implements single-photon Raman routing: the first incident photon is coherently scattered into an orthogonally polarised cavity mode while projecting the QD into a dark state, thereby blocking all subsequent photons.

Deterministic single-photon addition, by contrast, remains a forward challenge. Jiang *et al.* [21] proposed a universal three-level emitter-cavity protocol in which the emitter is first driven to an excited state and then brought into resonance with the cavity for a sub-nanosecond window, releasing exactly one photon before decoupling. Theory predicts $> 99\%$ addition fidelity provided the cavity tuning window is faster than both the emitter lifetime (< 300 ps for InAs QDs) and the residual κ and γ^* . Current electro-optic and optomechanical cavity shifters operate at sub-GHz bandwidths, and multi-photon gain channels further degrade fidelity. Progress hinges on > 10 GHz Pockels-class modulators or GHz-bandwidth acousto-optic shutters integrated within the photonic-crystal platform, together with spin-selective three-level QD schemes that suppress unwanted amplification.

V. SELECTED APPLICATIONS OF CLAVINA

A. Linear Optical Unit: Quantum Simulation and Sampling Tasks

Our linear optical system is designed to simulate the Hamiltonian

$$\hat{H} = \sum_{i,j} \kappa_{ij} (\hat{a}_i^\dagger \hat{a}_j + \hat{a}_j^\dagger \hat{a}_i), \quad (24)$$

which describes photon hopping between modes. In our approach, each term $\hat{a}_i^\dagger \hat{a}_j + \hat{a}_j^\dagger \hat{a}_i$ is implemented using a beam splitter operation. When two optical modes interfere at a beam splitter, the resulting transformation mirrors the exchange of photons between those modes. By adjusting the beam splitter's coupling strength and phase, we can reproduce the effect of the coupling parameter κ_{ij} .

At the experimental level, our setup employs a reconfigurable network of beam splitters and phase shifters. Each beam splitter is realised as an MZI, where one arm contains a phase modulator and the other a polarisation modulator. The internal phase shifter within the MZI controls the effective splitting ratio, while an external phase shifter adjusts the relative phase between the modes. This combination allows us to finely tune the amplitude and phase of the mode coupling, effectively mapping the optical transformation to the target Hamiltonian.

The overall interference network is constructed by arranging these tunable couplers in sequence. By configuring the network so that the optical paths experience the desired sequence of beam splitter operations over a duration Δt , we approximate the evolution operator

$$\exp(-i\hat{H}\Delta t), \quad (25)$$

which governs the Hamiltonian dynamics. In this way, the continuous optical interference produced by the network directly emulates the unitary evolution dictated by $\hat{H}$. Precise control over the beam splitter transmissivities and phase shifts ensures that the simulation accurately reflects the intended coupling strengths and phase relationships among the modes.

By interfacing the LOU, with a dual-core configuration, with the CU, alongside the squeezer and SNSPDs, we conduct a 100-mode GBS experiment, which requires adjustable squeezing and a fully reconfigurable interferometer [6].




FIG. S12: **Linear transform circuit.** **a**, An $N \times N$ mesh of tunable Mach-Zehnder interferometers implements an arbitrary unitary transformation on the incoming modes, delivering the transformed outputs to the detectors. **b**, The experimentally measured output probability distribution for a 100-node random network (blue) closely matches the theoretical prediction (yellow), with a total variation distance of 0.0953. The participation coefficient inferred for each node in the 100-node network is shown in Fig. 2a of the main text.

The measured detection patterns agree well with the theoretical predictions (Fig. S12b), confirming stable, high-quality interference across all 100 modes.

A clique in a graph $G = (V, E)$ is a subset of vertices in which every pair is connected by an edge. A complete list of maximal cliques identifies the exact, tightly knit subgroups to which each vertex belongs. We can then compute each node's participation coefficient. Conceptually, this coefficient characterises whether a given node interacts exclusively within a single tightly bound group or distributes its connections across multiple groups. Treating every detected clique as a “group”, we determine two values for node i : the total number (or weight) of its links, k_i , and the subset of those links that fall within group s , k_{is} . The participation coefficient is then defined as $P_i = 1 - \sum_s (k_{is}/k_i)^2$. Thus, $P_i = 0$ when all of node i 's links lie inside a single clique, while $P_i \rightarrow 1$ as its links are distributed evenly across many cliques. Consequently, clique finding supplies a high-resolution community map, and the formula above distills that map into a per-node indicator of how provincial ($P_i \approx 0$) or global ($P_i \approx 1$) each vertex is.

Clavina can also readily time-multiplex squeezed-light pulses to generate scalable cluster states [29–31]. Figs. S13a and b confirm the presence of quantum correlations across 8,000 temporal modes. With a multi-core LOU that maintains low loss and high stability as the system scales, Clavina achieves an unprecedented scale and is perfectly positioned for preparing large-scale cluster states.

Encoding a Weighted Graph in a Gaussian Boson Sampling Device

We consider a weighted undirected graph $G = (V, E)$ with $|V| \approx 100$ nodes (vertices) and real-valued, non-negative edge weights. Our goal is to encode G into a GBS apparatus consisting of M input modes (one per node), each initialised in a squeezed-vacuum state. These squeezed states pass through a fully programmable time-bin linear interferometer implementing an $M \times M$ unitary transformation $\mathbf{U}$, followed by an array of SNSPDs.

Mathematically, let $\mathbf{W}$ denote the (weighted) adjacency matrix of G . Within the GBS framework, the two-photon correlation matrix is engineered to be proportional to $\mathbf{W}$, ensuring that the probability of detecting photons in a particular combination of output modes encodes the substructure of G . In the idealised pure-squeezing scenario (zero displacement, no thermal noise), the probability $P(S)$ of observing a detection event $S \subseteq \{1, 2, \dots, M\}$ is given by

$$P(S) \propto \text{Haf}(\mathbf{W}_S), \quad (26)$$

where $\mathbf{W}_S$ is the submatrix of $\mathbf{W}$ indexed by the modes in S , and $\text{Haf}(\cdot)$ denotes the hafnian. Since the hafnian is larger for denser subgraphs, the GBS device preferentially samples cliques, particularly those with large total edge weights.

