nature portfolio

Peer Review File

Extensible universal photonic quantum computing with nonlinearity

Corresponding Author: Dr Shang Yu

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
Review of “Extensible universal photonic quantum computing with nonlinearity”

The authors of the paper report on a novel modular temporal-multiplexed architecture, centered around a (potentially “multi-core”) linear optics unit, in which non-linear optical elements can be straightforwardly connected to create non-Gaussian photonic states.

I think the paper is timely and incorporates key techniques that have been improved over the last 15 years to address the two most important problems in quantum information processing based on photons: loss and photon-photon interaction. The authors demonstrate the achievement of non-linear Kerr interactions to generate cat states and subsequent generation of GKP states, which have been demonstrated to effectively handle photon loss. On top of this, they can in principle generate GKP states deterministically, which is, in my view, the most interesting part of the work.

They also provide other simple yet convincing demonstrations of the capabilities of their architecture: simulation of a Bose-Hubbard Hamiltonian and Gaussian Boson Sampling.

The quality of the work is excellent, as well as the methodology. I therefore strongly recommend publication in Nature Photonics after minor revisions, which include the following:

Main Text:

The text is for the most part well-written, save for a few sentences:

-In page 3, the sentence starting by “By implementing time-bin encoding[12, 23, 24] within a loop-based design” is not easy to follow. It should be modified.
-In page 4, the last paragraph starting by “Beyond linear operations, Clavina’s extensibility offers two routes to introduce nonlinear” is halfway through a description of the hardware and a prescription for usage. The subject is missing in some sentences.
-In Fig 4, the sentence ending with “Kerr module to main QPU.” s missing a pronoun.
-In the Discussion, the sentence including “including either photon-number source and nonlinear gates,” is not grammatically correct.

I have the feeling the figures could be referenced in a tidier manner. Figure 1 is referenced after figure 3 in the main text. Panels 2(b) and 2(c) are not informative. It is unclear from the caption or from the main text what information they want to convey. I would probably explain GBS in the text before pointing to Fig 2(c). In the caption of Fig.2, it reads that error bars denote standard error. This is insufficient from a methodological viewpoint. In Fig. 4(a) the caption does not correspond to the panel. There is a Bose-Hubbard term, corresponding to a Kerr non-linearity, but no fidelity is shown.

I would also tend not to conflate the concepts of universal quantum computation with non-Gaussianity, as implied in the introduction. Although from a certain perspective, it is possible to equate the two for continuous variables, when it comes to real fault-tolerant computing, universality will arise not from non-linearity only, but from tailored codes that exploit gate transversality and presumably magic state distillation. This is not discussed in the text, which I think is unfortunate given the discussion of universality of quantum dynamics in the context of quantum processors.

Supplemental Material:

Many references are missing in the text. It looks like the Latex compilation did not work properly.

It is apparent that the main text was written by a non-English speakr, which is perfectly ok. However, grammatical errors should be corrected, which is easily achieved by a native or proficient English speaker going through the text. Some examples:

Above eq.5 : “derivation of how the Kerr nonlinearity realized through photon”
Above eq.18: “Not that”
Section VII: “NON-GAUSSIAN STATE PREPARING”

Reviewer #2

(Remarks to the Author)
The manuscript review is included in the attached file.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The manuscript has been substantially improved, and previously seemingly unrelated pieces of information (in particular about the Boson Sampling part) are now presented in a more coherent way.

I would now recommend this work for publication.

Comments:

I was surprised by the statement that "326 electro-optic modulation, the circuit can be reconfigured on the fly, making both the couplings Jij and the on-site interactions U tunable in time. Such capabilities are typically unavailable in superconducting qubit architectures". I would have thought that superconducting architectures are at least as versatile as photonic ones, in particular when it comes to tuning interactions (via fluxes). The fact that they are not modified on-the-fly responds more to the fact that photonic computation happens as a stream of photons rather than on a physical support.

As a minor comment, the ISU acronym is defined in the caption of a figure, rather than in the main text.

Reviewer #2

(Remarks to the Author)
The authors have adequately addressed all my comments and provided the necessary amendments in the newly submitted version of the manuscript. In particular, they added to the Methods section a fair comparison between their approach to deterministic GKP generation and the one presented by the Xanadu group (side note: the authors should correct a few typos between lines 653–656). They also included missing information in the figure captions and reorganized their content. Each figure is now properly referenced in the main text, and the captions are self-explanatory.

As stated during the first round of review, the technical quality of the work is impressive, and I have no doubt that the manuscript meets the high standards set by Nature Photonics.

Therefore, I recommend publication.

