

A manufacturable platform for photonic quantum computing

<https://doi.org/10.1038/s41586-025-08820-7>

PsiQuantum team*

Received: 30 July 2024

Accepted: 20 February 2025

Published online: 26 February 2025

Open access

 Check for updates

Although holding great promise for low noise, ease of operation and networking¹, useful photonic quantum computing has been precluded by the need for beyond-state-of-the-art components, manufactured by the millions^{2–6}. Here we introduce a manufacturable platform⁷ for quantum computing with photons. We benchmark a set of monolithically integrated silicon-photonics-based modules to generate, manipulate, network and detect heralded photonic qubits, demonstrating dual-rail photonic qubits with $99.98\% \pm 0.01\%$ state preparation and measurement fidelity, Hong–Ou–Mandel (HOM) quantum interference between independent photon sources with $99.50\% \pm 0.25\%$ visibility, two-qubit fusion with $99.22\% \pm 0.12\%$ fidelity and a chip-to-chip qubit interconnect with $99.72\% \pm 0.04\%$ fidelity, conditional on photon detection and not accounting for loss. We preview a selection of next-generation technologies: low-loss silicon nitride (SiN) waveguides and components to address loss, as well as fabrication-tolerant photon sources, high-efficiency photon-number-resolving detectors (PNRDs), low-loss chip-to-fibre coupling and barium titanate (BTO) electro-optic phase shifters for high-performance fast switching.

It has long been understood that useful quantum computers will require error correction for fault-tolerant operation and, therefore, on the order of millions of physical qubits⁸. Owing to their intrinsic low-noise properties, photons have been used to implement many of the foundational demonstrations of superposition, entanglement, logic gates, algorithms and so on¹. However, large-scale photonic quantum computing has so far been precluded by several outstanding and challenging requirements.

Since the earliest proposals for fault-tolerant optical quantum computers^{2–6}, it has been clear that a very large number of photonic components would be required for any useful system^{9,10}. Furthermore, to satisfy the requirements of error-correcting codes, these components should also perform beyond the state of the art of conventional integrated photonics^{9,10} and must also extend outside the scope of a typical photonics library, introducing non-standard devices—most notably high-efficiency single-photon detectors^{11,12}. The need for a very large number of near-identical devices motivates an emphasis on fabrication using conventional, high-volume semiconductor manufacturing processes⁷. Finally, these devices must be integrated in an extensive system—demanding fast control electronics, high-power cryogenic cooling to support the operation of superconducting detectors and low-loss, high-fidelity networking of qubits between modules.

In this paper, we describe a technology stack and basic building blocks for photonic quantum computing, demonstrating the crucial functionalities of qubit generation, manipulation, detection and networking, including single-photon sources, waveguide-integrated superconducting single-photon detectors, single-qubit state preparation and measurement (SPAM), chip-to-chip qubit interconnects, two-photon quantum interference and two-qubit fusion, all at telecommunications (C band) wavelengths. These constitute the basic operations required

for most approaches to photonic quantum computing^{2–6,9}, including fusion-based quantum computing (FBQC; recently introduced in ref. 13). These components are fabricated in a commercial semiconductor foundry¹⁴, using a fully integrated 300-mm silicon photonics process flow, with all operations on-chip.

To enter the fault-tolerant regime of operation will require a technology stack with improved component performance and further functionality. To this end, we have developed and present next-generation components, with SiN waveguide losses as low as $0.5 \pm 0.3 \text{ dB m}^{-1}$, splitters and crossings with $0.5 \pm 0.2 \text{ mdB}$ and $1.2 \pm 0.4 \text{ mdB}$ loss, respectively, and fibre-to-chip coupling losses as low as $52 \pm 12 \text{ mdB}$. The quantum benchmarking results presented here are conditional on photon detection and the photon production is heralded but non-deterministic. To overcome non-determinism in photonic quantum computation, a fast optical switch for multiplexing is required^{15–18}. Here we introduce BTO switches into our technology for this purpose, with a loss-voltage product of $0.33 \pm 0.02 \text{ dB.V}$. Furthermore, we demonstrate robust photon sources capable of indistinguishable photon generation over a $\pm 400\text{-pm}$ resonance shift and waveguide-integrated photon-number-resolution detectors with 98.9% single-photon median efficiency and up to four-photon resolution. When taken together, these new and improved components constitute a feature-complete set of photonic building blocks having all of the functionality necessary to enable future fault-tolerant photonic quantum computing systems.

Technology stack and building blocks

Silicon photonics is a mature manufacturing technology, built on decades of industrial development for established applications in the communications, medical and automotive sectors^{19,20}.

*A list of authors and their affiliations appear at the end of the paper.

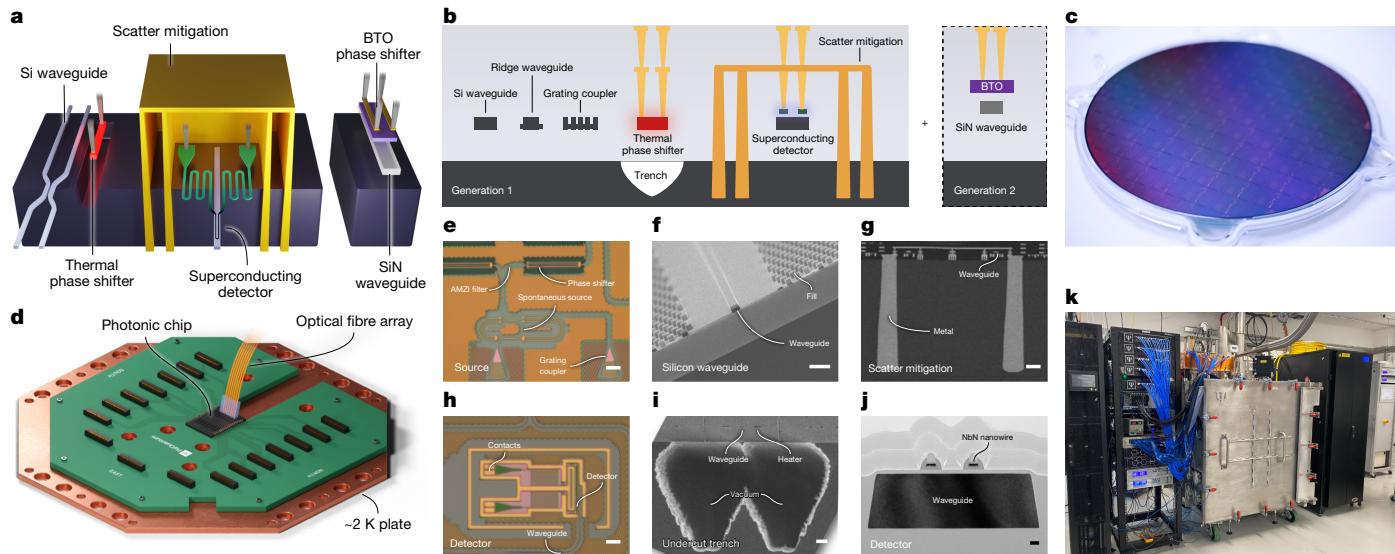


Fig. 1 | Manufacturable integrated quantum photonic stack. **a, b**, Schematics of key components and process modules. We highlight (on the right) further process steps included in our next-generation platform. **c**, A 300-mm wafer containing single-photon sources, superconducting single-photon detectors and quantum benchmarking circuits. **d**, An acryogenic assembly containing a photonic die, heat spreader, electronic PCB and 100-channel telecommunications fibre attach unit. **e–j**, Optical micrograph, scanning electron microscope or transmission electron microscopy images of: photon