Extraction of Maximum-Weight Cliques from GBS Samples

Although identifying maximum-weight cliques is an NP-hard problem in general, the GBS output statistically amplifies highly connected subgraphs, thus increasing the probability of sampling near-optimal or optimal solutions. We denote by S a raw sample: the set of output modes for which the detectors click. In practice, S may not necessarily form a perfect clique (some edges may be missing), so we apply a two-step post-processing algorithm (greedy shrinking

and local expansion).

Repeating many GBS samples and applying this two-step procedure yields a family of clique candidates. GBS probabilistically biases towards heavier cliques, therefore the true maximum-weighted clique often appears frequently among these candidates, making it easier to pinpoint the global optimum compared to uniform sampling.

Definition of Node Participation

To characterise how often a node contributes to the detected cliques, we define the notion of *node participation*. Let $\mathcal{C} = \{C_1, C_2, \dots, C_T\}$ be the set of cliques obtained from T GBS samples. For each node $i \in V$, we define the participation score p_i as:

$$p_i = \frac{1}{T} \sum_{j=1}^T \mathbb{I}(i \in C_j), \quad (27)$$

where the indicator function $\mathbb{I}(i \in C_j)$ is 1 if node i appears in the j th clique and 0 otherwise. A value of $p_i = 1$ implies that i is part of *every* sampled clique, while $p_i = 0$ means i was never found in any clique. We can also define a *weighted* participation score $p_i^{(w)}$ by weighting each clique by its total edge weight.

Computation of Participation and its Real-World Significance

Since post-processing yields many clique candidates, computing p_i or $p_i^{(w)}$ is straightforward: we record how frequently (or heavily) each node appears in the sampled cliques. High-participation nodes are those that appear consistently in dense subgraphs, indicating their central importance to the connectivity structure.

In drug discovery [6], for instance, nodes might represent molecular features or pharmacophore sites, and edges represent favourable interactions. High-participation nodes suggest features that consistently form part of strong binding configurations, pinpointing key interaction hotspots for medicinal chemistry. In finance, nodes may correspond to assets, and edges capture correlation strengths. A node with a large p_i participates in many highly correlated sub-portfolios, marking it as a systemic risk factor or a diversification bottleneck. Similarly, in social network analysis, high-participation individuals appear in multiple tightly knit communities, highlighting their role as influencers or information brokers.

Thus, the participation metric extends the GBS approach beyond merely finding the single best clique, enabling a richer analysis of structural importance and interaction patterns. Our results suggest that the GBS device, combined with efficient post-processing and the participation framework, offers a powerful tool for tackling NP-hard clique problems and generating deep insights into complex networked systems.

B. Linear Optical Unit: Cluster State Generation

Cluster States for Measurement-Based Quantum Computation (MBQC)

Cluster states are highly entangled multipartite quantum states serving as a resource for one-way quantum computing [25]. In MBQC, a computation proceeds by preparing a cluster state (with qubits or CV modes) and performing only single-particle measurements in adaptive bases. This scheme separates the creation of multi-qubit entanglement (done once, at the start) from the execution of quantum gates (implemented via measurements). It has proven especially attractive for photonic platforms, where deterministic two-photon gates are challenging.

General Techniques for Cluster State Generation

Early experiments in discrete-variable photonic systems combined entangled photon pairs on beam splitters, fusing small entangled states into larger clusters [26]. However, their inherently probabilistic nature limited the achievable size. In CV systems, squeezed-light sources and linear optical networks can deterministically generate cluster states [27, 28]. By mixing multiple squeezed modes on beam splitters, large-scale entanglement can emerge without the exponential overhead of post-selection. Techniques such as time-domain multiplexing reuse the same hardware to entangle many temporal modes, enabling thousands of correlated modes in a single apparatus.

Experimental Implementation

As demonstrated in Ref. [31], a CV cluster state with over 10,000 entangled modes can be created using this approach. Using two OPO squeezed-light sources, a continuous EPR-like entangled state is produced by interfering the modes on a 50:50 beam splitter (BS). An optical delay line is then employed, followed by a second BS, to link consecutive time slices. Each time bin of one beam is interfered with the subsequent time bin of the other beam, creating a one-dimensional “chain” of entangled modes. Verification is performed by measuring quadrature variances (nullifiers) below the shot-noise level across all modes, confirming large-scale entanglement. This approach illustrates how time-domain multiplexing can dramatically expand the size of cluster states in a practical, deterministic manner.




FIG. S13: **Schematic of cluster state entanglement.** Two synchronised squeezed-pulse trains first interfere at Core 1 (acting as a variable beam splitter) to form time-ordered EPR pairs. One rail is then sent through a Herriott-cell delay line (offsetting it by one time bin) before the two paths recombine at Core 2. This process stitches together the nearest- and next-nearest-neighbour links, completing the square-lattice cluster state at the output. **a** and **b**, Quadrature correlation plots indicate strong x - and p -quadrature correlations between modes A and B across the first 200 time-bin modes. The measurements below show nullifier variances consistently below the -3 dB inseparability threshold over 8,000 modes, confirming sustained multimode entanglement in both quadratures.

Applying a similar architecture in our work, Figs. S13a and b show the experimental results confirming quantum correlations across 8,000 temporal modes.

C. Quantum Optical Neural Network

Quantum neural networks (QNNs) fuse the trainable structure of classical deep learning with quantum superposition and entanglement, yielding models such as quantum perceptrons, convolutional layers, and Boltzmann machines that can be implemented on photonic hardware through measurement-induced, non-Gaussian operations.

Quantum optical neural networks (QONNs) use photonic circuits to implement neural-network-like computations on quantum states of light. The core idea is to alternate between linear optical transformations (e.g., tuneable beam splitter networks) and nonlinear processes (e.g., photon-photon interactions) in a layered structure reminiscent of classical deep neural networks. By tuning phase shifters and coupling strengths in these layers, one can train the QONN to perform tasks such as quantum state classification, quantum channel simulation, or even quantum reinforcement learning. Prior theoretical work (Steinbrecher *et al.* [22]) has demonstrated that QONNs can learn diverse quantum transformations, highlighting their potential for quantum machine learning.