Reply to Referee 1:

The authors of the paper report on a novel modular temporal-multiplexed architecture, centered around a (potentially "multi-core") linear optics unit, in which non-linear optical elements can be straightforwardly connected to create non-Gaussian photonic states.

I think the paper is timely and incorporates key techniques that have been improved over the last 15 years to address the two most important problems in quantum information processing based on photons: loss and photon-photon interaction. The authors demonstrate the achievement of non-linear Kerr interactions to generate cat states and subsequent generation of GKP states, which have been demonstrated to effectively handle photon loss. On top of this, they can in principle generate GKP states deterministically, which is, in my view, the most interesting part of the work.

They also provide other simple yet convincing demonstrations of the capabilities of their architecture: simulation of a Bose-Hubbard Hamiltonian and Gaussian Boson Sampling. The quality of the work is excellent, as well as the methodology. I therefore strongly recommend publication in Nature Photonics after minor revisions.

Reply: We sincerely appreciate and thank the referee for taking the time to review this manuscript and for the strong recommendation.
We have carefully considered the referee's comments and suggestions and have made substantial revisions accordingly. In response, we have revised the presentation, including the wording to improve the clarity and flow of the text throughout. We have reorganised parts of the figures, as suggested. We have also added more details on Gaussian boson sampling and the description of the error bars, and clarified our presentation and description regarding universality.
We believe that these changes have substantially improved the quality of the manuscript and brought it to the standard required for publication. Below, we provide a comprehensive point-by-point response to the referee's comments and suggestions. In the revised manuscript, these changes are highlighted in red.

-In page 3, the sentence starting by "By implementing time-bin encoding[12, 23, 24] within a loop-based design" is not easy to follow. It should be modified.

Reply: We have revised this sentence as follows:

"*The combination of time-bin encoding [13,25,26] and a loop-based architecture [11] allows functional modules to be located along a single optical path. Each module can be reused to perform operations on successive time bins, programmed by the CU. The system's capabilities can be further extended by adding additional modules via the CU's input and output ports. Quantum states of light can be selected from coherent, squeezed, or heralded photon-number states, and detectors can be chosen from superconducting nanowire single-photon detectors (SNSPDs), photon-number-resolving detectors (PNRDs) or homodyne detectors (HDs). In addition, a linear operation unit (LOU), supporting multiple computational cores, shortens the overall optical path length for the same circuit depth and, together with phase-stabilised delay lines, implements large-scale linear transformations.*"

This revision can be found on page 3, lines 130-146.

-In page 4, the last paragraph starting by "Beyond linear operations, Clavina's extensibility offers two routes to introduce nonlinear" is halfway through a description of the hardware and a prescription for usage. The subject is missing in some sentences.

Reply: We have revised this paragraph as:

"*Beyond linear operations, Clavina's extensibility offers two routes to introducing the tunable nonlinear resources required for universality. First, nonlinear resource states, such as photon-number states, can be injected at the input and then subjected to inline squeezing via the ISU. Second, the NLU can apply a Kerr gate to the circulating states. The loop-based architecture allows individual time-bins to be arbitrarily switched into either unit, depending on the application. This flexibility in interleaving the ISU and the NLU with LOU executions enables a variety of computational tasks and substantially advances the realisation of future universal fault-tolerant photonic quantum computers. In the following sections, we demonstrate the utility of these resources for GKP state engineering and Bose-Hubbard model simulation.*"

The above revision can be found on page 3, lines 178-192. We reserve further descriptions of the usage of these nonlinear resources for the sections that follow.

-In Fig 4, the sentence ending with "Kerr module to main QPU." s missing a pronoun.

Reply: We have rewritten this sentence as:

"*This model can be simulated on Clavina by interfacing the NLU with the LOU, via the central CU. The LOU implements the hopping term (J) by coupling time-bin modes, analogous to tunnelling between neighbouring sites, while the Kerr module realises the on-site interaction, U.*".

We also refer specifically to the LOU for implementing the hopping term, rather than the "QPU". We consider Clavina the QPU architecture, comprising individual functional units such as the LOU, CU, ISU, NLU etc.

This revision can be found in the caption of Fig. 4, on page 6.

-In the Discussion, the sentence including "including either photon-number source and nonlinear gates," is not grammatically correct.

Reply: We have revised this sentence as:

"*To overcome the longstanding limitation of photonic quantum computing to linear operations, we propose an extensible architecture that integrates scalable linear networks with diverse nonlinear functional modules. We demonstrate an inline squeezer applied to Fock states and a nonlinear Kerr gate, both of which are crucial for advancing toward universality.*"

The above revision can be found on page 7, lines 338-344.