source (top-down) (e); optical waveguide (cross-section) (f); deep/shallow trench scattered light shield (cross-section) (g); single-photon detector (top-down) (h); thermal isolation trench (cross-section) (i); single-photon detector on waveguide (cross-section) (j). **k**, Custom cryostat used in benchmarking experiments with >10 W cooling power at 2.2 K. Scale bars, 20 μm (e,h), 1 μm (f,i), 10 μm (g), 40 nm (j). AMZI, asymmetric Mach–Zehnder interferometer.

We modified an established silicon photonics manufacturing flow to include high-performance single-photon detection and photon-pair generation (Fig. 1). To our knowledge, this is the first realization of an integrated photonic technology platform capable of on-chip generation, manipulation and detection of photonic qubits.

Our baseline quantum photonic technology stack was developed in partnership with GlobalFoundries and is fabricated in their 300-mm, state-of-the-art, high-volume semiconductor foundry. By making use of industrial unit process steps from semiconductor manufacturing in combination with foundry design services, such as optical proximity correction and optimized process design rules, the technology inherits the scalability and performance of a high-volume commercial environment. The manufacturing flow includes more than 20 photolithography levels and hundreds of processing and in-line measurement steps. Critical process modules developed include passive silicon-on-insulator photonic waveguides, a niobium nitride (NbN) superconducting layer for single-photon detection, deep metal-filled trenches for optical noise reduction, resistive heaters for phase control and optical circuit reconfigurability, grating couplers for optical input/output, back-end-of-line copper electrical interconnects and aluminium redistribution layers.

Using this stack, we build quantum photonic integrated circuits using standard silicon photonic waveguide components, including directional couplers, crossings and thermal phase shifters. We combine these components to produce key building blocks: high-fidelity spontaneous photon-pair sources; interferometers for circuit reconfigurability, qubit manipulation and filtering; and waveguide-integrated single-photon detectors (Fig. 2a). We now outline the performance of each of these building blocks.

Photon sources

To construct entangled resource states and, in turn, an error-correcting code, photonic quantum computers consume many single photons, which must be generated with high efficiency, well-defined timing and a high repetition rate, whilst also being spectrally pure and

indistinguishable²¹. Our single-photon sources use spontaneous four-wave mixing (SFWM)²² driven by a pulsed laser pump, for which the generation of a single photon is probabilistic but heralded by the detection of its pair—making a heralded single-photon source (HSPS).

The visibility of two-photon quantum interference, a key operation in photonic quantum computation, is limited by the spectral purity of the heralding single photons, which is determined by the joint spectral intensity of the photon pairs. We use resonator-based optical waveguide structures to tailor the spectral properties of the photon sources to achieve high spectral purity. The pump is aligned to a resonance frequency and single photons are generated at resonant frequencies spaced symmetrically around it, as illustrated with shaded bands in Fig. 2c. Single-ring resonator sources are intrinsically limited to a heralded photon purity of about 93% (ref. 23). We circumvent the spectral purity limitation of single resonator sources using interferometrically coupled resonator designs²³, which we characterized to have a measured spectral purity of $99.5\% \pm 0.1\%$ without spectral filtering (Fig. 2b and Supplementary Information).

Photon detection

Photonic quantum computing relies on heralding the creation of quantum states by the detection of correlated photons. Examples include single-photon heralding from pair sources, heralded probabilistic resource state generation and fusion measurements. For fault tolerance, these functions require near-unit-efficiency single-photon detection. We introduced a NbN layer into our photonic stack to enable high-performance manufacturable superconducting nanowire single-photon detectors (SNSPDs)^{11,12}.

We use a hairpin-shaped SNSPD design²⁴, as depicted in Fig. 2a, with a film thickness of roughly 5 nm, nanowire width of about 90 nm and detector length of ≥ 80 μm . When operated at approximately 2 K temperature, these detectors exhibit clear plateaus in the photon count rate versus bias current (Fig. 2e and wafer maps in the Supplementary Information), indicating high internal detection efficiency. The on-chip detection efficiency is measured through cryogenic electro-optical