A typical QONN layer consists of a multi-mode interferometer—realising an arbitrary linear unitary—followed by single-mode nonlinear elements serving as “activation functions”. The linear networks are often designed using programmable meshes of beam splitters and phase shifters, which can be fabricated on integrated photonic chips. Nonlinear elements are more challenging: common approaches include exploiting Kerr effects in waveguides or using quantum emitters in cavities to induce photon-photon interactions. Measurement-based or post-selected nonlinearities are also possible, and typically require feed-forward architectures. These nonlinear blocks are essential for enabling the network’s computational power, as purely linear optics cannot replicate the full functionality of a neural network.

In practice, QONNs rely on reliable photon-number states or squeezed states as photon sources, low-loss photonic circuits, and stable interferometer meshes. Photons passing through a programmed series of beam splitters and phase shifters experience a specific unitary transform, then undergo a photon-number-dependent phase shift (or other nonlinear effect) in each mode. This process can be repeated across multiple layers. The output is measured with




FIG. S14: **Implementation of a QONN in a CV framework.** Within the CV paradigm, a QONN is built from four functional blocks [23, 24]: I. unitary transformations enacted by linear optical elements; II. inline squeezing gates; III. displacements driven by local oscillators; and IV. nonlinear layers supplied by non-Gaussian operations—which can currently be realised through measurement-induced Kerr operations.

single-photon detectors, and a classical optimiser adjusts the optical parameters to minimise a training cost function. While large-scale QONN demonstrations are still in development, smaller prototypes have been built, showcasing tasks like quantum cloning and quantum state discrimination. As integrated photonics, single-photon sources, and on-chip nonlinearities continue to improve, QONNs are expected to play an increasingly important role in quantum computing and quantum machine learning research.

VI. NON-GAUSSIAN STATE PREPARATION

A. Cat State Breeding

A single-mode squeezed vacuum $\hat{S}(r)|0\rangle$ is tapped on a highly unbalanced beam splitter; detection of a single reflected photon heralds the action of the annihilation operator $\hat{a}$ acting on the transmitted mode. For moderate squeezing, $r \lesssim 0.8$, the heralded state is well approximated by an odd Schrödinger cat

$$|\mathcal{C}_-(\alpha)\rangle \propto |\alpha\rangle - |-\alpha\rangle, \quad \alpha \simeq \frac{1}{2}(e^r - e^{-r}) \approx \sqrt{\langle \hat{n} \rangle}, \quad (28)$$

where α sets the coherent amplitude. Two-photon subtraction (or phase-randomised squeezing) produces the even analogue $|\mathcal{C}_+(\alpha)\rangle$.

To amplify the amplitude, two identical small cats are overlapped on a 50:50 beam splitter, and one output is measured with homodyne detection. Conditioning on a near-zero x -quadrature outcome ($|x| \lesssim \varepsilon$) projects the remaining port onto a cat with

$$\alpha \longrightarrow \alpha' = \sqrt{2}\alpha. \quad (29)$$

Iterating this breed-and-herald step enlarges α after each round while suppressing higher-order photon components, provided the acceptance window ε is gradually narrowed. As both photon subtraction and homodyne conditioning are passive, the protocol is compatible with existing CV photonic platforms and can be repeated until optical loss becomes the dominant limitation.

As a complement to the in-line single-photon squeezing scheme discussed in the main text, here we describe a more traditional route to optical Schrödinger cat states that relies on conditional photon subtraction.

We start with a pulsed squeezed-vacuum state generated by an OPO operated with a squeezing parameter of $r = 0.5$. Approximately 5% of the field is tapped off on a highly transmissive beam splitter and directed to a PNRD. A single-photon click heralds the successful subtraction of one photon, thereby projecting the remaining mode onto an odd Schrödinger cat state. The ideal and experimentally reconstructed Wigner functions of the heralded state are shown in Figs. S15a and d.

In the first breeding stage, two such heralded cat states, generated in successive pump cycles, are interfered on a balanced beam splitter implemented in Core 1. Post-selection on the vacuum outcome in one output port then probabilistically yields, in the other output port, a cat state with an enlarged coherent-state separation, as displayed in Figs. S15b and e.

In the second breeding stage, two independently prepared first-stage cat states interfere on another balanced beam splitter realised in Core 2. A second probabilistic vacuum projection then heralds a larger cat state, depicted in Figs.

Ideal:




FIG. S15: **Numerical simulation and experimental results of the cat-state breeding process.** Here, we compare the lossless simulation results and the experimental outcomes. The small cat state is prepared by probabilistically subtracting one photon from a squeezed-vacuum state. Progressive breeding steps yield cat states with larger amplitudes.

S15c and f.

In the main text, by contrast, we use photon addition to generate the cat state. A single-photon Fock state is generated near-deterministically with the previously described boosting protocol. In the cat-state breeding experiments, the input squeezing is set to $r = 1.2$. To favour single-photon events, we intentionally introduce loss, reducing the overall transmittance to 0.25. Each pulse then heralds a photon with a probability of 16%. By multiplexing 15 consecutive pulses, the cumulative success probability rises to $1 - 0.84^{15} \approx 0.88$.

Next, we select a single-photon and inject it into the ISU with a squeezing parameter of $r = 0.3$. This squeezed Fock state can be regarded as a kitten state $\mathcal{C}$. In the experiment, we prepare four such kitten states, $\mathcal{C}_1$ – $\mathcal{C}_4$. First, pairs $\mathcal{C}_1$ – $\mathcal{C}_2$ and $\mathcal{C}_3$ – $\mathcal{C}_4$ interfere in Core 1; the two outputs, $\mathcal{C}_{1'}$ and $\mathcal{C}_{3'}$, subsequently interfere in Core 2. A breeding event is considered successful when no photon is detected in the heralding arm. Alternatively, a compass state (as shown in Fig. 2b in the main text) is generated when a two-photon event is heralded.

We perform homodyne measurements while scanning the LO phase continuously from 0 to 2π in increments of $\pi/6$, recording a quadrature distribution at each phase setting. The Wigner function is then reconstructed using maximum-likelihood techniques [33] from these quadrature measurements. The fidelities of the three states in Fig. 3a (from top to bottom) are 0.950 ± 0.037 , 0.959 ± 0.032 , and 0.938 ± 0.033 when benchmarked with numerical simulations including photon loss; the corresponding fidelities relative to the ideal (lossless) states are 0.733 ± 0.027 , 0.580 ± 0.025 , and 0.380 ± 0.020 , respectively. The compass state in Fig. 3b shows a 0.917 fidelity against photon-loss-included simulations and a 0.466 fidelity relative to the ideal lossless state.