I have the feeling the figures could be referenced in a tidier manner. Figure 1 is referenced after figure 3 in the main text. Panels 2(b) and 2(c) are not informative. It is unclear from the caption or from the main text what information they want to convey. I would probably explain GBS in the text before pointing to Fig 2(c). In the caption of Fig.2, it reads that error bars denote standard error. This is insufficient from a methodological viewpoint. In Fig. 4(a) the caption does not correspond to the panel. There is a Bose-Hubbard term, corresponding to a Kerr non-linearity, but no fidelity is shown.

Reply: We thank the referee for identifying these errors and suggestions. We have modified the order in which the figures are referenced. Per Referee 2's suggestion, we have also split Fig. 2f into two, and relocated the sub-figures to Fig. 3a and Fig. 4a., to better introduce the applications that utilise each unit.

Regarding the referencing of figures, Figure 1 is now referenced at the beginning of the section "Extensible photonic quantum processor" on Page 3. Figure 2 is referenced in lines 163-165. Figure 3 is first mentioned in the next section "Non-Gaussian state engineering" (line 207), and Figure 4 is first referenced in following section "Photonic quantum many-body simulations" (line 284).

- In the caption of Fig. 2b, we modified "Error bars denote the standard error" as:

  "*The points show the mean fidelity at each time; error bars denote* ± 1 *standard error of the mean over ten independent experimental runs.*"

- In the caption of Fig. 2c, we added the following information:

  "*The participation coefficients for eight-cliques associated with all nodes in a 100-node network are extracted from Gaussian boson sampling data with more than three photons.*"

- In the main text, lines 162-172, we describe the GBS experiment before pointing to Fig. 2c:

  "*Clavina's scalability and full programmability are demonstrated using a dual-core LOU (Fig. 2a), which maintains high phase stability over long operating times (Fig. 2b). As a benchmark, we performed a 100-mode Gaussian boson sampling experiment. This sampling task relies on tunable squeezed-state inputs and a fully reconfigurable interferometer, thereby providing a stringent test of adjustable squeezing and programmable linear-optical control. We used this experiment to evaluate the participation coefficients of a network [11,27], whose results are shown in Fig. 2c.*".

- The layout of Fig. 4 has been modified to now illustrate the operation of the Kerr gate in the nonlinear unit. The panels have been rearranged to ensure the captions match the sub-figure.

I would also tend not to conflate the concepts of universal quantum computation with non-Gaussianity, as implied in the introduction. Although from a certain perspective, it is possible to equate the two for continuous variables, when it comes to real fault-tolerant computing, universality will arise not from non-linearity only, but from tailored codes that exploit gate transversality and presumably magic state distillation. This is not discussed in the text, which I think is unfortunate given the discussion of universality of quantum dynamics in the context of quantum processors.

Reply: We thank the referee for highlighting this important distinction. We entirely agree that while non-Gaussianity (combined with a complete Gaussian toolkit) mathematically completes the universal gate set for continuous-variable (CV) systems at the physical hardware level, achieving true, fault-tolerant universal quantum computation (FTQC) requires a much more rigorous framework. As the reviewer rightly points out, universality of FTQC, at scale, requires bridging this physical layer to the logical layer through tailored quantum error-correcting codes, the execution of transversal logical gates, and the integration of magic state distillation.

Our original phrasing in the Introduction focused heavily on the historical hardware bottleneck for photonic systems—namely, that without non-Gaussian resources, photonic platforms are strictly bounded to classical simulability (non-universality). However, we recognise that this hardware-level milestone should not be conflated with the full paradigm of fault-tolerant universality.

- We have adjusted the opening paragraphs to clarify that integrating nonlinearities with linear circuitry achieves a hardware-level universal gate set, which is a prerequisite—but not a complete solution—for FTQC. Lines 25-38 read:

"The long-term goal of realising a large-scale fault-tolerant photonic quantum computer requires transversal logic, and magic state distillation for the implementation of tailored error-correcting codes such as the surface code [12]. An immediate experimental challenge lies in developing architectures consisting of a complete physical gate set of linear and nonlinear operations.

Recent experiments have separately demonstrated scalability [1, 3, 13], programmability in linear optics [3, 11], and access to nonlinear operations [14, 15]. Historically, optical circuitry in these architectures remains non-universal at the hardware level [16-18], as single-photon-level nonlinearities are intrinsically difficult to realise..."