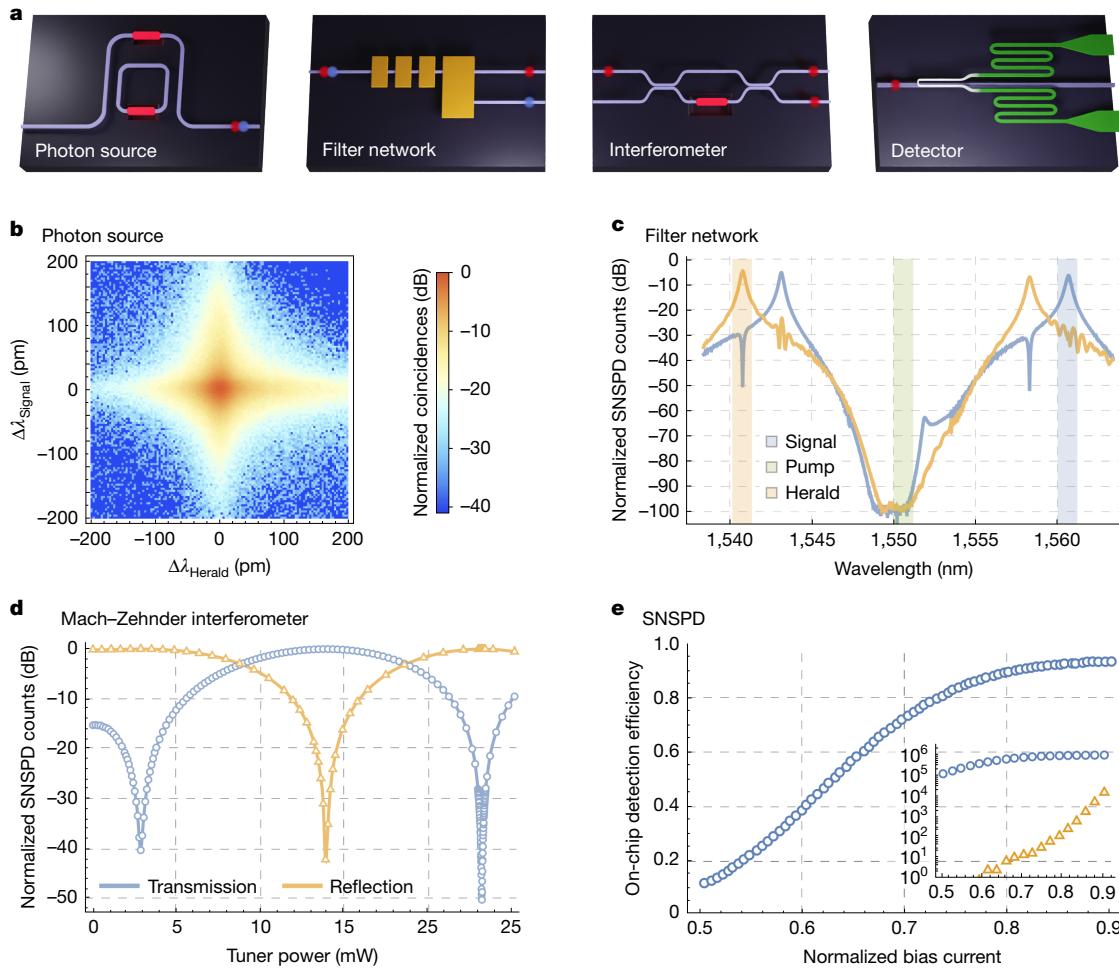


Fig. 2 | Key building blocks of the platform. **a**, Schematics of photon source, filter network, interferometer and detector. **b**, Measured joint spectral intensity of an interferometrically coupled resonator photon source, indicating a spectral purity of 99.5% (Supplementary Information). **c**, Response of our pump filter network. We shade the pump, signal and herald frequency bands and show the measured herald (orange) and signal (blue) filter spectrum, characterized with on-chip SNSPDs. **d**, Measured response of a Mach–Zehnder interferometer to heralded single-photon illumination on a fully integrated platform.

The extinction ratio at the transmission port is >50 dB. The asymmetry in the Mach–Zehnder interferometer response is an artefact of a non-constant step size, which is finer around one feature only. There is no marked variation in performance across a circuit or among different circuits. **e**, Measured on-chip detection efficiency as a function of detector bias current (I_b) normalized by the detector switching current (I_{sw}) and the detector count rate (blue) and dark count rate (orange) per second (inset) (Supplementary Information).

measurements of waveguide-integrated SNSPDs (Supplementary Information). Testing of screened SNSPDs yielded a median on-chip efficiency of 93.4% and average value of $88.9\% \pm 3.5\%$ (Supplementary Information), limited by the hairpin design of the detector.

Interferometers and filters

Interferometers are a key building block of integrated photonic quantum computing, enabling qubit state preparation and projection, pump filtering, switching networks, resource state generation and fusion measurements. We use combinations of directional couplers, crossings and rings to construct ring resonators and Mach–Zehnder-type interferometers. These components have been optimized through design–test cycles and provide predictable performance guaranteed by strict fabrication process control. An example high-contrast Mach–Zehnder interference fringe, measured with co-integrated HSPS and SNSPDs, is shown in Fig. 2d, with a >50 dB extinction ratio.

Such passive circuits are reconfigurable at low frequencies using thermal phase shifters, which are commonplace in silicon photonics. Although the circuit itself is cooled to cryogenic temperature (to support the operation of integrated superconducting single-photon detectors), the thermal phase shifters reach local temperatures well in

excess of 100 °C during operation. Given the available cooling power per unit area at about 2 K cryogenic temperatures, thermal insulation of the phase shifter—using undercut regions etched from the silicon substrate (Fig. 1a,b,i)—is critical to achieve sufficient efficiency. Most if not all of these heaters are made redundant and will ultimately be removed thanks to the second-generation technologies described later in this manuscript—in particular, the fabrication-tolerant source, which will not require tuning, as well as the electro-optic phase shifter, which can operate at GHz rates.

Integrated heralded single-photon generation and quantum benchmarking circuits

So far, photonic quantum computing platforms have depended on off-chip single-photon sources, off-chip single-photon detectors or both. Although sufficient for demonstration purposes, it is very challenging to achieve the heralding efficiency and component density required for practical fault-tolerant quantum computing without co-integration of the source, filter and heralding detector. Through integration of our key building blocks into our semiconductor platform, we have developed the world's first fully integrated

Table 1 | Single-qubit and two-qubit performance metrics. Not accounting for loss

Metric	Experiment value (%)
Single-qubit	SPAM fidelity 99.98±0.01
	99.996±0.003*
Two-qubit	Chip-to-chip fidelity 99.72±0.04
	Quantum interference visibility 99.50±0.25
	Bell fidelity 99.22±0.12

*Second SPAM fidelity listed above is measured with bright light and off-chip detectors; see main text.

HSPS—including source, filtering and heralding on the same chip. Using this, we construct benchmarking quantum circuits to quantify single-qubit, two-qubit and chip-to-chip qubit interconnect performance, which is summarized in Table 1.

We selected photonic dies from 300-mm wafers using high-volume in-line and end-of-line electric, optical and electro-optical room-temperature testing, as well as cryogenic electro-optic testing for select parts. For our most complex systems, we package these dies into assemblies (Fig. 1d) together with thermal heat sinks, more than 1,000 electrical connections and up to 200 optical input/output. We house these packages in cryostats with approximately 2 K base temperature and up to 20 W cooling capacity (Fig. 1k).

HSPS

A high-performance HSFS requires engineered SFWM sources, heralding detectors, as well as a high-performance filter network on-chip, which we now describe. To separate the bright laser pump from the single photons, we require about 100 dB suppression of the pump photons. To achieve this in an integrated circuit, we combine both interferometric in-guide filtering and shielding of the detectors from out-of-guide scattered pump light. In-guide filtering uses a series of first-order and third-order asymmetric Mach–Zehnder interferometers combined with add-drop resonators to select single-source resonances for the herald and the signal photons. Optimizing the free spectral range and coupling values of each element, we achieve pump rejection of 99.1 ± 1.2 dB (Fig. 2c and Supplementary Information) and the simultaneous rejection of unwanted broadband parametric processes. The signal and herald photons are transmitted through filter networks with approximately 1 dB of loss. To suppress scattered light, we locally shield the detectors by encasing them in metal (Fig. 1a,b,g). The shields are constructed from deep and shallow metal-filled trenches and back-end-of-line metals. We observe approximately 115 dB pump power suppression between the pump input and the SNSPDs.

The integrated filters and scattered light shielding, combined with co-integration of SFWM and SNSPDs, allowed for the first demonstration, to our knowledge, of successful on-chip integrated heralded single-photon production, with coincidences-to-accidentals ratios²⁵ of up to 3,000 (Supplementary Information).

SPAM

We prepare a path-encoded qubit^{26,27} using a heralded photon and two-mode interferometers, as illustrated in Fig. 3a. We measure the path-encoded qubit using a two-mode interferometer and SNSPDs. The state of the single photon in two optical modes is controlled by two thermal phase shifters, which enable the encoding of arbitrary qubit states. We observe an average SPAM fidelity of $99.98\% \pm 0.01\%$ (Fig. 3e), conditional on the photon being detected (Supplementary Information). Aiming to separate the impact of the HSFS's signal-to-noise ratio, we repeat the measurement on a different but equivalent chip, using bright coherent light and off-chip photodetectors, achieving a fidelity of

$99.996\% \pm 0.003\%$ (Supplementary Information), showing that higher SPAM fidelity will be possible with improved HSFS signal-to-noise ratio.