B. GKP State Engineering

The quasi-deterministic generation of an approximate GKP grid state can be achieved by using four identical cat resources $|\text{cat}\rangle \propto |\alpha\rangle + |-\alpha\rangle$. The cats are first interfered pairwise on a cascade of balanced beam splitters, creating a four-mode entangled state in which the individual $\pm\alpha$ lobes combine to form a coarse, multi-peaked interference pattern. Three of the four output modes are then subjected to independent HD, chosen along suitable quadratures




FIG. S16: **Generation of GKP states through cat-state interference.** **a**, The cat states are first generated by squeezing single-photon Fock states. Next, four such cat states interfere, and the homodyne-detection outcomes are used to feed-forward an appropriate displacement operation to the output mode. **b**, The corresponding circuits for realising the detection and feed-forward operations.

so that the remaining mode is conditionally projected onto a grid-like wavefunction. The homodyne outcomes impart random phase-space displacements, so their electronic readouts are processed in real time and applied as feed-forward displacements $D(\delta x, \delta p)$ on the remaining mode, thereby cancelling the measurement-dependent offsets. After this single-shot correction, the remaining optical mode is left in a state whose position- and momentum-space representations exhibit the periodic comb structure characteristic of a GKP logical $|1_L\rangle$. The protocol relies solely on Gaussian operations, PNRDs are not required, and in principle its success probability is unity, making it a practical route to scalable, fault-tolerant GKP qubits.

Fig. S16b outlines our experimental sequence for engineering a GKP state. Four single-photon Fock states are first individually squeezed and then interfered inside the LOU cores according to the circuit in Fig. S16a. During syndrome readout, an EOM rotates the polarisation of each output so that the ancilla modes are directed to HDs




FIG. S17: **Detailed HD results.** Quadrature distributions measured by homodyne detection with local-oscillator phases θ spanning 0 to π .

and measured in the x -quadrature. The remaining signal mode is displaced by an LO whose amplitude is updated in real time from the homodyne outcomes. This feed-forward operation completes the correction step and produces the target GKP state; the corresponding homodyne traces are plotted in Fig. S17.

The GKP state in our work is defined as a square grid. This discretised lattice structure in phase space is known to have potential for reliably detecting and correcting small photon-loss and dephasing errors, thereby enabling robust fault-tolerant quantum computing. In the ideal limit—requiring infinite squeezing and perfectly lossless circuits—the logical basis states in the position representation are

$$|0_L\rangle \propto \sum_{k \in \mathbb{Z}} |x = 2k\sqrt{\pi}\rangle, \quad |1_L\rangle \propto \sum_{k \in \mathbb{Z}} |x = (2k+1)\sqrt{\pi}\rangle. \quad (30)$$

In experiments, the ideal Dirac-delta peaks are replaced by a comb of narrow Gaussian peaks of width Δ (smaller Δ corresponds to stronger effective squeezing) within a wide Gaussian envelope function of width $1/\kappa$:

$$|0_L^{(\Delta, \kappa)}\rangle \propto \int dx \sum_{n \in \mathbb{Z}} e^{-2\pi\kappa^2 n^2} e^{-(x-2n\sqrt{\pi})^2/(2\Delta^2)} |x\rangle, \quad (31)$$

and similarly for $|1_L^{(\Delta, \kappa)}\rangle$ with the shift $2n \rightarrow 2n+1$. Smaller Δ and κ give higher code fidelity but require more squeezing and limited loss.

The phase-space Pauli operators are defined as $\hat{X} = \exp(-i\sqrt{\pi}\hat{p})$ and $\hat{Z} = \exp(i\sqrt{\pi}\hat{x})$, corresponding to displacements of $\sqrt{\pi}$ along the x - and p -axes, respectively. The logical GKP states reside in the subspace stabilised by $\hat{S}_x = \hat{X}^2 = \exp(-i2\sqrt{\pi}\hat{p})$ and $\hat{S}_p = \hat{Z}^2 = \exp(i2\sqrt{\pi}\hat{x})$. In the ideal limit, they satisfy $\langle \hat{S}_x \rangle = \langle \hat{S}_p \rangle = 1$, meaning the wavefunction is strictly periodic with period $2\sqrt{\pi}$ along both quadratures. These ideal states are formally non-normalisable. In practice, each position eigenket is replaced by a finitely squeezed state, and the superposition is enveloped by a Gaussian to obtain finite energy while preserving the required x - and p -symmetry.

In this work, photonic approximate GKP states are generated via cat-state breeding for potential fault-tolerant quantum error correction. In principle, the protocol can be made deterministic if large-amplitude cat states are available on demand [40]. In our implementation, we first create a small cat state by squeezing a single-photon Fock state with an inline squeezing level of $r \approx 0.5$. We then interfere four cat states using linear optics and feed-forward based on specific measurement outcomes to build up the multi-peaked grid structure that defines an optical GKP state.

Due to imperfections such as limited squeezing and photon loss, the ideal grid spacing $2\sqrt{\pi}$ shrinks to a smaller value β . Fitting the separation between successive positive (or negative) peaks yields $\beta = 2.3123$. Replacing $2\sqrt{\pi}$ by

this fitted β in the stabiliser definitions, we evaluate $\langle \hat{S}_p \rangle = 0.4462 \pm 0.0077$, and determine $\langle \hat{S}_x \rangle$ as well as $\langle \hat{S}_{|L} \rangle$ in the main text from the reconstructed Wigner function. The corresponding effective squeezing of the peaks is then calculated as: $\Delta_x^2 = -\ln |\langle \hat{S}_x \rangle|^2 = 1.6451 \pm 0.0709$, $\Delta_p^2 = -\ln |\langle \hat{S}_p \rangle|^2 = 0.5137 \pm 0.0109$, and $\Delta_{\text{sys}}^2 = (\Delta_x^2 + \Delta_p^2)/2 = 1.0794 \pm 0.0374$. In addition, we calculate the degree of dispersion (Δ_x^2 and Δ_p^2) of the central peak, which is the small red circle surrounding the background of the Gaussian envelope in Fig. 3c. Error bars are calculated using bootstrapping methods.