- We altered the sentence (lines 79-84) "This achieves universal gate sets, and thus completes the final piece required for universal photonic computing." to read:

  "*By realising a universal gate set at the physical level, this architecture enables the future implementation of tailored codes, such as the surface-GKP code [12], exploiting transversal gates and magic state distillation for fault-tolerant universal photonic computing.*"

- We modified the sentence (lines 147-151) "Notably, this architecture lets us complete the final piece required for universal photonic computing" as

  "*Notably, this architecture enables us to fulfil the hardware requirements for a universal physical gate set with a nonlinear unit (NLU), incorporating a Kerr gate [15], coupled via the CU to the LOU, thus overcoming the limitations of solely linear operations.*"

- We have enhanced the final discussion to explicitly bridge our experimental demonstrations (particularly the quasi-deterministic GKP state generation) with the requirements of FTQC. Lines 380-399 explicitly discuss how our platform's capabilities lay the groundwork for tailored error-correcting codes, gate transversality, and magic state distillation:

  "*In this work, we demonstrated the generation of GKP states, laying the groundwork for producing high-quality GKP states characterised by large amplitudes and high squeezing levels. Clavina's approach to temporal encoding is naturally suited for generating large-scale GKP cluster states and constructing surface codes—a crucial requirement for concatenating with topological error-correcting codes to achieve fault tolerance. Consequently, our architecture offers a highly promising route toward active quantum error correction (see Supplementary Information Section VII.C for details).*

  To move toward fault-tolerant universality, the physical non-Gaussian resources demonstrated here must be integrated with tailored error-correcting codes to fully exploit gate transversality. Specifically, high-fidelity GKP states will serve as the foundational logical Pauli states, while our NLU will supply the essential non-Clifford physical operations required to generate and distill logical magic states, thereby completing the universal logical gate set that transversal operations alone cannot provide."

Supplemental Material:

Many references are missing in the text. It looks like the Latex compilation did not work properly.

We have fixed the compilation issue and all references are now properly displayed in the revised version.

It is apparent that the main text was written by a non-English speaker, which is perfectly ok. However, grammatical errors should be corrected, which is easily achieved by a native or proficient English speaker going through the text. Some examples:

Above eq.5 : "derivation of how the Kerr nonlinearity realized through photon"
Above eq.18: "Not that"
Section VII: "NON-GAUSSIAN STATE PREPARING"

Reply: In this revision, we have thoroughly checked and corrected the grammar and flow throughout the main text and Supplementary Information, in addition to the corrections highlighted by Referee 1:

- Above eq.5: In this section, we present a detailed derivation of how Kerr nonlinearity can be realised through photon addition and subtraction. (can be found on page 13)
- Above eq.18: Note that (can be found in page 16)
- Section VII: Non-Gaussian state preparation (can be found in page 21)

We believe that these revisions have substantially improved the quality of the manuscript and brought it to the standard required for publication.

Reply to Referee 2:

Manuscript review of: "Extensible universal photonic quantum computing with nonlinearity", by Shang Yu et. al, submitted for consideration to Nature Photonics.

The authors report a general-purpose photonic platform capable of performing a variety of quantum computing tasks using continuous-variable encoding. The architecture consists of an extensible unit in which a central computing core can be connected to multiple plug-in modules providing distinct functionalities, enabling the construction and execution of experiments in a modular, "piece-by-piece" fashion.

The system comprises a 100-mode, fully connected linear operation unit (LOU), a post-selected Kerr gate, sources of both Gaussian and non-Gaussian states, an inline squeezing gate, and multiple detection stages (SNSPDs, PNRDs, and homodyne detection). The architecture exploits time-bin encoding of optical qumodes, assisted by polarization for implementing both linear and nonlinear gates.

The versatility of the machine, termed "Clavina" by the authors, is demonstrated through several tasks, including the generation of an 8000-mode cluster state, the generation and breeding of cat states, the near-deterministic production of GKP states, and the simulation of many-body dynamics. The experimental results are solid, supported by numerical simulations, and thoroughly detailed in the Supplementary Materials.

The work integrates several elements that have previously been demonstrated as stand-alone experiments, such as the mechanism for generating and validating square-lattice cluster states, the Kerr gate, the inline squeezing gate, and the LOU (here presented in a boosted dual-core implementation), with new applications. The primary novelties include the use of nonlinear resources for near-deterministic GKP state generation and the implementation of non-Gaussian operations (notably the Kerr gate) for simulating the Bose–Hubbard Hamiltonian. Overall, I believe that the combination of these new results, the high fidelity of the passive operations, and the reconfigurability and versatility afforded by the extensible architecture meets the high standards of Nature Photonics.

Reply: We sincerely thank the referee for the thorough review of our manuscript and for highlighting our contributions and innovations. We have carefully considered the referee's comments and provided a detailed comparison between our approach to GKP state generation and the method reported by Larsen et al. Additionally, we have addressed all other comments, including those regarding long-term fidelity drift, further reduction of photon loss, the presentation of the inline squeezing and Kerr-gate figures, as well as the figure and reference issues in the Supplementary Information. In the following, we provide a point-by-point response to the reviewer's comments and suggestions, and the corresponding changes are highlighted in red.

