

Real-time optical signal equalization with a silicon photonic spatially distributed reservoir computer

Received: 25 March 2025

Accepted: 23 June 2026

Published online: 21 July 2026

 Check for updatesRuben Van Assche^{1,2}✉, Sarah Masaad^{1,2}, Emmanuel Gooskens^{1,2},
Stijn Sackesyn^{1,2}, Joris Van Kerrebrouck^{2,3}, Xin Yin^{2,3} & Peter Bienstman^{1,2}

Communication through optical fibres experiences limitations due to chromatic dispersion and nonlinear Kerr effects that degrade the signal. Mitigating these impairments is typically done using complex digital signal processing algorithms. However, these equalization methods require substantial power consumption and introduce high latencies. Photonic reservoir computing (a subfield of neural networks) offers an alternative solution—processing signals in the analogue optical domain. Here we present, to our knowledge, the first experimental demonstration of the real-time equalization of fibre distortions using a silicon photonics chip that combines the recurrent reservoir and the programmable read-out layer. We successfully equalize a 28-Gbps on–off keying signal across varying power levels and fibre lengths, even in the highly nonlinear regime. We obtain bit error rates that are orders of magnitude below previously reported optical equalization methods, reaching as low as 4×10^{-7} , far below the generic forward error correction limit of 5.8×10^{-5} used in commercial Ethernet interfaces.

Global data traffic has increased exponentially over the past 20 years¹. Although optical networks were initially primarily designed for long-distance, high-speed transmission, emergent novel applications such as the Internet of Things, cloud services, artificial intelligence and edge computing have drastically increased the amount of short-length inter- and intra-data-centre traffic. For long-reach optical networks, wavelength-division multiplexing and coherent modulation formats are frequently used and can achieve high data rates and excellent spectral efficiency. Examples of these formats are polarization-division multiplexing quadrature phase shift keying and polarization-division multiplexing quadrature amplitude modulation. However, the transceivers for these formats are costly, which makes them unsuitable for short-haul networks in which the rapid increase in demand for network capacity has made deployment and operational costs a critical factor². Currently, intensity modulation with direct detection (IM/DD) is the transceiver technology used for short-reach networks. However, this

limits the achievable bit rate, since IM/DD only encodes data on the intensity of the signal, compared with coherent communication in which information is encoded on the amplitude, phase and/or polarization of the signal. To meet the increasing demands of the industry, it is essential to develop solutions that enhance data rates and power efficiency of IM/DD systems. Coherent and IM/DD transceivers occupy different points in the link design space. Coherent detection provides high spectral efficiency and strong robustness to dispersion and polarization effects through powerful digital compensation, but generally relies on substantial mixed-signal and digital signal processing (DSP) complexity as well as more complex receiver hardware. In many short-reach regimes, IM/DD remains attractive due to simpler optics and packaging, yet its performance can be limited by fibre-induced memory and nonlinear distortions that demand increasingly complex equalization. The present work targets this bottleneck by demonstrating a reconfigurable, integrated optical-domain equalizer for

¹Photonics Research Group, INTEC, Ghent University–imec, Ghent, Belgium. ²imec, Leuven, Belgium. ³IDLab, INTEC, Ghent University–imec, Ghent, Belgium. ✉e-mail: ruben.vanassche@ugent.be

IM/DD signals that can provide strong impairment mitigation without requiring coherent detection. Although coherent architectures will remain essential in regimes that demand higher spectral efficiency and full field recovery, our contribution is positioned as a practical route to extend IM/DD capability using integrated optical processing. System-level comparisons of IM/DD and coherent transceiver architectures in short-reach and data-centre regimes—including the underlying cost, power and complexity trade-offs—are discussed elsewhere^{2–6}.

When considering fibre-optic communication systems, a large constraint for the throughput of IM/DD systems stems from the distortions arising from propagation through optical fibres. These distortions can be subdivided into linear distortions, due to chromatic dispersion and polarization-mode dispersion, as well as nonlinear distortions due to the Kerr effect. These Kerr nonlinearities, in particular, pose challenges because of the self-phase modulation they induce. In a single channel, it is the interplay between chromatic dispersion and self-phase modulation that severely degrades the quality of the signal coming out of the fibre. State-of-the-art approaches utilize DSP techniques to mitigate these issues. Although effective, these algorithms introduce additional communication latency, as they require processing in the electrical domain. Moreover, their computational complexity leads to substantial power consumption, presenting challenges for energy efficiency and operational costs in short-haul optical networks. Their performance for signals with high nonlinear distortions is limited as well^{6,7}.