The GKP generation in Clavina is currently limited by two technical factors: (i) the available single-mode squeezing (≈ 5 dB) broadens the GKP lattice peaks, adding noise to the homodyne readout; and (ii) the overall optical transmittance is ≈ 0.81 , so loss-induced random displacements further blur the syndrome signal. To guarantee correct feed-forward, we therefore reject high-noise measurements, leaving a small residual probabilistic overhead. In our experiment, we achieve a 2 kHz success rate in GKP state generation when operating the experiments at 250 kHz. Although imperfections preclude on-demand GKP state generation, the near-deterministically generated cat states and displacement feed-forward (incorporating both X - and Z -operations) boost the generation efficiency roughly two to three orders of magnitude higher than the latest probabilistic approaches [41, 42]. Raising the squeezing to ≥ 10 dB, reducing the total loss below 5%, and using higher-efficiency HDs in the future would enable near-deterministic photonic GKP state preparation.

Furthermore, Clavina also supports the emission of a stream of GKP states at a fixed clock rate in a straightforward way. Since each pulse occupies a distinct time bin while circulating in the same spatial mode, pairwise beam splitter operations applied synchronously at the loop output entangle successive bins into a one-dimensional GKP cluster. Cascading two identical loops with an additional 50:50 coupler transforms the resource into a two-dimensional GKP suitable for fault-tolerant surface-code encoding. This time-bin encoding circumvents the loss and alignment overheads of spatial multiplexing, enabling large-volume cluster states with only a single set of optics and detectors. By delivering a deterministic, or quasi-deterministic, supply of GKP qubits—embedded in a scalable cluster state—using the same hardware, Clavina provides an experimentally viable path from state preparation to universal, fault-tolerant photonic quantum computation.

Comparison with Gaussian-boson-sampling-based GKP state engineering

Probability.

The method introduced in Larsen *et al.* [42] use few-mode GBS circuits on a PIC with photon-number-resolving measurements on all but one output to provide the non-Gaussian resource for post-selected GKP-state engineering. The input squeezing, interferometer settings, and photon-number pattern are optimised in advance to balance fidelity and success probability. Although the ideal success probability is < 0.05 , the architecture relies on many PICs operating in parallel and a network of EOMs to multiplex successful outputs into a “GKP refinery”. Using a single chip, a success probability of 2.9×10^{-4} and a 58 Hz success rate was achieved before discarding noisy outcomes.

By contrast, our scheme introduces the non-Gaussian resource at the input, using squeezed single-photon states as kitten states. While single-photon generation is intrinsically probabilistic with SPDC, our “boosting method” effectively multiplexes successfully heralded photons into the ISU, enabling near-deterministic preparation (see Supplementary Information, Section II.C). Once four squeezed single photons are available, the subsequent cat/GKP breeding is deterministic and requires only Gaussian homodyne measurements at room temperature; the remaining probabilistic element comes solely from the near-deterministic heralded-photon source. This offers two practical advantages: multiplexing heralded single photons is much simpler than multiplexing full GKP states, which requires many PICs and photon-number-resolving detectors, and temporal multiplexing is native to Clavina’s fully time-bin-encoded architecture, making it naturally compatible with large time-domain cluster states. By comparison, Larsen *et al.*’s [42] approach combines path encoding on PICs, temporal multiplexing of GKP states, and time-bin encoding for cluster-state generation.

Size and quality.

The size and quality of the generated GKP states are determined mainly by the amount of nonlinear resource used and by the achievable squeezing and loss, respectively. In Larsen *et al.* [42], the current method can produce a two-dimensional 3×3 grid conditioned on 3-3-3 photon-sampling outcomes, corresponding to nonlinear resources equivalent to nine photons. Its main advantage is that larger GKP states can, in principle, be obtained by post-selecting higher-photon-number events, though its main drawback is the very low success probability.

Instead, our approach starts from four squeezed single-photon states used as small-amplitude cat (kitten) states, corresponding to nonlinear resources equivalent to four photons. We report the first demonstration of two rounds of cat breeding, and the scheme is scalable: more squeezed single photons and more breeding rounds would lead to full two-dimensional GKP states. However, in the current implementation, we aim to demonstrate a multifunctional photonic quantum-computing platform rather than a device optimised solely for GKP state generation. Our single apparatus supports, but is not limited to, large-scale GBS, cluster-state generation, Cat/GKP state engineering, and Bose-Hubbard simulation. A dedicated device built only to produce one GKP state, as in Larsen *et al.* [42], could




FIG. S18: **Generating GKP cluster states and surface codes with Clavina.** Building on the GKP state generation circuits developed in this work, each GKP state can be injected into the LOU where programmed interference using a modified variable delay line produces the required entanglement. In this way, Clavina is capable of generating large-scale GKP cluster states and surface-code lattices.

likely reduce loss and yield a slightly larger two-dimensional GKP state. Instead, we trade some size for higher fidelity, orders-of-magnitude improvements in success rate, and broader functionality.

Photon-loss error correction with the GKP code

The GKP code can be employed to protect quantum information against photon loss entirely within a single bosonic mode [32]. In this framework, photon loss is modelled by a beam splitter of reflectivity γ that couples the mode to the vacuum. This process contracts each $\sqrt{\pi}$ -spaced lattice peak of the GKP state towards the phase-space origin, while preserving the state’s intrinsic $\sqrt{\pi}$ -periodicity. Error recovery is achieved via ideal GKP error correction as follows:

- (1) Non-destructively measure the modular position and momentum quadratures ($\hat{x} \bmod \sqrt{\pi}$ and $\hat{p} \bmod \sqrt{\pi}$) using an ancillary GKP qubit and Gaussian entangling gates;
- (2) Decompose each measurement outcome m into an integer multiple of $\sqrt{\pi}$ plus a fractional part $\delta m \in (-\sqrt{\pi}/2, \sqrt{\pi}/2]$;
- (3) Apply corrective phase-space displacements to shift each fractional remainder to zero (“standard binning”).

Crucially, the pure-loss channel commutes with these modular displacement corrections. As a result, this procedure projects the corrupted state back onto the two-dimensional GKP codespace, converting the continuous phase-space contraction from loss into an equivalent discrete Pauli error on the logical qubit. The effective error probability (Pauli flip rate) is determined by the loss rate γ and by the finite-energy damping parameter β of the approximate GKP state.