However, I have one major concern that should, in my view, be addressed before the manuscript is considered for publication.

As a non-expert in continuous-variable quantum computation, I found it difficult to compare the approach to deterministic GKP state generation presented here with the strategy reported by Larsen, Mikkel V., et al., "Integrated photonic source of Gottesman–Kitaev–Preskill qubits," Nature 642, 587–591 (2025). To my understanding, the approach of Larsen et al. relies on post-selection of specific photon-number patterns in a subset of modes to herald the GKP state. As such, it is intrinsically probabilistic and requires multiplexing to enhance the success probability. Furthermore, their reported Wigner functions exhibit three lobes along both the x-and p-quadratures, which is claimed to be a necessary condition for fault tolerance.

In contrast, the present work starts from cat states deterministically generated via inline squeezing of single photons, which are in turn heralded with near unit probability by using time-multiplexing. Through interference of cat states and measurement feed-forward, the authors produce GKP states deterministically with three lobes along the p-quadrature and only one lobe along the x-quadrature. Subsequent post-selection on low-noise states renders the generation "quasi-deterministic."

I would therefore ask the authors to clarify the distinctions and commonalities between these two approaches, discussing their respective advantages and disadvantages, particularly in terms of loss tolerance and resource optimization. In addition, do the GKP states reported here exhibit the full multi-peak grid structure required for fault tolerance?

Reply: As the referee rightly points out, our approach bears key distinctions from that of Larsen *et al.*, which we clarify below.

Probability:
Xanadu's approach in Larsen *et al.* utilises Gaussian boson sampling, where squeezed states are interfered in a few-mode interferometer on a photonic integrated circuit (PIC), and photon-number-resolving (PNR) measurements are made on all but one output. These PNR measurements (Fock state projections) introduce the non-Gaussian resource required for GKP state engineering. The squeezing parameters of the initial states, the parameters of the interferometer, and the photon-number pattern are all optimised before the experiment is run to achieve an appropriate balance between the fidelity and the success rate of the desired post-selected GKP state. In theory, the probability of successfully generating the correct state is < 0.05. Their architecture relies on running many PICs simultaneously and using a network of switches (EOMs) to multiplex successfully generated states (a so-called GKP refinery).

In practice, this multiplexing was **not** performed with actual GKP states in Larsen *et al.*. Their experiment ran a single chip with a success probability of 2.9×10^{-4} and a 58 Hz success rate before discarding noisy outcomes.

In contrast, our approach introduces the non-Gaussian resources at the start of the process in the form of squeezed single-photon states acting as kittens. The generation of single photons here is inherently probabilistic, but we achieve near-deterministic preparation through our "boosting method," which effectively multiplexes successfully heralded single photons (see Supplementary Information Section II.C) that are then fed into the inline squeezing unit (ISU). Provided four squeezed single photons are present, the subsequent cat/GKP breeding process is deterministic (see our later reply for details) with Gaussian homodyne measurements at room temperature. The overall process becomes slightly probabilistic due to the near-deterministic nature of the initial heralded single-photon state production.

We believe our approach is more favorable in two key ways. First, it is easier to multiplex a supply of single photons by heralding with a single detector than it is to multiplex entire GKP states, as the latter requires many PICs and PNRDs. Second, temporal multiplexing is natural to Clavina, a fully time-bin-encoded architecture that lends itself well to generating large cluster states in time. Xanadu's approach requires a combination of path encoding on PICs, temporal multiplexing, and time-bin encoding for cluster state generation.

Size and quality:
The size and quality of the generated GKP states are determined by the amount of nonlinear

resources introduced and by the achievable squeezing and loss levels, respectively. The method of Larsen *et al.* can currently achieve a two-dimensional 3×3 grid conditioned on 3-3-3 photon sampling outcomes, which is equivalent to introducing nonlinear resources corresponding to nine photons. The advantage of their method is that high-photon-number sampling results can be readily post-selected to create a larger GKP state; however, its primary disadvantage is that the state can only be generated probabilistically with a very low success rate.

Our approach uses four squeezed single-photon states as small-amplitude cat states (kitten states) at the input, corresponding to nonlinear resources equivalent to four photons. We report the first demonstration of two rounds of cat breeding. Our method is highly scalable—more squeezed single-photons, and more rounds of breeding would result in full two-dimensional GKP states.

In the current implementation, we aim to demonstrate a photonic quantum computing platform that is multifunctional. Our singular apparatus can perform large-scale GBS, cluster-state generation, Cat/GKP state engineering, and Bose-Hubbard simulation. If one were to build a dedicated apparatus, with the sole purpose of producing one GKP state similar to Larsen *et al.*, then some loss could be mitigated to produce a slightly larger two-dimensional GKP state. Instead, we sacrifice some size in favour of fidelity, orders of magnitude gains in success rate, and demonstrations of other applications.