Extensive research has been conducted to explore alternative methods for addressing these challenges. For compensating linear dispersion, solutions have been developed in both electric and optical domains, many of which are now widely adopted. Examples include dispersion-compensated/shifted fibres⁸ and fibre Bragg gratings⁹ in the optical domain, and fractionally spaced equalizers¹⁰ and feed-forward equalizers¹¹ in the digital electronic domain.

Although the optical solutions for linear dispersion perform extremely well and achieve bit error rates (BERs) well under any practical threshold needed for post-processing, they do not mitigate Kerr nonlinearity and cannot adapt to different operating conditions without having to perform major interventions in the link. These passive solutions, moreover, suffer from bandwidth limitations in the modulator/transceiver, a limitation for which the reservoir chips in part compensate through training.

For the nonlinear part of the problem, electrical and optical solutions have been found as well. In the electric domain, digital back-propagation¹² or nonlinear Volterra series¹³ are notable solutions. These implementations are power-hungry, however, due to the complexity of their algorithms, and introduce a latency due to their digital implementation. Optically, the most important examples are optical phase conjugation¹⁴ and phase-conjugated twin waves¹⁵. However, optical phase conjugation relies on an a priori known fibre length, which means it cannot be used for an all-purpose solution, and phase-conjugated twin waves sacrifice half the transmission capacity.

In addition to these implementations, research has been conducted to try and solve these problems using machine learning and neural networks, both in the digital electronic domain and in the optical domain^{16–23}. An intuitive approach for machine learning in this context is to implement a neural network after the optical fibre or after the receiver and to consider this neural network as a post-processor that retrieves the original undistorted signal. A possible approach towards equalizing IM/DD signals is to use a complex perceptron in which the signal is split into multiple delayed copies before being recombined. This method has recently been demonstrated for dispersion compensation in a 20-Gbps PAM-4 signal over 125 km (ref. 24), and for orthogonal frequency-division multiplexing pre-compensation²⁵, showing strong performance in the linear regime. Another promising candidate for an optical implementation of this concept is reservoir computing (RC). This is a machine learning paradigm that uses a fixed, randomly




Fig. 1 | Silicon photonics integrated reservoir. **a**, Photograph of the photonic chip on a printed circuit board with a 1 € coin next to it for scale. **b**, Diagram of signal equalization using RC. PIC, photonic integrated circuit.

initialized recurrent neural network (RNN), followed by a simple trainable linear read-out. In practice, the random initialization typically stems from fabrication imperfections. Moreover, the simple trainable linear read-out has the advantage of being easy to train and to help keep the power consumption limited.

Within the reservoir paradigm, there are two possible implementations: time-delay-based reservoirs (TDRCs) and spatially distributed reservoirs²⁶. Experiments and simulations utilizing TDRCs have already shown the ability for optically equalizing dispersion^{27–31}. However, these systems require time multiplexing, which introduces challenges at high speeds due to the response time of the nonlinear neuron. This means that for TDRCs, a trade-off exists between the computational capacity and the processing rate^{32–34}.

This issue is not present for spatially distributed reservoirs in which we previously showed that a linear and passive integrated circuit as a reservoir, in combination with a separate digital read-out layer, forms a suitable solution for signal equalization³⁵. Additionally, in terms of bandwidth, a single chip with a single set of weights can simultaneously equalize signals at multiple wavelengths³⁶. However, both of these demonstrations did not have functional optical weights and all output channels needed to be measured sequentially and processed digitally offline. As a result, these systems did not operate in real time, despite the high-speed optical computation in the reservoir itself. In this Article, ‘real time’ refers to inference at the full line rate of the optical signal (no buffering or time stretching), with training performed separately via iterative hardware-in-the-loop optimization. There was an initial attempt to integrate a read-out layer, which, however, was limited in scale and only tested on simple academic benchmarks as opposed to on a complex, industrially relevant highly nonlinear task³⁷. Beyond integration, another scientific contribution is to demonstrate that a recurrent coherent photonic processor with a programmable on-chip read-out can be trained in hardware and used to mitigate fibre-induced memory and nonlinear waveform distortions at line rate, establishing a physical computing approach that outperforms a conventional linear feed-forward equalization (FFE) baseline under matched conditions. After publishing a preprint of this paper, related research was made public in ref. 38 in which a deep RC architecture was used as an optical signal processor with DSP-free operation at data rates as high as 100 Gbaud for a PAM-4 signal, but only for a 5-km link.