Chip-to-chip qubit interconnect

Networking of quantum modules has seen growing interest as various technologies seek to scale beyond the boundary of a single chip, trap or reticle. Telecommunications-wavelength photonic qubits are naturally suited for transmission through optical fibre, without the need for quantum transduction²⁸. Furthermore, optical-fibre-based networking can enable further new functionality, such as interleaving²⁹ and active-volume compilation³⁰, leading to large resource savings for fault-tolerant algorithms. To demonstrate the networking capability of our photonic qubits, we build a point-to-point qubit network (Fig. 3b) and assess the fidelity of qubits after propagating between modules. We prepare high-fidelity single-qubit states using the same qubit state preparation circuit as described above and convert to polarization encoding using a two-dimensional grating-coupler-based path-to-polarization converter³¹. We transmit the qubit over 42 m of standard telecommunications-grade optical fibre, before converting to path encoding at the receiving module and performing on-chip qubit state measurement. The transmission and receiving modules both use on-chip superconducting detectors and operate at liquid helium temperature. We determined the Pauli transfer matrix³² fidelity between the physical channel and the identity operation, conditional on photon arrival, to be $99.72\% \pm 0.04\%$ (Fig. 3g and Supplementary Information). The system exhibits high loss associated with fibre-to-chip coupling by grating couplers (about 3 dB loss), which will be overcome in future systems using edge-coupled devices (discussed below).

Two-photon quantum interference

To benchmark our integrated single-photon sources, we measure HOM quantum interference between heralded photons from two independent sources integrated on the same chip (Fig. 3c). The measured visibility depends on many factors, including indistinguishability, spectral purity, number purity, signal-to-noise ratio and system detection efficiency. To control these, we implement a single system that integrates the technologies described above: high-purity, tunable photon-pair sources; high extinction filter network; and high-efficiency and shielded SNSPDs.

The on-chip HOM quantum interference between heralded photons from different sources, without substantial spectral filtering, was $99.50\% \pm 0.25\%$ (Fig. 3f), which—to our knowledge—is the highest measured in any platform. The experiment was performed at a pump repetition rate of 125 MHz, with a source coincidences-to-accidentals ratio of 929 ± 4 , a heralded $g^{(2)}(0) = 0.00358 \pm 0.00024$ and a maximum Klyshko efficiency of approximately 26% (Supplementary Information).

Two-qubit fusion

Bell fusion is a projective measurement onto two-qubit Bell states and is the prototypical example of the class of measurements that underpins the FBQC model¹³. We implement Bell fusion using type II fusion measurements⁶ on dual-rail qubits. Type II fusion uses a four-mode linear optical circuit followed by photon detection. It requires both single-qubit interference and interference between qubits, enabled by high-performance qubit preparations and high-visibility two-photon quantum interference, respectively.

We demonstrate that the fusion operation can perform a high-fidelity projection onto a Bell state, using the benchmarking circuit in Fig. 3d. Two independent path-encoded single qubits are prepared in the product state $|+-\rangle$. Using a reconfigurable fusion-measurement network, paths are then exchanged between the qubits and the resulting state is measured through single-qubit measurements. When a photon is detected in each pair of detectors, we measure a fidelity of $99.22\% \pm 0.12\%$ with the ideal Bell state. The density matrix is shown in Fig. 3h.

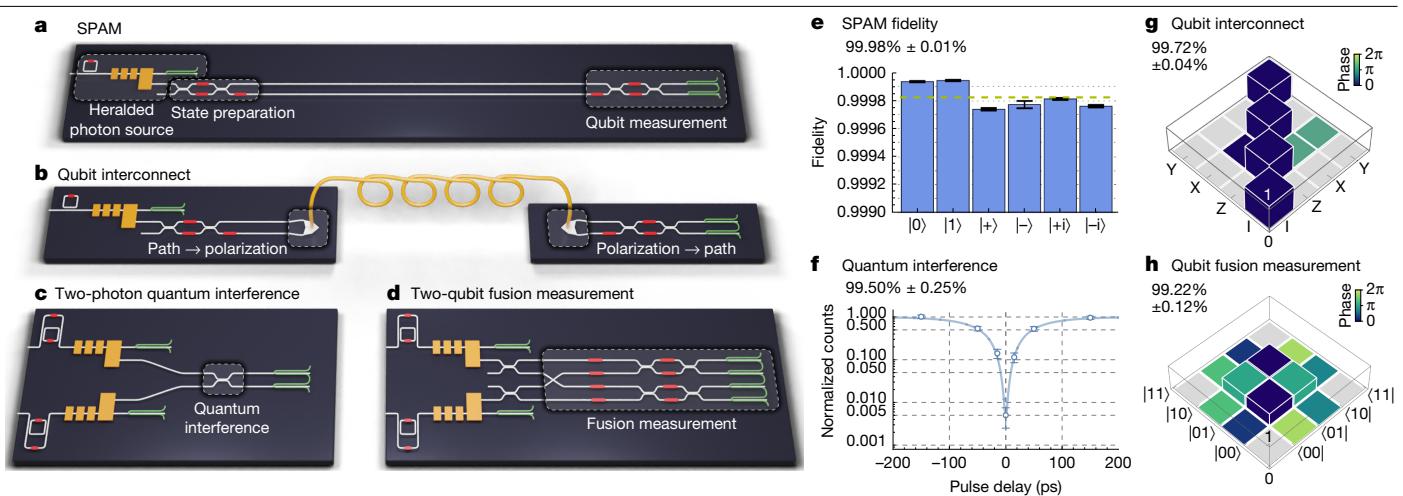


Fig. 3 | Quantum benchmarking circuits. These circuits are reconfigurable by means of thermal phase shifters indicated in red in the schematics. **a–d**, Schematics of: quantum state preparation and measurement (**a**); point-to-point qubit network (**b**); two-photon quantum (HOM) interference (**c**); two-qubit fusion measurement (**d**). **e**, SPAM fidelity of the reconstructed

state with the target state for Pauli eigenstates. **f**, HOM interference. **g**, Measured Pauli transfer matrix³² of chip-to-chip qubit interconnect channel. **h**, Reconstructed two-qubit density matrix after fusion (grey bars indicate magnitude below the 0.01 threshold).

Next-generation technologies

The performance of the baseline technology described above is still not sufficient for useful photonic quantum computing. In particular, silicon waveguides incur too much propagation loss for fault tolerance, photon sources require complex and power-hungry tuning and high-speed optical switching is unavoidably necessary to overcome the intrinsic non-determinism of the spontaneous single-photon sources.

We now describe some of the critical developments towards higher performance and further functionality in our next-generation-technology platforms, derived from several process flows. We focus on advanced photon sources, high-efficiency photon-number-resolving detection, low-loss waveguides, high-efficiency fibre-to-chip coupling and on-chip electro-optic phase shifters.