Fault-tolerance requirement:
Regarding fault tolerance, the states demonstrated in the present work are one-dimensional GKP states rather than full two-dimensional GKP code states. Even so, they already enable the detection and correction of displacement errors along the p-quadrature, thereby demonstrating an essential error-correction capability relevant to fault-tolerant quantum computing. Further afield, one-dimensional GKP states find utility in quantum sensing [K. Duivenvoorden et al., Physical Review A, 95(1), 012305 (2017) and C. Valahu, C. et al., Science Advances, 11(39), eadw9757 (2024)].

Although our GKP states do not yet exemplify the full two-dimensional grid structure required for a general fault-tolerant GKP qubit, our hardware architecture is inherently compatible with the generation of such two-dimensional GKP states and provides a clear path toward this goal. For example, one can use eight squeezed single-photon states as input to generate two one-dimensional GKP states, and then interfere these states with appropriate feedback to obtain a two-dimensional GKP state with multiple lobes along both the x- and p-quadratures. We therefore view the present demonstration as an important step toward fault-tolerant photonic quantum computing.

In this revision, we have added this discussion to the Methods section ("GKP state preparation") on page 10, and in the Supplementary Information on pages 25-26.

Additional questions and comments

The authors state that "the next challenge on the path to a large-scale practical quantum photonic computer is suppressing photon loss and other noise mechanisms." What strategies do the authors envision for further reducing losses? The transmissivities of the optical opponents reported in the Supplementary Materials already appear extremely high, with seemingly limited room for further improvement.

Reply: We thank the reviewer for highlighting the high component transmissivities achieved in our current setup. While the transmittance is indeed already very high by conventional optics standards for individual optical components, there remains significant room to improve the *cumulative* system efficiency.

In the near term, one strategy is to adopt higher-bandwidth electro-optic modulators (EOMs), which would allow us to increase the repetition rate of the sources and shorten the temporal period between time bins. Doing so allows us to shorten the Herriott cells and replace fiber delays with purely free-space delays, thereby negating fiber-coupling losses. Furthermore, advanced photodetectors with efficiencies up to 99% and SNSPDs and PNRDs with detection efficiencies exceeding 95% are continuing to mature and become commercially available. Integrating these will directly reduce readout losses.

In the long term, we envision the Clavina architecture translating seamlessly to integrated photonic circuits. In particular, thin-film lithium niobate (TFLN) platforms can achieve modulation at rates of 100–200 GHz, corresponding to temporal delays on the order of picoseconds or path lengths on the order of millimeters. With propagation losses currently reaching levels of 0.01 dB/cm and continuing to improve, monolithically integrating the ISU and NLU onto a single chip would substantially mitigate the coupling losses that currently exist between bulk-optic units.

In this revision, we expanded the discussion on loss in the "Discussion and outlook" section (page 8, line 404-406), and Methods (page 9-10, line 628-647)

In Supplementary Section VII-A, the authors state that when interfering two cat states, breeding succeeds when no photons are detected in the heralding arm. Why is this stage not considered probabilistic, given that it relies on post-selection of a specific measurement outcome?

Reply: We would like to clarify that cat-state breeding is indeed a probabilistic process. In that protocol, successful breeding is heralded by the absence of detected photons in the heralding arm. By contrast, our GKP-state generation procedure follows a different process to cat-state breeding. Instead of conditioning on a vacuum event in the heralding arm, we perform homodyne measurement on the heralding mode and use the measurement outcome to apply a rapid feedforward displacement to the signal mode. In this sense, once the required input resource states are prepared, the final step of our GKP-state generation is deterministic rather than probabilistic.

In this revision, we have rewritten the corresponding sentence in Supplementary Information Section VII.A as: "*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 probabilistic yields, in the other output port, a cat state with an enlarged coherent-state separation, as displayed in Fig. S15****b****,****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 Fig. S15****c****,****f****.*"

These modifications can be found in Supplementary Information, Page 21-22.

What causes the fidelity drift observed in Fig. 2(b), given that all optical paths in the LOU are actively stabilized? Moreover, does reusing the hardware within the LOU introduce systematic biases in the mode transformations as the optical depth increases (e.g., polarization-dependent losses or small waveplate misalignments)?