In this work, we present a novel silicon photonics chip (Fig. 1) that incorporates a fully integrated optical read-out layer, enabling real-time dispersion within a single-chip architecture with optical input and output. Our experiments demonstrate the chip’s capability to equalize a 28-Gbps on-off keying (OOK) signal at 10 dBm for fibre lengths up to 50 km and a 17-dBm signal over 25 km, achieving, in all

cases, a BER below the generic standard forward error correction (GFEC) threshold of 5.8×10^{-5} (ref. 39).

Although RC and photonic equalization have both been explored previously, this work goes beyond an incremental combination by (1) realizing a spatially distributed recurrent reservoir with an integrated, programmable optical read-out that performs phase-sensitive (complex) weighting on chip, (2) enabling single-shot line-rate inference without time multiplexing and without sequential multichannel state acquisition, and (3) experimentally validating the resulting reconfigurable processor on a challenging fibre equalization task with nonlinear memory, benchmarked against a carefully matched electrical FFE baseline. This closes the gap between high-speed optical reservoir transformations demonstrated with offline read-out and a practical integrated equalizer core with reconfigurable, on-chip read-out suitable for real-time operation.

This paper highlights several benefits of our integrated silicon photonics equalization core:

1. Real-time optical signal processing: the chip processes signals online (that is, in real time) entirely within the optical domain, eliminating the latencies associated with electronic computation or inter-domain connectivity as is the case for electro-optic solutions. The only latency is the intrinsic light propagation time within the chip, measured by reflectometry to be equal to 243.32 ps, which is negligible in comparison. Importantly, the absence of electrical overhead further reduces power loss compared with electro-optic alternatives.
2. High modulation speeds: the on-chip delay lines synchronize the reservoir's timescale with the bit period. In the current design, we targeted compatibility with 28-Gbaud signals. However, as we will show, removing these delays allows compatibility with single-channel modulation rates nearing 900 Gbaud. Even though this is far beyond the capabilities of current transceivers and detectors, this illustrates that our approach is highly future-proof.
3. Linear and passive device design: the reservoir comprises exclusively linear and passive components, such as multimode interferometers (MMIs) and waveguides. This reduces the fabrication and operational complexity, as integrated nonlinear devices are either difficult to integrate (for example, if they require novel materials) or difficult to control (for example, if they rely on resonances). The only required nonlinearity in our implementation is that of the final photodiode, which converts complex-valued fields to real-valued intensities; this is an intrinsically nonlinear operation⁴⁰. Moreover, this nonlinearity does not require a minimum power level, paving the way to low power consumption (our current implementation uses integrated tungsten heaters, which act as complex weights; however, replacing these with non-volatile weights, for example, based on barium titanate⁴¹, opens pathways to ultralow power consumption).
4. Agnostic to fibre length and input power: the chip demonstrates robust performance across diverse test cases involving both linear and nonlinear distortions. This was tested by experiments with fibre lengths up to 50 km for a 10-dBm input power into the fibre, and input powers up to 17 dBm for a 25-km fibre. This adaptability underscores its potential as a versatile solution for modern high-speed optical communication systems.
5. Complementary metal–oxide–semiconductor compatibility and scalability: the chip is based on silicon photonics, ensuring compatibility with complementary metal–oxide–semiconductor fabrication processes. This feature also enables seamless integration with silicon electronics and supports scalable production for widespread deployment and low cost.

The rest of this paper is structured as follows. In 'Integrated photonic RC', we will give a brief introduction to the RC formalism,

followed by the design of the chip. 'Experimental results' explains the experimental results. Subsequently, 'Future outlook' discusses these results and provides a future outlook. Finally, 'Conclusion' concludes the paper.

Integrated photonic RC

RC is an RNN approach introduced independently by ref. 42 and ref. 43. A reservoir computer is an RNN made up of three layers:

1. An input layer in which (analogue) signals are introduced to the reservoir.
2. The reservoir itself, typically a nonlinear dynamical system.
3. An output layer in which signals processed by the reservoir undergo linear combination.

Unlike traditional RNNs in which all interconnections are trained, RC retains fixed internal connections inside the reservoir. The random and static nature of the reservoir network in RC makes it well-suited for analogue computing, as it relaxes fabrication tolerances and complexity compared with physical RNNs.