Harris *et al.* [32] derive closed-form expressions for the Kraus operators and process matrices of this induced logical channel, considering both the unconditional loss case (no information about the lost photons) and the heralded loss case (the lost light is measured via photon counting or heterodyne detection before GKP correction). Their results indicate that for realistic parameters ($\beta \approx 0.1$, $\gamma \lesssim 0.2$), the residual logical noise is dominated by small Pauli X - or Z -flip errors. Furthermore, by conditioning on the measured loss (i.e., using heralding), one can deliberately engineer non-Pauli logical noise channels that are useful for tasks such as magic-state distillation.

In summary, the GKP code translates analogue photon-loss errors into a discrete syndrome that can be corrected with Gaussian operations and classical feed-forward, enabling fault-tolerant error correction against photon loss without requiring additional physical qubits.

C. GKP Cluster States and Surface Code

Our Clavina is perfectly suited to preparing GKP cluster states and implementing surface codes—tasks that are difficult to achieve with prior architectures. As shown in the left panel of Fig. S18, the GKP state preparation proceeds exactly as introduced in the main text. A squeezer generates a train of squeezed-vacuum pulses. Each pulse interacts with a non-Gaussian ancilla from a heralded single-photon source. A homodyne measurement then conditionally projects the output into an approximate GKP state. Repeating this sequence at a fixed clock rate yields a time-bin stream of single-mode GKP qubits, supplying scalable resources without the need for multiple sources.

The right panel of Fig. S18 illustrates a time-bin multiplexing network that entangles the sequential GKP pulses into a large-scale cluster state. Here, components are reused in the loop. The optical pulses are fed through the I/O module and interfered multiple times using the same optical elements. The EOMs act as variable beam splitters, and a variable delay line allows each GKP mode to interfere with others at designated time offsets. For example, a short delay (one time bin) can entangle each pulse with the next, forming a one-dimensional cluster chain, while a longer delay (e.g., connecting each pulse to the N th subsequent pulse) “fuses” the chain into a two-dimensional lattice. A fast optical I/O switch synchronises the injection of new GKP states into the loop and the extraction of the entangled cluster output at the proper time. Clavina enables the deterministic production of massive, topologically structured cluster states using only a handful of hardware components. By leveraging EOM-driven reconfigurable delay lines and a compact set of optical devices, this platform is capable of generating scalable GKP cluster states for fault-tolerant quantum computing.

In preparing a surface-GKP code, each optical mode is encoded as a GKP qubit, and frequent GKP syndrome extraction discretises small phase-space displacements caused by loss and noise into correctable shifts. Arrays of such GKP qubits populate a two-dimensional surface-code lattice, where stabiliser measurements detect and repair residual Pauli-type errors. Logical computation interleaves entangling Gaussian interactions (e.g., SUM or controlled-phase gates) with GKP correction and repeated surface-code cycles. Finally, incorporating analogue, real-valued GKP syndromes into the surface-code decoder significantly improves error thresholds and reduces overall computational overhead.

VII. BOSE-HUBBARD HAMILTONIAN SIMULATION

The Bose-Hubbard model and its simulation results

The Bose-Hubbard model on an L -site lattice is described by the Hamiltonian

$$\hat{H} = -J \sum_{\langle i,j \rangle} (\hat{a}_i^\dagger \hat{a}_j + \hat{a}_j^\dagger \hat{a}_i) + \frac{U}{2} \sum_{i=1}^L \hat{n}_i (\hat{n}_i - 1), \quad (32)$$

where $\langle i, j \rangle$ indicates a sum over adjacent sites i and j , and $\hat{n}_i = \hat{a}_i^\dagger \hat{a}_i$. The competition between the tunnelling term (proportional to J) and the on-site repulsive interaction (controlled by U) gives rise to rich quantum phases. Most notably, the Bose-Hubbard model hosts the celebrated superfluid-Mott insulator transition: for small U/J , bosons delocalise and form a superfluid state with long-range phase coherence, whereas for large U/J , they localise, yielding a Mott-insulating state with integer occupation per site and an excitation gap. This transition has been experimentally verified in ultracold atomic gases loaded into optical lattices, establishing the Bose-Hubbard framework as a cornerstone in quantum many-body physics.

In our experiments, we consider three distinct sites. Although the system size is relatively small, it captures several important quantum many-body phenomena, including Rabi oscillations in the non-interacting limit ($U \approx 0$), Josephson oscillations for finite interactions, and the onset of macroscopic self-trapping when U dominates over J . Specifically, if the initial state is prepared with a population imbalance between the modes, a moderate interaction strength induces nonlinear Josephson oscillations in the population and phase difference, reflecting bosonic analogues of superconducting Josephson junctions. In the strong-interaction regime, the model exhibits self-trapping: tunnelling is suppressed, and one mode remains macroscopically occupied while the others remain nearly empty. Such dynamical regimes have been widely studied in double-well Bose-Einstein condensates and provide a valuable testbed for comparing theoretical predictions with experimental observations.

Experimental realisation

In our experiments, the two-photon Fock states $|200\rangle$ and $|110\rangle$ are produced using the boosting technique discussed in Fig. S2. During the BHM simulations, the system is operated at a clock rate of 2 MHz. A variable beam splitter, realised with a half-wave plate followed by a polarising beam splitter, is set to an 80:20 transmittance-reflectance




FIG. S19: **Numerical simulation results of Bose-Hubbard dynamics versus interaction strength.** Ideal simulation results (without Trotter error) for a three-site Bose-Hubbard chain with the dimensionless interaction strength U/J ranging from 0.5 to 3.5, plotted after an evolution time $t = 0.5$. The coloured traces give the normalised populations $P_{n_1 n_2 n_3}$ of the main Fock configurations $|n_1 n_2 n_3\rangle$ as a function of U/J . **a**, Dynamics starting from the imbalanced state $|200\rangle$. **b**, Dynamics starting from the two-site state $|110\rangle$.