Reply: In our experiment, the observed fidelity reduction over extended operation is likely not caused by a single dominant mechanism, but rather by the cumulative effect of several small drifts and imperfections. These include a slight decrease in coupling efficiency, a modest degradation of the interference extinction ratio in the LOU even with active phase locking enabled, small drifts in the EOM driver accuracy during long runs possibly associated with thermal effects, and slight variations in the SNSPD detection efficiency over time, which may also be related to heating. While each of these effects is relatively minor, their combined impact can become non-negligible during prolonged measurements.

In this revision, we briefly summarised our hardware in the Methods (Page 9, line 617-627), and added the above discussion as:

"*Over extended operation, the observed fidelity reduction (as shown in Fig. 2**b**) is likely caused by the cumulative effect of several minor drifts and imperfections. These include slight reductions in coupling efficiency, modest degradation of the interference extinction ratio in the LOU (despite active phase locking), small drifts in EOM driver accuracy during long runs (possibly due to thermal effects), and slight variations in SNSPD detection efficiency over time. Although each of these effects is minor in isolation, their combined impact can become noticeable during longer measurements.*"

In Fig. 2(d), partitions A and B, as well as the meaning of the blue curve (individual quadratures x_k and p_k?), are not specified in the main text.

Reply: The partitions A and B corresponds to the odd and even squeezed input, as the red pluses (corresponds to A) and blue pluses (corresponds to B) shown in Fig. 1**b**, rightmost panel. The blue curve denote the measured variances of shot noise.

In this revision, we added the explanation in the caption of Fig. 2**d** and **e** as (page 4): "*The quadrature correlations of the cluster state indicate strong x- and p-quadrature correlations between modes A and B, corresponding to the odd and even input states shown in the rightmost panel of Fig. 1**b**, across the first 200 time-bin modes. The lower measurements (orange) show that the nullifier variances remain consistently below the -3 dB inseparability threshold over 8,000 modes, while the reference shot-noise variances are shown in blue, confirming sustained multimode entanglement in both quadratures.*"

In Fig. 2(c), what is the size of the cliques (i.e., how many photons determine the size of the induced subgraph) considered in the analysis?

Reply: In the GBS experiment and the participation coefficient analysis, we considered eight-cliques identified from sampling events containing more than three photons. In the revised manuscript, we have added this information, which can be found in the main text, caption of Fig. 2**c**, Page 4.

The conceptual schematics of the inline squeezing gate and the Kerr gate currently shown in Fig. 2(f) would be better placed in Figs. 3 and 4, respectively. Figure 2 should focus exclusively on linear optical operations. Additionally, the Wigner functions shown in the Kerr-gate panel are unclear: what is the input state on which the gate acts? What is the conceptual meaning of the rotated Gaussians shown in red?

Reply: We have moved the conceptual schematic of the inline squeezing unit from Fig. 2**f** (upper panel) to Fig. 3**a**, and the conceptual schematic of the nonlinear unit from Fig. 2**f** (bottom panel) to Fig. 4**a**. Fig. 2 now focuses exclusively on the linear optical operations, which makes the presentation clearer.

We have also revised the Kerr interaction panel and its caption to clarify this nonlinear process. The Wigner function of the input state (that is, the coherent state) is now shown at the beginning, and the Kerr interaction with different nonlinearity ($\Phi = \pi/15$ and $\pi/2$) are explicitly specified. In addition, the rotated Gaussian shapes shown in red are now clearly express that the zero- and one-photon input cases do not acquire a phase shift, whereas only the two-photon input case acquires a phase shift with angle $\Phi = \kappa t$, which directly exhibits the nonlinearity of this measurement-based Kerr gate.

In this revised version, the caption of Fig. 4**a** now reads:
"*A measurement-induced Kerr interaction ($e^{-i\hat{H}t/\hbar}$) [14], with Hamiltonian $\hat{H} = \hbar\kappa\hat{n}(\hat{n}-1)$ (where $\hat{n}$ is the bosonic number operator), directly introduces nonlinearity by imparting a phase shift of $\Phi = \kappa t$ to the two-photon component relative to the zero- and one-photon components, as evidenced by the Wigner negativity (blue regions) of the output state.*", and can be found in Page 6.

We have updated the corresponding figure and caption to make this point explicit and to avoid ambiguity.

The authors should explicitly state in the main text that only three modes are considered in the Bose–Hubbard Hamiltonian simulation shown in Fig. 4(c). Maybe I missed it, but this information seems currently to appear only in the Supplementary Materials.

Reply: We have made the following revision: "*As a demonstration, we simulate a three-site Bose-Hubbard model using a three-mode version of the circuit in Fig. 4***c***...*" in this revision, and can be found in main text page 6-7, line 291-293.

In Fig. 4, the descriptions of panels (a) and (b) appear to be swapped.