Our photonic RC architecture is implemented in a silicon nitride platform and follows a four-port design⁴⁴. At each output of the reservoir, the signal will be a superposition of many different signals with different delays and different phases. In practice, the randomness of the reservoir is caused by fabrication-induced variations in the surface roughness and in the length of the waveguides. This results in a random (but fixed for any given chip) set of phase changes for the propagating signals. The nonlinearity of the photodiode, which acts as the receiver, will transform the relative phases of the different interfering signals into power variations. It turns out that this nonlinearity by itself gives us enough processing power to perform signal equalization, as shown in previous experiments^{35,36}. A reservoir consisting of eight nodes was chosen. This is the smallest number of nodes following a four-port architecture showing non-trivial recurrent behaviour. When a reservoir of a higher number of nodes would be fabricated, the memory capabilities of the reservoir would increase, leading to better performance for longer links at the cost of a larger footprint. The eight nodes of the reservoir are implemented using MMIs, with passive waveguides forming the interconnections. A microscope image of the fabricated chip is shown in Fig. 2a, and a schematic is provided in Fig. 2b. The four-port architecture usually uses 3×3 MMIs as nodes. Indeed, each node features three inputs and three outputs: one input connects to the input layer, one output links to the read-out layer and the remaining ports facilitate interconnections between nodes. However, to be able to use more standard building blocks, in our current design, each node is constructed by cascading two 2×2 MMIs (Fig. 2c), rather than using a single 3×3 MMI. This approach introduces asymmetry in the node transfer functions, but this can enhance the RNN performance. The input layer distributes the chip's single input to two reservoir nodes. Two inputs for the reservoir are used instead of one, to have a more even signal distribution across the network.

A key design parameter is the length of the interconnecting waveguides, or delay lines, which must be tuned to the data rate of the signal. Simulations indicate that an interconnection delay of half a bit period optimally balances signal propagation.

The read-out layer consists of eight Mach–Zehnder interferometers (MZIs), each corresponding to a reservoir node and tuned via two thermoelectric heaters (Fig. 2d). Adjusting the heater currents modifies the local temperature, thereby controlling both amplitude and phase of each MZI's output. This enables weighting of the reservoir signals before they are recombined in a summation tree. Functionally, the reservoir splits the single optical input into multiple delayed and phase-shifted replicas that are coherently recombined at successive MMI couplers. The combination of distinct delay lengths and relative phases produces a set of temporal features with memory extending over several bit periods, whereas recurrence causes these features to




Fig. 2 | Design of photonic integrated reservoir. **a**, Microscope image of the chip design, with the blue rectangle showing a node and the red rectangle showing a weight implemented as an MZI with heaters. **b**, Schematic of the chip. The blue circles indicate nodes, and arrows indicate the propagation direction

of light. Note that the detector is off-chip. w , amplitude of the complex readout weights; θ , phase of the complex readout weights. **c**, Close-up image of a node composed of two 2 x 2 MMIs. **d**, Close-up image of an MZI with heaters.

persist and gradually decay due to insertion loss. A fraction of the optical field at each node is routed to the integrated read-out, whose trained weights combine these states to mitigate distortions. More information on the internal state dynamics and memory of this four-port architecture is available elsewhere^{35,36,40,45}. Due to design and fabrication imperfections, the unpackaged chip introduces substantial losses, measured as 12-dB pure insertion loss, and an added 6–13-dB loss when equalizing, caused by having weights with amplitudes smaller than 1. Integrating a semiconductor optical amplifier could mitigate this loss and introducing additional nonlinearities, which may enhance performance for complex tasks⁴⁶.

Experimental results

The experiments evaluated the system's performance and two parameters were varied: the input power to the fibre and the fibre length. The input power was swept from 10 to 17 dBm for a fixed fibre length of 25 km, whereas the fibre length was varied from 25 to 50 km at a constant input power of 10 dBm. An increase in fibre length extends the interaction distance for both linear and nonlinear dispersion effects, whereas increasing the input power primarily enhances the impact of Kerr nonlinearity, thereby increasing the nonlinear complexity of the machine learning task. For the complete setup and processing steps, we refer to the Methods, Supplementary Figs. 1 and 2 and Supplementary Table 1. We emphasize that the results below are intended to benchmark equalization capability under controlled fibre-induced distortions (relative to a matched FFE baseline), and are not meant as a full system-sensitivity study including received-power sweeps, back-to-back penalties or complete noise budgeting.