Cascaded resonator source

The key performance metrics for photon sources are two-photon interference visibility and photon efficiency. However, there are other characteristics that must be addressed to enable the operation of devices at the scale of useful quantum computers. Two important considerations are the pump power required to drive the SFWM process and the thermal power dissipated at cryogenic temperatures to control and tune the source. We have implemented a cascaded resonator source that addresses these aspects simultaneously.

The source comprises several integrated resonators coupled to a single bus waveguide (Fig. 4a). Through joint optimization of the resonator–bus coupling, the resonance wavelengths and the pump spectral amplitude, the joint spectral intensity of the source can be engineered. Our 24-resonator device has a measured upper-bounded purity of 99.35% (Supplementary Information), assuming flat spectral phase (Fig. 4b), whilst using an order of magnitude less pump power than the interferometrically coupled source design. Even with the lower n_2 of SiN, our optimized cascaded resonator source achieves 5% pair probability with approximately 100 pJ of pump pulse energy, which is within the range of scalable erbium amplifiers at GHz repetition rates.

This cascaded resonator source addresses indistinguishability in a new way. The spectrum of the photon pairs is fixed by the pump wavelength and not by the resonant wavelength of the device. Thus, global resonance shifts (for example, from fabrication variations) have minimal impact on the spectral indistinguishability of photons generated from different devices. Figure 4c shows the measured

indistinguishability between two sources as a function of resonance shift (Supplementary Information). Using thermal tuners, we aligned two devices to the optimal operating point and applied a controlled global resonance shift to one cascaded resonator source, to simulate the impact of fabrication variation. In this implementation, we achieve $>99\%$ two-source indistinguishability over a ± 400 -pm resonance shift window, compared with less than ± 40 pm for a single-ring source. The built-in tolerance of the cascaded resonator source to device-to-device global wavelength variation, together with state-of-the-art fabrication control, can enable tunerless indistinguishable photon sources.

PNRDs

The waveguide-integrated manufacturable single-photon detectors presented earlier, although transformative, lack the photon-number-resolving capability required for FBQC. The ability to distinguish low photon numbers in detection, and to herald on that information, allows for both the removal of higher-order photon number states generated in SFWM sources and the identification of unwanted events in fusion-based entangled state generation and computation²¹.

Spatial multiplexing³³ of many SNSPD-like detector elements, as shown in Fig. 4d, can be used to assemble a scalable detector with effective photon-number resolution. In these PNRDs, the number of detected photons is approximately proportional to the amplitude of the detector output voltage. To validate this concept, we have produced waveguide-integrated PNRDs with 4 and 5 unit cells, with the best performing designs yielding on-chip detection efficiencies of 98.9% (median) and $96.2\% \pm 4.3\%$ (mean) (Supplementary Information) (Fig. 4e). These detectors have the ability to resolve 0, 1, 2, 3 and 4+ photons, as shown in the histogram of Fig. 4f (Supplementary Information).

Low-loss SiN waveguides, directional couplers and crossings

Silicon-on-insulator waveguides are limited in waveguide propagation loss owing to their large refractive index contrast³⁴. SiN waveguides, on the other hand, have lower refractive index contrast, offering a good compromise between confinement and sensitivity to manufacturing variations³⁴. We have demonstrated single-mode SiN waveguide loss of 1.8 ± 0.2 dB m⁻¹ and multimode waveguide loss of 0.5 ± 0.3 dB m⁻¹ (Fig. 5a), measured using a cutback technique (Supplementary Information). In this same platform, we have implemented waveguide crossings

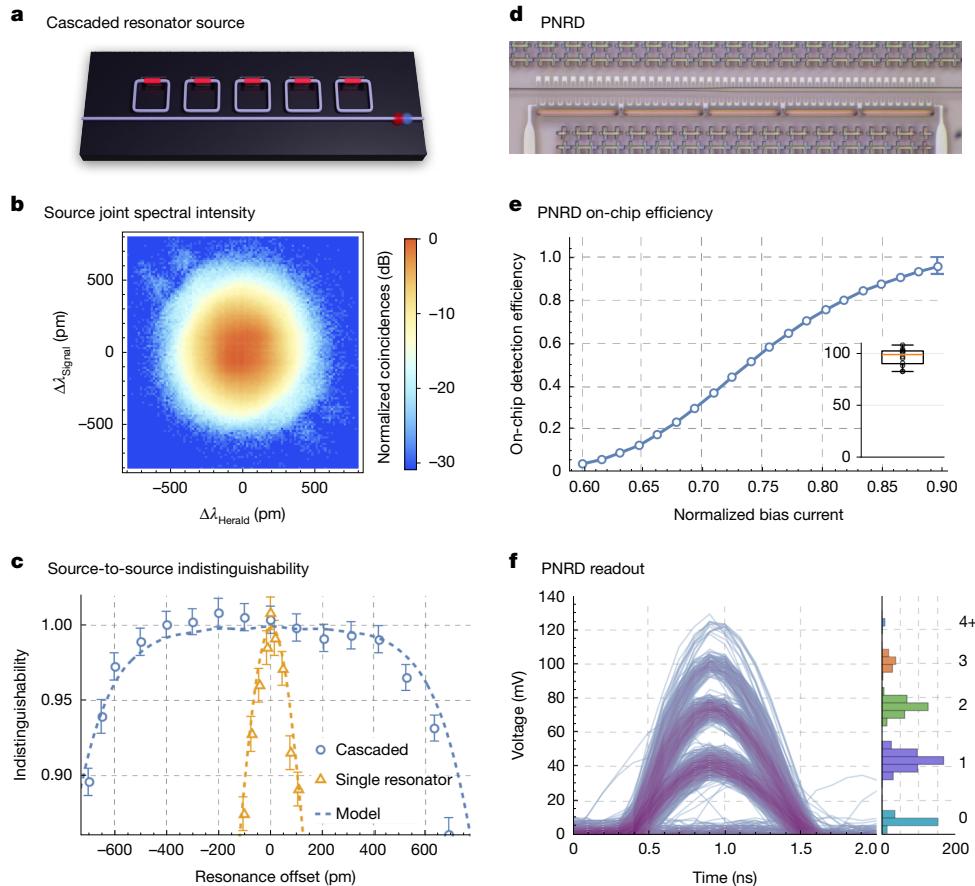


Fig. 4 | Cascaded resonator source and PNRD. **a**, Schematic of the source. **b**, Measured joint spectral intensity of a cascaded resonator source showing up to 99.35% purity, assuming flat spectral phase (Supplementary Information). **c**, Measured indistinguishability of two source copies as a function of the resonance wavelength offset (Supplementary Information). **d**, Top-down optical micrograph of a SiN-waveguide-coupled PNRD, in which single-photon detectors (SNSPDs) are crossing a waveguide and absorb light from the waveguide through evanescent coupling. Sets of SNSPDs are connected through on-chip resistors to comprise a unit cell. Identical unit cells are

connected in series. **e**, On-chip detection efficiency for the PNRD shown in **d** as a function of normalized bias current, showing the average across six unique devices (Supplementary Information). **e**, Inset, distribution of single-shot detection efficiency for each of the unique devices biased at roughly $0.9 I_{\text{SW}}$ at two input power levels. **f**, Left, persistent plot of the electrical photodetection signal (voltage traces) of a four-unit-cell PNRD. The traces were recorded using a cryogenic amplifier. The voltage traces show five distinct levels, corresponding to 0, 1, 2, 3 and 4 unit cells detecting photons simultaneously. Right, voltage trace histogram.

with 1.2 ± 0.4 m dB loss and waveguide splitters with 0.5 ± 0.2 m dB loss (Supplementary Information) (Fig. 5b). These component losses are about two times away from our target value, whereas the waveguide losses are on target.