Reply: As noted in our response to Referee 1, we have split Fig. 2**f** into two, and relocated the sub-figures to Fig. 3**a** and Fig. 4**a**, to better introduce the applications that utilise each unit. We have modified to the captions to match the correct panels. The revision can be found in the main text, on page 4, 5, and 6.

The sentence "Weak interactions keep the photons apart, while increasing U/J favors rebunching into doubly occupied states" seems inconsistent with Fig. 4(c), which shows the opposite trend for the |110> state.

Reply: Indeed, the sentence quoted by the referee is the result for $|200\rangle$, while the result for $|110\rangle$ is reversed. In this revision, we revised the whole sentence as:

"*Starting from the bunched state $|200\rangle$ at a fixed evolution time of $t = 0.5$, weak interactions populate both bunched and separated configurations, whereas increasing U/J stabilises double occupancy and suppresses singly occupied states. The trend is reversed for the initial state $|110\rangle$, where weak interactions favour doubly occupied states, while increasing U/J keeps the photons apart.*"

We also carefully checked the results descriptions to minimise these types of inconsistencies. The corresponding modifications can be found on page 7, line 306-313.

In Fig. 4(a), what does μ represent? I could not find its definition in either the main text or the Supplementary Materials.

Reply: Here, μ denotes the on-site potential term in the Bose-Hubbard model, which is implemented experimentally as a local phase shift. In our work, for simplicity, we set $\mu = 0$. The caption of Fig. 4**c** has now been revised to reflect this.

Several references are missing in the Supplementary Materials and appear as "?".

Reply: We have corrected the reference citations and the compilation issue, and all references are now properly displayed in the revised version.

Reply to Referee 1:

I was surprised by the statement that "326 electro-optic modulation, the circuit can be reconfigured on the fly, making both the couplings Jij and the on-site interactions U tunable in time. Such capabilities are typically unavailable in superconducting qubit architectures". I would have thought that superconducting architectures are at least as versatile as photonic ones, in particular when it comes to tuning interactions (via fluxes). The fact that they are not modified on-the-fly responds more to the fact that photonic computation happens as a stream of photons rather than on a physical support.

Reply: We thank the referee for this useful comment. First, we explain why reconfiguration is difficult to realise in superconducting transmon qubits. As demonstrated in [A. H. Karamlou et al. *Nature* 629, 561-566 (2024)], capacitively coupled superconducting transmon qubits are typically fabricated using a flip-chip process, and the native Hamiltonian is given by:

$$\hat{H} = \sum_i \epsilon_i \hat{n}_i + \sum_i \frac{U_i}{2} \hat{n}_i(\hat{n}_i - 1) + \sum_{\langle i,j \rangle} J_{ij} \hat{a}_i^\dagger \hat{a}_j, \tag{1}$$

where $\hat{a}_i^\dagger$ ($\hat{a}_i$) is the creation (annihilation) operator for excitations at site i, and $\hat{n}_i = \hat{a}_i^\dagger \hat{a}_i$. The first term represents the site energies, which can be tuned within a certain range. However, the second term (i.e., the on-site interaction) is typically tied to the fixed anharmonicity of transmon qubit i, making it less naturally reconfigurable in real time. In [A. H. Karamlou et al. *Nature* 629. 561-566 (2024), Z. Yan et al. *Science* 364, 753-756 (2019)], for example, the anharmonicity cannot be freely adjusted once the transmon qubits are fabricated. Specifically, in [A. H. Karamlou et al. *Nature* 629, 561-566 (2024)], $U = -2\pi \times 218$MHz, which is significantly larger than J and ϵ (where $J \sim 5$ MHz).

According to the referee's comment, we have revised the corresponding text as: "Benefiting from fast electro-optic modulation, the circuit can be reconfigured on the fly, making both the couplings J_{ij} and the on-site interactions U tunable in time. Such capabilities are typically challenging to realize in superconducting-qubit architectures [44], where U is generally fixed at a large value and cannot be tuned, thereby restricting simulations to the hard-core boson limit (within a perturbative regime)."

This can be found in the left column, on page 4.

As a minor comment, the ISU acronym is defined in the caption of a figure, rather than in the main text.

Reply: We have now included the definition of ISU in the main text, which can be found in the left column, on page 3.

Reply to Referee 2:

The authors have adequately addressed all my comments and provided the necessary amendments in the newly submitted version of the manuscript. In particular, they added to the Methods section a fair comparison between their approach to deterministic GKP generation and the one presented by the Xanadu group (side note: the authors should correct a few typos between lines 653–656). They also included missing information in the figure captions and reorganized their content. Each figure is now properly referenced in the main text, and the captions are self-explanatory.

Reply: We appreciate the referee's recognition of our work. The typos between lines 653-656 have been corrected and can now be found on pages 8-9, in the left column.