Figure 3 provides a visual comparison of the equalization effectiveness through eye diagrams. These diagrams illustrate the signal degradation and subsequent restoration for a transmission distance of 50 km at an input power of 10 dBm. Each eye diagram corresponds to 90,000 bits. The histograms of these eyes and eye diagram of the optical back-to-back signal are provided in Supplementary Figs. 3 and 4, respectively. The first eye diagram, corresponding to the raw distorted signal at the output of the fibre, exhibits substantial inter-symbol interference due to chromatic dispersion, leading to severe eye closure. The second diagram shows a partially restored bitstream after applying digital FFE. However, residual nonlinear distortion effects remain, and the overall eye opening is still limited. By contrast, the third diagram—after equalization by the photonic reservoir—demonstrates a substantially wider eye opening, because in contrast to the FFE, the reservoir can also perform nonlinear operations. Not only is the black region of the eye wider than for the FFE, signifying a well-defined transition region, the eye is higher as well, indicating that the overall distinction between logical ones and zeros is considerably improved.

The eye diagrams also show the statistical BER, as calculated based on the assumption of a Gaussian distribution of ones and zeros. This serves as an initial quality metric for the eye diagrams. However, for subsequent plots, we will show the more accurate BER obtained through direct error counting over a larger dataset.

These BER results for different experimental conditions are summarized in Fig. 4. Figure 4a presents BER trends as fibre length increases at a fixed input power of 10 dBm. In this regime, chromatic dispersion is the dominant performance-limiting factor. For longer fibre spans, performance degrades due to the limited memory of both equalization




Fig. 3 | Experimental eye diagrams. a–c, Eye diagrams and the corresponding statistical BERs of the distorted signal after propagating through 50 km of optical fibre with an input power of 10 dBm (a), after applying a digital FFE with 17 taps (b) and after applying the reservoir chip (c). UI, unit interval.




Fig. 4 | Experimental BERs. a,b, Measured BERs for different lengths of the optical fibre for an input power of 10 dBm (a) and different input powers for a fibre length of 25 km (b). The grey dashed line signifies the resolution limit up to which predictions can accurately be made for the reservoir, and the GFEC boundary. SNR, signal-to-noise ratio.

systems. In particular, despite the FFE possessing a larger memory capacity than the photonic reservoir, the reservoir outperforms the FFE by at least two orders of magnitude in BER at any fibre length. This is because the reservoir performs rich operations on the input, compared with the FFE, where only linear combinations of delayed versions of the same signal are used. Another advantage compared with the FFE is that the reservoir processes the signal before the photodetector, avoiding problems due to frequency fading that digital implementations suffer from. The equalized decision variable is then obtained from the photodetector output in the electrical domain. Figure 4b shows the BER evolution as the input power increases for a fixed fibre length of 25 km. In this case, nonlinear dispersion and its interaction with chromatic dispersion impose performance constraints. Even with the same number of degrees of freedom, the photonic reservoir again surpasses the FFE, achieving at least a two-order-of-magnitude improvement in BER.

It is interesting to note that the Kerr effect is distributed along the fibre, whereas the square-law nonlinearity of the photodiode is localized at the detector. Accordingly, the observed mitigation should not be interpreted as all-optical regeneration or full cancellation of distributed Kerr nonlinearity. Instead, the recurrent reservoir reshapes the distorted waveform such that the subsequent square-law nonlinearity becomes sufficient to separate symbol classes and to enable compensation. This practical mitigation is a direct consequence of the RC architecture, rather than of the photodiode alone. The demonstrated system is, therefore, best described as an optical-domain analogue processor whose nonlinear decision variable is produced on photodetection (+ processor with an optoelectronic read-out).

In both experimental conditions, the photonic reservoir achieves equalization well below the GFEC limit³⁹ across all tested configurations, whereas the FFE fails to do so. This implies that the reservoir-based system requires a notably less computationally intensive forward error correction stage compared with the FFE. Table 1 provides a more

general comparison with state-of-the-art optical implementations of equalization and DSP. Note that non-machine learning optical solutions such as dispersion-compensating fibre and fibre Bragg gratings are not included in the comparison. Although they work very well in the linear regime, they fail in the nonlinear regime and, more importantly, lack the reconfigurability that all solutions in this comparison have.