SiN also provides advantages for photon generation. The ultralow loss combined with its Kerr nonlinearity supports SFWM with high signal-to-noise ratio. Further, there is an absence of nonlinear loss, allowing sources to operate with low loss at high pair rates, unlike silicon, for which two-photon absorption degrades performance³⁵.

Fibre-to-chip coupling

Low-loss coupling of light from optical fibres to our quantum photonic chips is required to make fibre networking practical. We implement new edge coupler designs that minimize mode overlap and mode conversion loss, enabling high-performance fibre-to-chip coupling. A key challenge is to convert the highly confined on-chip waveguide mode to match the much larger mode of optical fibre. To measure the insertion loss of the edge coupler, a chip is positioned between input and output optical fibres using high-precision optical alignment stages. Figure 5c shows repeated measurements from two of our best chip-to-fibre coupler designs, with coupling loss to standard telecommunications-grade optical fibre (SMF-28) of 127 ± 18 m dB and coupling loss to high-numerical aperture fibre (UHNA4) of 52 ± 12 m dB.

Electro-optic switching

To overcome the intrinsic non-determinism of both spontaneous sources and fusion gates, photonic quantum computing will require beyond-state-of-the-art high-speed optical switches, to enable large optical networks that can be rapidly reconfigured on the basis of the results of previous heralded photon generation, entangling gates and fusion outcomes¹⁸. The key component required for such switching networks is a high-speed, low-loss electro-optic phase shifter. Complex $N \times M$ networks may be constructed by embedding this phase shifter into passive interferometers constructed from the beam splitter and crossing devices previously described¹⁸.

The performance of the phase shifter is fundamentally constrained by the choice of electro-optic material. We incorporate BTO³⁶ into our photonic stack as the electro-optic phase shifter. We have developed a proprietary process for the growth of high-quality BTO films using molecular-beam epitaxy, compatible with foundry processes, on full 300-mm silicon wafers. We achieved a 3σ thickness uniformity of $<3\%$ across the entire 300-mm wafer, with electro-optic Pockels values of $>1,000$ pm V⁻¹ (compared with about 30 pm V⁻¹ for lithium niobate³⁷), measured through free-space Pockels measurements (Fig. 5d).

The fabricated 2×2 BTO Mach-Zehnder switches include a 2-mm-long phase shifter section, with a propagation loss of 53 ± 3 dB m⁻¹

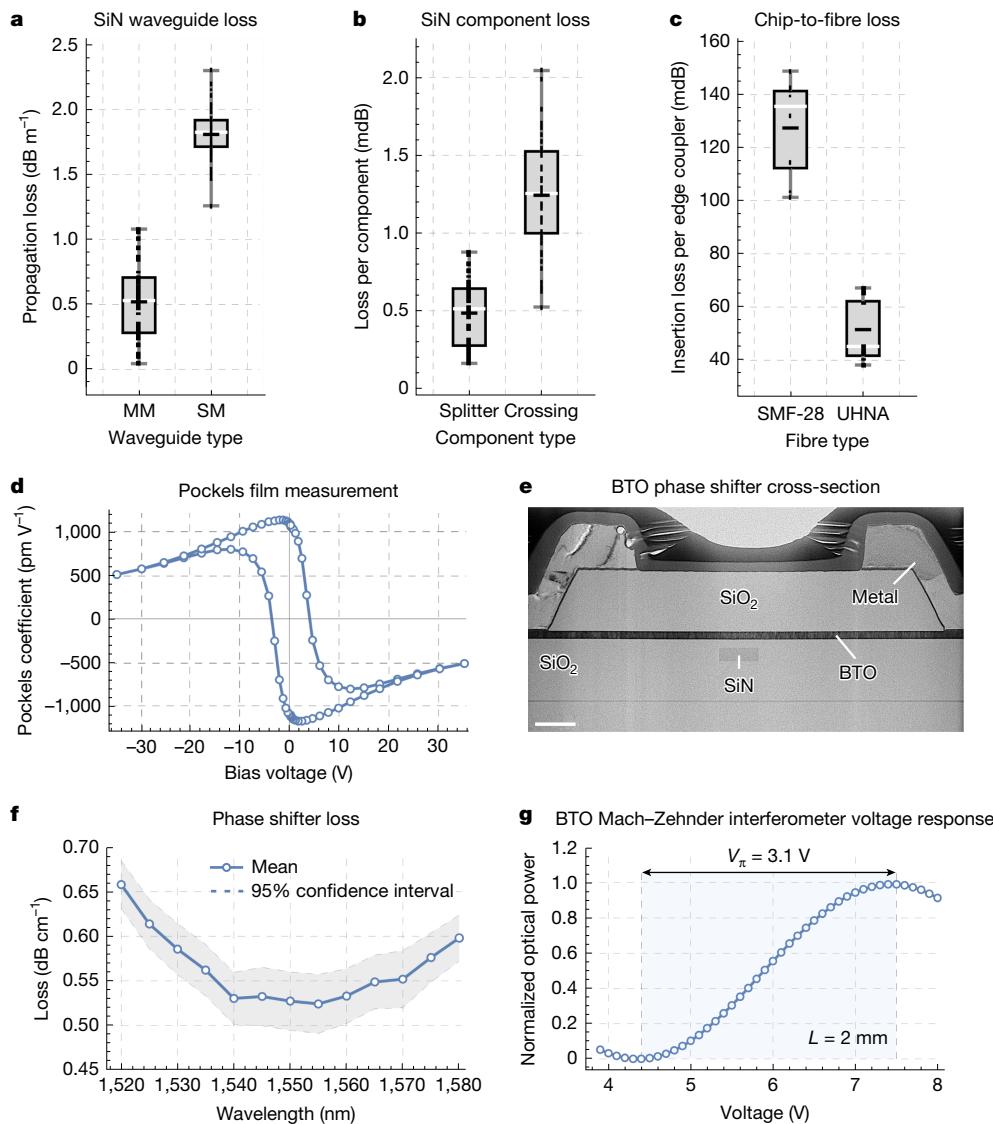


Fig. 5 | Waveguide and component loss and BTO optical switch. **a–c**, Loss of SiN-based components with mean (black line) and median (white line). **a**, SiN waveguide loss measurement, showing results across example wafers for both multimode (MM) and single-mode (SM) waveguides (Supplementary Information). **b**, SiN component loss for waveguide splitters and crossings (Supplementary Information). **c**, Chip-to-fibre loss. The fibre-to-chip coupling is measured in the low-loss regime using repeated transmission measurements on two exemplary devices designed for SMF-28 fibre and an exemplary device designed for UHNA fibre (Supplementary Information). **d**, Free-space electro-optic measurement of the effective Pockels coefficient of a BTO film grown

(Fig. 5f) and a d.c. $V_{\pi}L$ of 0.62 V.cm (Fig. 5g and Supplementary Information). This gives a phase shifter insertion loss of about 100 dB and a phase shifter half-wave loss-voltage product ($\alpha V_{\pi}L$) of 0.33 ± 0.02 dB.V, enabling a path to the construction of larger $N \times M$ low-loss switching networks required for photonic quantum computing. The insertion loss of this device is about two times away from our target value.