FIG. S20: **Numerical simulation results of Bose-Hubbard model evolution over time.** Ideal simulation results (without Trotter error) when the dimensionless interaction strength is fixed at $U/J = 1$. **a**, Dynamics starting from the imbalanced initial state $|200\rangle$. The population undergoes Rabi-like oscillations between $|200\rangle$, $|110\rangle$, and $|020\rangle$, demonstrating coherent tunnelling of two bosons. **b**, Dynamics starting from the state $|110\rangle$. Here the three dominant configurations $|110\rangle$, $|101\rangle$, and $|011\rangle$ exchange population in a cyclic fashion, corresponding to pairwise hopping of single bosons between neighbouring sites. The simulations highlight the high-contrast, interaction-dependent oscillations that will be mapped onto photon number in the experiment.

ratio. The effective Kerr strength can be tuned from approximately $\pi/25$ to $\pi/3.5$, implemented by tuning EOMb in the NLU (Fig. S1). We use two Trotter steps in this experiment, and the final Kerr-gate layer can be omitted because we immediately perform PNRD afterwards. Including optical loss throughout the circuit, the overall heralded success rate is approximately 86.5 ± 30 Hz. For each experimental setting, we record > 1000 successful events, from which the output photon-number distributions are calculated. Numerical simulations of the ideal results without the Trotter error for varying interaction strengths and evolution times, corresponding to Fig. 4d and Fig. 4e (bottom panel), are shown in Figs. S19 and S20, respectively.

Experimental imperfections and photon-number distribution post-processing

In this section, we discuss how noiseless linear amplification introduced by the NLU, and unbalanced detection efficiencies—affects the results.

We first analyse the noiseless amplification arising from measurement-induced Kerr gates [4]. Since this gate is

applied equally to all modes:

$$G_1(g) = G_2(g) = \dots = G_L(g) = \sum_{n=0}^{\infty} g^n |n\rangle\langle n|, \quad g > 1. \quad (33)$$

Hence, every photon-number state $|n\rangle$ acquires a purely real amplitude factor g^n after compensating for the Kerr phase with an external phase shifter. As the same gain is applied to all these modes, acting on a Fock basis vector yields:

$$\mathcal{G}(g) |n_1, n_2, \dots, n_L\rangle = g^{\sum_{j=1}^L n_j} |n_1, n_2, \dots, n_L\rangle. \quad (34)$$

In our experiments, we use a three-mode circuit and input a total of two photons to simulate the dynamics of two bosonic particles across three sites. The experiment post-selects the subspace with a fixed total photon number $N = n_1 + n_2 + n_3 = 2$, and the relevant basis vectors are:

$$|200\rangle, |020\rangle, |002\rangle, |110\rangle, |101\rangle, |011\rangle. \quad (35)$$

For all these states, the exponent in the previous equation equals two, giving:

$$\mathcal{G}(g) |n_1, n_2, n_3\rangle = g^2 |n_1, n_2, n_3\rangle. \quad (36)$$

Thus, the same amplitude factor g^2 multiplies every basis state in the two-photon manifold.

Consequently, we find that when all three spatial modes share the same gain g and the analysis is restricted to a fixed total photon number $N = 2$, the noiseless amplifier acts as an overall scalar multiplier on the six-dimensional subspace. As a result, the relative probabilities of the states $|200\rangle$, $|020\rangle$, $|002\rangle$, $|110\rangle$, $|101\rangle$, and $|011\rangle$ remain exactly those produced by the preceding interferometric network. The raw histogram obtained with a PNRD already represents the true Bose-Hubbard output distribution; no physical noiseless attenuation or numerical rescaling is required.

Next, we discuss the effect of unbalanced detection efficiencies. As our multi-pixel PNRD exhibits photon-number-dependent efficiencies— $\epsilon_2 \simeq 0.80$ when both photons land in the same mode and $\epsilon_{1,1} \simeq 0.94$ when the photons are split between two modes—we model the experiment with the transfer matrix $\mathbf{M} = \mathbf{D}\mathbf{L}(\eta)\mathbf{G}(g)$. The detector-response matrix is $\mathbf{D} = \text{diag}(\epsilon_2, \epsilon_2, \epsilon_2, \epsilon_{1,1}, \epsilon_{1,1}, \epsilon_{1,1})$, acting on the six two-photon events $(2, 0, 0)$, $(0, 2, 0)$, $(0, 0, 2)$, $(1, 1, 0)$, $(1, 0, 1)$, and $(0, 1, 1)$. The channel-loss matrix $\mathbf{L}(\eta)$ is a 6×6 diagonal matrix whose k th entry equals the product of the calibrated transmittances of the ports involved; for equal transmittances, it reduces to $\mathbf{L}(\eta) = \eta^2 \mathbf{I}_6$. The gain matrix is the identity, $\mathbf{G}(g) = \mathbf{I}_6$, because the electronic amplifier introduces no port- or photon-number-dependent weighting, and any overall gain factor is absorbed into the total number of triggers N_{tot} . Hence, for the six-component ideal probability vector $\mathbf{p}_{\text{ideal}}$ and the observed count vector $\mathbf{n}_{\text{obs}}$, we have $\mathbf{n}_{\text{obs}} = N_{\text{tot}} \mathbf{M} \mathbf{p}_{\text{ideal}}$. We recover $\mathbf{p}_{\text{ideal}}$ by non-negative maximum-likelihood (or non-negative least-squares) inversion of $\mathbf{M}$ and propagate Poisson shot-noise uncertainties via Monte Carlo resampling, yielding an unbiased, loss- and efficiency-corrected estimate of the ideal two-photon distribution.

-
- [1] M. Mehmet *et al.* Squeezed light at 1550 nm with a quantum noise reduction of 12.3 dB. *Opt. Exp.* **19**, 25763 (2011).
 - [2] L. S. Madsen *et al.* Quantum computational advantage with a programmable photonic processor. *Nature* **606**, 75-81 (2022).
 - [3] A. I. Lvovsky, Squeezed light. Preprint at <https://arxiv.org/pdf/1401.4118> (2016).
 - [4] L. S. Costanzo *et al.* Measurement-induced strong Kerr nonlinearity for weak quantum states of light. *Phys. Rev. Lett.* **119**, 013601 (2017).
 - [5] Biagi, N., Francesconi, S., Zavatta, A. & Bellini, M., Photon-by-photon quantum light state engineering. *Progress in Quantum Electronics* **84**, 100414 (2022).
 - [6] S. Yu *et al.* A universal programmable Gaussian boson sampler for drug discovery. *Nat. Comput. Sci.* **3**, 839-848 (2023).
 - [7] S. Takeda, K. Takase, A. Furusawa, On-demand photonic entanglement synthesizer. *Sci. Adv.* **5**, eaaw4530 (2019).
 - [8] Y. Enomoto, K. Yonezu, Y., Takase, K. & Takeda S. Programmable and sequential Gaussian gates in a loop-based single-mode photonic quantum processor. *Sci. Adv.* **7**, eabj6624 (2021).
 - [9] S. Sempere-Llagostera, R. B. Patel, I. A. Walmsley, W. S. Kolthammer, Experimentally finding dense subgraphs using a time-bin encoded Gaussian boson sampling device. *Phys. Rev. X* **12**, 031045 (2022).
 - [10] S. Yu *et al.* A von-Neumann-like photonic processor and its application in studying quantum signature of chaos. *Light: Sci. & App.* **13**, 74 (2024).
 - [11] W. R. Clements *et al.* Optimal design for universal multiport interferometers. *Optica* **12**, 1460-1465 (2016).