We emphasize that Table 1 aggregates results obtained under differing experimental conditions and reporting conventions (for example, test set size and whether statistical BER is evaluated). Therefore, we avoid claiming superiority. Nevertheless, the strong BER observed here can be interpreted as arising from a specific combination of architectural characteristics. First, the recurrent, spatially distributed reservoir produces a coherent superposition of multiple delayed replicas with distinct phases, providing an intrinsic temporal memory and a high-dimensional dynamical representation that is well matched to dispersive channels. Second, the optical read-out is phase sensitive: the integrated interferometric weighting realizes a complex-valued linear regression on the optical field before photodetection, enabling the system to exploit phase information and to shape the waveform into an advantageous form before square-law detection. Third, this coherent mixing followed by square-law photodetection yields structured quadratic cross-terms, such that the effective nonlinearity used for symbol separation arises from the reservoir-plus-detector composition rather than from the photodiode alone. Finally, because this phase-sensitive interference system leads to a strongly non-convex optimization landscape, robust hardware-in-the-loop optimization using the covariance matrix adaptation evolution strategy is a practical enabler for discovering high-performing complex weight configurations under experimental non-idealities, taking every distortion and thermal cross-talk into account. Taken together, these factors provide the performance demonstrated here.

Although all but one BER measurement of the reservoir falls below the resolution limit of the experimental setup, further increasing the

Table 1 | Comparison of different state-of-the-art experimental optical equalization solutions for intensity modulation

Reference	<i>B</i> (Gbaud)	Modulation format	<i>L</i> (km)	<i>P</i> (dBm)	Real time	BER
Optical solutions						
Argyris et al. ³²	25	PAM-2	45	10	No	1.8×10^{-4}
Photonic TDRC using response laser with feedback loop; digital read-out						
Argyris et al. ³³	28	PAM-4	27	14	No	2.0×10^{-3}
Photonic TDRC using response laser with feedback loop; digital read-out						
Sackesyn et al. ³⁵	32	PAM-2	25	18	No	1.0×10^{-3}
Photonic split digital reservoir computing using four-port architecture; digital read-out						
Estébanez et al. ³⁴	56	PAM-4	100	6	No	3.0×10^{-4}
Photonic TDRC using response laser with feedback loop; digital read-out						
Li et al. ⁵⁰	56	PAM-4	100	0	No	$<3.0 \times 10^{-4}$
Integrated silicon tapped-delay FIR using microring/MZI tap bank; digital read-out						
Sarantoglou et al. ⁵¹	28	PAM-4	40	–	No	4.2×10^{-2}
Photonic extreme learning machine using Fabry–Pérot laser; digital read-out						
Sozos et al. ⁵²	64	PAM-4	50	3	Yes	$<1.0 \times 10^{-2}$
Recurrent optical spectrum splicer using Mach–Zehnder delayed interferometers; integrated read-out						
Staffoli et al. ⁵³	10	PAM-2	125	0	Yes	$<1.0 \times 10^{-4}$
Integrated optical perceptron using optical tapped delay line; integrated read-out						
This work	28	PAM-2	25	17	Yes	5.7×10^{-6}
This work	28	PAM-2	50	10	Yes	$<5.7 \times 10^{-6}$
DSP baselines						
Tang et al. ⁵⁴	28	PAM-4	80	–	No	3.1×10^{-3}
Receiver-side Memory Polynomial Equalizer–Decision Feedback Equalizer						
Yin et al. ⁵⁵	28	PAM-4	80	6	No	$\leq 3.8 \times 10^{-3}$
Transmitter-only pre-equalization (no post-equalization)						
Wang et al. ⁵⁶	64	OOK	100	–	No	$\leq 3.8 \times 10^{-3}$
Receiver-side Adaptive Channel-Matched Detection (Polynomial Nonlinear Equalizer + Decision Feedback Equalizer + Maximum Likelihood Sequence Estimation) only						

Fibre length and launch power are included because higher values for these parameters result in a more nonlinear task that is drastically more difficult to solve^{48,49}. Most experiments are not in real time, in the sense that either the input signal needed to be buffered and slowed down, or that the output signals were captured and subsequently processed offline on a PC. Note that for ref. 50, a number was not provided in the paper; therefore, an estimation of the best result was made based on the provided graphs. For ref. 53, the provided value is adjusted for the size of the dataset according to the resolution limit of 100 bits. *P*, launch power.

test set size was not pursued due to hardware constraints. Indeed, the arbitrary waveform generator used in these experiments has a limited bitstream capacity per transmission cycle of approximately 10^5 bits. Achieving a resolution limit of 10^7 would require a test set of 10^9 bits, necessitating 10^4 arbitrary waveform generator transmission cycles, compared with the 200 cycles performed for each experiment in this study. However, the close agreement between the statistical BER estimation and the directly measured BER strongly suggests that additional measurements would yield minimal deviation from the reported results.