Conclusion

We have described modifications made to an industrial semiconductor manufacturing process for integrated quantum photonics, demonstrating record performance. Through the addition of new materials, designs and process steps, we have enabled volume manufacturing of heralded photon sources and superconducting single-photon

by molecular-beam epitaxy, with hysteresis. **e**, Scanning electron microscope cross-section of a fully fabricated BTO-on-SiN phase shifter. Scale bar, 1,000 nm. **f**, Cutback-based propagation loss measurement of a BTO-on-SiN phase shifter (data points and guide line), with 95% confidence intervals provided (dashed lines). **g**, Measured optical transmission of a Mach-Zehnder interferometer with a $L = 2$ -mm-long BTO phase shifter. A voltage was applied to one arm of the Mach-Zehnder interferometer, resulting in a $V_{\pi}L = 0.62$ V.cm in a non-push-pull configuration (Supplementary Information), in which V_{π} is the voltage required to change the phase by π radians. Wafer maps of these results can be found in the Supplementary Information.

detectors, together with photon manipulation by means of interferometry, tunability and control of unwanted light. We have also described higher-performing devices, towards a resolution of the outstanding limitations of this baseline platform.

FBQC supports fault-tolerant protocols that can tolerate on the order of 10% total accumulated optical loss between photon emission and detection, with per-qubit errors in the fusion network on the order of 1% (refs. 13,38,39). Here we have demonstrated a feature-complete set of optical components for FBQC, each with optical losses at the several-percent or below level, as well as fully integrated circuits demonstrating high-visibility interference, distribution and measurement functionalities of photonic qubits, all with sub-percent error levels.

Improvements to the platform and processes are still required. It will be necessary to further reduce SiN materials and component losses,

improve filter performance and increase detector efficiency to push overall photon loss and fidelity. Some specific examples of the remaining challenges are: implementation of low-loss $N \times M$ fast switches towards a multiplexed photon source; repeatable alignment and packaging of ultralow-loss chip-to-fibre edge connects; and improved targeting and robustness of photonic designs to minimize the need for tuning and trimming with heaters, thus further reducing the heat load at cryogenic temperatures.

Finally, we note that the platforms we have developed, and their future improvements, are highly flexible. Component arrangements are highly configurable, making the system suitable for different variations of quantum computer architectures, different quantum technology applications and, indeed, other photonic technologies. The ability to connect chips by fibre with very low loss makes the system technologically scalable across large numbers of photonic dies and allows for future networking or connections between different systems in a range of application spaces. Although the singular intent of our development is a useful fault-tolerant quantum computer, we hope that the influence of our industrially manufacturable quantum photonic platform will be broad and substantial.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-025-08820-7>.

1. O'Brien, J. L. Optical quantum computing. *Science* **318**, 1567–1570 (2007).
2. Knill, E., Laflamme, R. & Milburn, G. J. A scheme for efficient quantum computation with linear optics. *Nature* **409**, 46–52 (2001).
3. Ralph, T. C., White, A. G., Munro, W. J. & Milburn, G. J. Simple scheme for efficient linear optics quantum gates. *Phys. Rev. A* **65**, 012314 (2001).
4. Yoran, N. & Reznik, B. Deterministic linear optics quantum computation with single photon qubits. *Phys. Rev. Lett.* **91**, 037903 (2003).
5. Nielsen, M. A. Optical quantum computation using cluster states. *Phys. Rev. Lett.* **93**, 040503 (2004).
6. Browne, D. E. & Rudolph, T. Resource-efficient linear optical quantum computation. *Phys. Rev. Lett.* **95**, 010501 (2005).
7. Geng, H. *Semiconductor Manufacturing Handbook* 2nd edn (McGraw-Hill, 2018).
8. Fowler, A. G., Devitt, S. J. & Hollenberg, L. C. L. Implementation of Shor's algorithm on a linear nearest neighbour qubit array. *Quantum Inf. Comput.* **4**, 237–251 (2004).
9. Li, Y., Humphreys, P. C., Mendoza, G. J. & Benjamin, S. C. Resource costs for fault-tolerant linear optical quantum computing. *Phys. Rev. X* **5**, 041007 (2015).
10. Rudolph, T. Why I am optimistic about the silicon-photonic route to quantum computing. *APL Photonics* **2**, 030901 (2017).
11. Gol'tsman, G. N. et al. Picosecond superconducting single-photon optical detector. *Appl. Phys. Lett.* **79**, 705–707 (2001).
12. Marsili, F. et al. Detecting single infrared photons with 93% system efficiency. *Nat. Photon.* **7**, 210–214 (2013).
13. Bartolucci, S. et al. Fusion-based quantum computation. *Nat. Commun.* **14**, 912 (2023).
14. Giewont, K. et al. 300-mm monolithic silicon photonics foundry technology. *IEEE J. Sel. Top. Quantum Electron.* **25**, 1–11 (2019).
15. Migdall, A. L., Branning, D. & Castelletto, S. Tailoring single-photon and multiphoton probabilities of a single-photon on-demand source. *Phys. Rev. A* **66**, 053805 (2002).
16. Pittman, T. B., Jacobs, B. C. & Franson, J. D. Single photons on pseudodemand from stored parametric down-conversion. *Phys. Rev. A* **66**, 042303 (2002).
17. Meyer-Scott, E., Silberhorn, C. & Migdall, A. Single-photon sources: approaching the ideal through multiplexing. *Rev. Sci. Instrum.* **91**, 041101 (2020).
18. Bartolucci, S. et al. Switch networks for photonic fusion-based quantum computing. Preprint at <https://arxiv.org/abs/2109.13760> (2021).
19. Siew, S. Y. et al. Review of silicon photonics technology and platform development. *J. Light. Technol.* **39**, 4374–4389 (2021).
20. Shekhar, S. et al. Roadmapping the next generation of silicon photonics. *Nat. Commun.* **15**, 751 (2024).
21. Bartolucci, S. et al. Creation of entangled photonic states using linear optics. Preprint at <https://arxiv.org/abs/2106.13825> (2021).
22. Silverstone, J. W., Bonneau, D., O'Brien, J. L. & Thompson, M. G. Silicon quantum photonics. *IEEE J. Sel. Top. Quantum Electron.* **22**, 390–402 (2016).
23. Vernon, Z. et al. Truly unentangled photon pairs without spectral filtering. *Opt. Lett.* **42**, 3638–3641 (2017).
24. Hu, X., Holzwarth, C. W., Masciarelli, D., Dauler, E. A. & Berggren, K. K. Efficiently coupling light to superconducting nanowire single-photon detectors. *IEEE Trans. Appl. Supercond.* **19**, 336–340 (2009).
25. Biennfang, J. C. et al. Single-photon sources and detectors dictionary. NIST Internal Report NIST IR 8486. <https://doi.org/10.6028/NIST.IR.8486> (2023).
26. Politi, A., Cryan, M. J., Rarity, J. G., Yu, S. & O'Brien, J. L. Silica-on-silicon waveguide quantum circuits. *Science* **320**, 646–649 (2008).
27. Wang, J., Sciarri, F., Laing, A. & Thompson, M. G. Integrated photonic quantum technologies. *Nat. Photon.* **14**, 273–284 (2020).
28. Lauk, N. et al. Perspectives on quantum transduction. *Quantum Sci. Technol.* **5**, 020501 (2020).
29. Bombin, H. et al. Interleaving: modular architectures for fault-tolerant photonic quantum computing. Preprint at <https://arxiv.org/abs/2103.08612> (2021).
30. Litinski, D. & Nickerson, N. Active volume: an architecture for efficient fault-tolerant quantum computers with limited non-local connections. Preprint at <https://arxiv.org/abs/2211.15465> (2022).
31. Llewellyn, D. et al. Chip-to-chip quantum teleportation and multi-photon entanglement in silicon. *Nat. Phys.* **16**, 148–153 (2020).
32. Chow, J. M. et al. Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits. *Phys. Rev. Lett.* **109**, 060501 (2012).
33. Jahamirnejad, S. & Fiore, A. Proposal for a superconducting photon number resolving detector with large dynamic range. *Opt. Express* **20**, 5017–5028 (2012).
34. Sacher, W. D. et al. Monolithically integrated multilayer silicon nitride-on-silicon waveguide platforms for 3-D photonic circuits and devices. *Proc. IEEE* **106**, 2232–2245 (2018).
35. Rosenfeld, L. M. et al. Mid-infrared quantum optics in silicon. *Opt. Express* **28**, 37092–37102 (2020).
36. Eltes, F. et al. An integrated optical modulator operating at cryogenic temperatures. *Nat. Mater.* **19**, 1164–1168 (2020).
37. Zhu, D. et al. Integrated photonics on thin-film lithium niobate. *Adv. Opt. Photonics.* **13**, 242–352 (2021).
38. Bombin, H., Dawson, C., Nickerson, N., Pant, M. & Sullivan, J. Increasing error tolerance in quantum computers with dynamic bias arrangement. Preprint at <https://arxiv.org/abs/2303.16122> (2023).
39. Bombin, H. et al. Fault-tolerant complexes. Preprint at <https://arxiv.org/abs/2308.07844> (2023).