- [12] Y. Zhu *et al.* A hybrid single quantum dot coupled cavity on a CMOS-compatible SiC photonic chip for Purcell-enhanced deterministic single-photon emission. *Light: Sci. & App.* **14**, 86 (2025).
- [13] S. Takeda *et al.* Deterministic quantum teleportation of photonic quantum bits by a hybrid technique. *Nature* **500**, 315-318 (2013).
- [14] G. Harder *et al.* Single-mode parametric-down-conversion states with 50 photons as a source for mesoscopic quantum optics. *Phys. Rev. Lett.* **116**, 143601 (2016).
- [15] C. Taballione *et al.* 8×8 reconfigurable quantum photonic processor based on silicon nitride waveguides. *Optics Express* **27**, 26842 (2019).
- [16] A. Eckstein, A. Christ, P. J. Mosley, C. Silberhorn, Highly efficient single-pass source of pulsed single-mode twin beams of light. *Phys. Rev. Lett.* **106**, 013603 (2011).
- [17] B. A. Bell *et al.* Testing multi-photon interference on a silicon chip. *Optics Express* **27**, 35646 (2019).
- [18] A. I. Lvovsky, Squeezed light. arXiv:1401.4118 (2016).
- [19] S. Sempere-Llagostera *et al.* Experimentally finding dense subgraphs using a time-bin encoded Gaussian boson sampling device. *Phys. Rev. X* **12**, 031045 (2022).
- [20] J. Du, W. Li, M. Bajcsy, Deterministic single-photon subtraction based on a coupled single quantum dot-cavity system. *Opt. Exp.* **28**, 6835-6845 (2020).
- [21] S.-Q. Jiang *et al.* Deterministic single-photon adder and subtractor. preprint on arXiv:2502.11376 (2025).
- [22] Steinbrecher, G. R. *et al.* Quantum optical neural networks. *npj Quantum Inf.* **5**, 60 (2019).
- [23] N. Killoran *et al.* Continuous-variable quantum neural networks. *Phys. Rev. Res.* **1**, 033063 (2019).
- [24] S. Yu *et al.*, Shedding light on the future: Exploring quantum neural networks through optics. *Adv. Quan. Tech.* 2400074 (2024).
- [25] R. Raussendorf, H. J. Briegel, A one-way quantum computer. *Phys. Rev. Lett.* **86**, 5188-5191 (2001).
- [26] P. Walther *et al.*, Experimental one-way quantum computing. *Nature* **434**, 169-176 (2005).
- [27] N. C. Menicucci *et al.*, Universal quantum computation with continuous-variable cluster states. *Phys. Rev. Lett.* **97**, 110501 (2006).
- [28] R. Ukai *et al.* Demonstration of unconditional one-way quantum computations for continuous variables. *Phys. Rev. Lett.* **106**, 240504 (2011).
- [29] Asavanant, W. *et al.* Generation of time-domain-multiplexed two-dimensional cluster state. *Science* **366**, 373-376 (2019).
- [30] Larsen, M. V., Guo, X., Breum, C. R., Neergaard-Nielsen, J. S. & Andersen, U. L., Deterministic generation of a two-dimensional cluster state. *Science* **366**, 369-372 (2019).
- [31] S. Yokoyama *et al.*, Ultra-large-scale continuous-variable cluster states multiplexed in the time domain, *Nat. Photon.* **7**, 982-986 (2013).
- [32] T. B. Harris *et al.* Logical channel for heralded and pure loss with the Gottesman-Kitaev-Preskill code. Preprint at arXiv:2504.13497 (2025).
- [33] A. I. Lvovsky, Iterative maximum-likelihood reconstruction in quantum homodyne tomography. *J. Opt. B Quantum Semi-classical Opt.* **6**, S556-S559 (2004).
- [34] Sotier, F. *et al.* Femtosecond few-fermion dynamics and deterministic single-photon gain in a quantum dot. *Nat. Phys.* **5**, 352-356 (2009).
- [35] Steindl, P. *et al.* Artificial coherent states of light by multiphoton interference in a single-photon stream. *Phys. Rev. Lett.* **126**, 143601 (2021).
- [36] Lund, M. M., Yang, F., Christiansen, V. R., Kornovan, D. & Mølmer, K., Subtraction and addition of propagating photons by two-level emitters. *Phys. Rev. Lett.* **133**, 103601 (2024).
- [37] Pasharavesh, A. & Bajcsy, M., Generation of non-Gaussian states of light using deterministic photon subtraction. *New J. Phys.* **26**, 113022 (2024).
- [38] M. Arcari, *et al.* Near-unity coupling efficiency of a quantum emitter to a photonic crystal waveguide. *Phys. Rev. Lett.* **113**, 093603 (2014).
- [39] S. Shi, *et al.* High-fidelity photonic quantum logic gate based on near-optimal Rydberg single-photon source. *Nat. Commun.* **13**, 4454 (2022).
- [40] Winnel, M. S., Guanzon, J. J, Singh, D. & Ralph, T. C, Deterministic preparation of optical squeezed cat and Gottesman-Kitaev-Preskill states. *Phys. Rev. Lett.* **23**, 230602 (2024).
- [41] Konno, S. *et al.* Logical states for fault-tolerant quantum computation with propagating light. *Science* **383**, 289-292 (2024).
- [42] Larsen, M. V. *et al.* Integrated photonic source of Gottesman-Kitaev-Preskill qubits. *Nature* **642**, 587-591 (2025).