Future outlook

The presented platform demonstrates important potential for future advancements, particularly in terms of throughput. Simulations based on the current work (Supplementary Figs. 5 and 6) indicate that by making the waveguides as short as possible to still allow for a four-port architecture in which the length and width of the MMIs is the limiting factor, the reservoir becomes compatible with a data rate of 896 Gbps in the case of ideal signals with an ideal optical signal-to-noise ratio. Although no modulators or detectors support these data rates, simulations show that hypothetically, this platform could perform dispersion compensation over a distance of 24 m at a launch power of 10 dBm. This configuration has a $B^2 \times L$ product approximately equivalent to our experiments

of a 28-Gbps signal transmitted over 25 km, where *B* is the data rate and *L* is the length of the optical fibre. Note that the lengths of the delay lines in the reservoir change accordingly to keep the ratio of waveguide length and data rate equivalent, ensuring similar characteristics to the fabricated device. A simulation showing the relationship between performance and data rate is shown in Supplementary Fig. 7. These results underline the platform's capacity to support future innovations in single-channel transceiver technology, by trivial modifications of the architecture. To estimate the ultimate throughput of this design, we consider the combination of these simulations with prior experimental and simulation results related to an electric read-out implementation of the same architecture for wavelength-division multiplexing^{36,47}. These studies have demonstrated that the four-port architecture can use the same set of weights to separately equalize tens to hundreds of distinct wavelength channels. Assuming the utilization of 100 channels, the achievable aggregate throughput would amount to 89.6 Tbps. In terms of metrics like tera-operations per second, the number would be even higher, since inside the network, each bit undergoes many elementary operations like splitting, combining and weighting. However, it is neither straightforward nor advisable to use this metric in this case, given the analogue nature of the computations and the recurrent character of the architecture. It should be noted that this is based on idealized scenarios in terms of signal generation, propagation and detection,

and for an actual transmission setup at such high speeds, further considerations may be necessary, enabled by future research findings.

These findings highlight the potential of this architecture equalization at ultrahigh data rates. Furthermore, the architecture's inherent versatility supports various modulation formats, including those used in coherent and self-coherent communication systems, as we have previously shown in ref. 45. This demonstrates its adaptability across different data rates, wavelengths and modulation schemes.

Conclusion

The cost-efficient and accurate mitigation of linear and nonlinear dispersion in fibre-optic communication remains a long-standing challenge in both industry and research. In this work, we experimentally demonstrate a photonic RC chip capable of equalizing a 28-Gbps OOK signal well below the GFEC threshold of 5.8×10^{-5} for fibre lengths up to 50 km and input powers up to 17 dBm. Furthermore, at the considered data points at which launch powers of at least 10 dBm were utilized, it is seen that the reservoir chip outperforms a conventional digital FFE method, which performs at least two orders of magnitude worse, and it particularly fails in reaching the same GFEC limit for more than one experiment. Crucially, this architecture enables signal equalization across different modulation formats and wavelengths and preserves an optical output, distinguishing it from many electro-optic alternatives that require optical-to-electrical conversion. More broadly, the same design principle can be generalized by scaling the reservoir dimensionality and by adjusting its effective memory depth via the number of nodes and the length of delay paths, enabling adaptation across data rates, wavelengths and modulation formats—including coherent modulation regimes.

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/s41566-026-01968-2>.