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



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025

PsiQuantum team

Koen Alexander¹, Avishai Benyamin¹, Dylan Black¹, Damien Bonneau¹, Stanley Burgos¹, Ben Burridge^{1,2}, Hugo Cable¹, Geoff Campbell¹, Gabriel Catalano¹, Alejandro Ceballos¹, Chia-Ming Chang¹, Sourav Sen Choudhury¹, C. J. Chung¹, Fariba Danesh¹, Tom Dauer¹, Michael Davis¹, Eric Dudley¹, Ping Er-Xuan¹, Josen Fargas¹, Alessandro Farsi¹, Colleen Fenrich¹, Jonathan Frazer^{1,2}, Masaya Fukami¹, Yogeeswaran Ganesan¹, Gary Gibson¹, Mercedes Gimeno-Segovia¹, Sebastian Goeldi¹, Patrick Goley¹, Ryan Haislmaier¹, Sami Halimi¹, Paul Hansen¹, Sam Hardy¹, Jason Horng¹, Matthew House¹, Hong Hu¹, Mehdi Jadi¹, Vijay Jain¹, Henrik Johansson¹, Thomas Jones^{1,2}, Vimal Kamineni¹, Nicholas Kelez¹, Ravi Koustuban¹, George Kovall¹, Peter Krogen¹, Nikhil Kumar¹, Yang Liang¹, Nicholas LiCausi¹, Dan Llewellyn^{1,2}, Kimberly Lokovic¹, Michael Lovelady¹, Vitor Riseti Manfrinato¹, Ann Melnichuk¹, Gabriel Mendoza¹, Brad Moores¹, Shaunak Mukherjee¹, Joseph Munns¹, Francois-Xavier Musalem¹, Faraz Najafi¹, Jeremy L. O'Brien¹, J. Elliott Ortmann¹, Sunil Pai¹, Bryan Park¹, Hsuan-Tung Peng¹, Nicholas Penthorn¹, Brennan Peterson¹, Gabriel Peterson¹, Matt Poush¹, Geoff J. Pryde¹, Tarun Rampasad^{1,2}, Gareth Ray^{1,2}, Angelita Viejo Rodriguez¹, Brian Roxworthy¹, Terry Rudolph¹, Dylan J. Saunders¹, Pete Shadbolt¹, Deesha Shah¹, Andrea Bahgat Shehata¹, Hyungki Shin¹, Jeffrey Sinsky¹, Jake Smith¹, Ben Sohn¹, Young-Ik Sohn¹, Gyeongho Son¹, Mario C. M. Souza¹, Chris Sparrow¹, Matteo Staffaroni¹, Camille Stavrakas¹, Vijay Sukumaran¹, Davide Tamborini¹, Mark G. Thompson^{1,2}, Khanh Tran¹, Mark Triplett¹, Maryann Tung¹, Andrzej Veitia¹, Alexey Vert¹, Mihai D. Vidrighin^{1,2}, Ilya Vorobeichik¹, Peter Weigel¹, Matthew Wingert¹, Jamie Wooding¹ & Xinran Zhou¹

¹PsiQuantum Corp., Palo Alto, CA, USA. ²PsiQuantum Ltd., Daresbury, UK. [✉]e-mail: mark@psiquantum.com

Data availability

The data that support the findings in this study are available at http://github.com/PsiQ/2404_17570.

Acknowledgements We thank the continual support of our investors, partners and the government agencies.

Author contributions The platform and benchmarking experiments were conceived and planned by the PsiQuantum team. The system architecture and engineering divisions at PsiQuantum created the data-analysis methodology, modelling and metrological tools that powered the experiments, assembled the system, handled calibrations and collected data.

The component teams were responsible for developing photonic components and the process teams established the manufacturing platform and BTO devices. The manuscript and Supplementary Information were composed and edited by all authors.

Competing interests The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-025-08820-7>.

Correspondence and requests for materials should be addressed to Mark G. Thompson.

Peer review information *Nature* thanks Stephane Clemmen and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>.