References

- Winzer, P. J., Neilson, D. T. & Chraplyvy, A. R. Fiber-optic transmission and networking: the previous 20 and the next 20 years. *Opt. Express* **26**, 24190–24239 (2018).
- Zhong, K. et al. Digital signal processing for short-reach optical communications: a review of current technologies and future trends. *J. Lightwave Technol.* **36**, 377–400 (2018).
- Perin, J. K., Shastri, A. & Kahn, J. M. Coherent data center links. *J. Lightwave Technol.* **39**, 730–741 (2021).
- Kahn, J. M., Perin, J. K. & Shastri, A. Data center links beyond 100 Gb/s per wavelength. In *Proc. Optical Fiber Communications Conference and Exhibition Th3C.5* (Optica Publishing Group, 2020).
- Hirokawa, T. et al. Analog coherent detection for energy efficient intra-data center links at 200 Gbps per wavelength. *J. Lightwave Technol.* **39**, 520–531 (2021).
- Che, D. & Chen, X. Modulation format and digital signal processing for IM-DD optics at post-200G era. *J. Lightwave Technol.* **42**, 588–605 (2024).
- Cartledge, J. C., Guimomar, F. P., Kschischang, F. R., Liga, G. & Yankov, M. P. Digital signal processing for fiber nonlinearities (invited). *Opt. Express* **25**, 1916–1936 (2017).
- Gruner-Nielsen, L. et al. Dispersion-compensating fibers. *J. Lightwave Technol.* **23**, 3566–3579 (2005).
- Petruzzini, P., Lowry, C., Goldhar, J. & Sivanesan, P. Dispersion compensation using only fiber Bragg gratings. In *Proc. Optical Fiber Communication Conference and the International Conference on Integrated Optics and Optical Fiber Communication FA5* (Optica Publishing Group, 1999).
- Ip, E. & Kahn, J. M. Digital equalization of chromatic dispersion and polarization mode dispersion. *J. Lightwave Technol.* **25**, 2033–2043 (2007).
- Bohn, M. & Xia, C. Electrical and optical equalization strategies in direct detected high-speed transmission systems. *AEU Int. J. Electron. Commun.* **63**, 526–532 (2009).
- Rafique, D., Zhao, J. & Ellis, A. D. Digital back-propagation for spectrally efficient WDM 112 Gbit/s PM m -ary QAM transmission. *Opt. Express* **19**, 5219–5224 (2011).
- Liu, L. et al. Intrachannel nonlinearity compensation by inverse Volterra series transfer function. *J. Lightwave Technol.* **30**, 310–316 (2012).
- Yariv, A., Fekete, D. & Pepper, D. M. Compensation for channel dispersion by nonlinear optical phase conjugation. *Opt. Lett.* **4**, 52–54 (1979).
- Liu, X., Chraplyvy, A. R., Winzer, P. J., Tkach, R. W. & Chandrasekhar, S. Phase-conjugated twin waves for communication beyond the Kerr nonlinearity limit. *Nat. Photon.* **7**, 560–568 (2013).
- Zhang, S. et al. Field and lab experimental demonstration of nonlinear impairment compensation using neural networks. *Nat. Commun.* **10**, 3033 (2019).
- Fan, Q., Zhou, G., Gui, T., Lu, C. & Lau, A. P. T. Advancing theoretical understanding and practical performance of signal processing for nonlinear optical communications through machine learning. *Nat. Commun.* **11**, 3694 (2020).
- Amari, A., Lin, X., Dobre, O. A., Venkatesan, R. & Alvarado, A. A machine learning-based detection technique for optical fiber nonlinearity mitigation. *IEEE Photon. Technol. Lett.* **31**, 627–630 (2019).
- Pan, X., Wang, X., Tian, B., Wang, C., Zhang, H. & Guizani, M. Machine-learning-aided optical fiber communication system. *IEEE Netw.* **35**, 136–142 (2021).
- Häger, C. & Pfister, H. D. Deep learning of the nonlinear Schrödinger equation in fiber-optic communications. In *Proc. IEEE International Symposium on Information Theory* (eds Guruswami, V. et al.) 1590–1594 (IEEE, 2018).
- Zibar, D., Piels, M., Jones, R. & Schaeffer, C. G. Machine learning techniques in optical communication. *J. Lightwave Technol.* **34**, 1442–1452 (2016).
- Sozos, K. et al. Experimental investigation of a recurrent optical spectrum slicing receiver for intensity modulation/direct detection systems using programmable photonics. *J. Lightwave Technol.* **42**, 7807–7815 (2024).
- Huang, C. et al. A silicon photonic–electronic neural network for fibre nonlinearity compensation. *Nat. Electron.* **4**, 837–844 (2021).
- Staffoli, E., Maddinelli, G. & Pavesi, L. A silicon photonic neural network for chromatic dispersion compensation in 20 Gbps PAM4 signal at 125 km and its scalability up to 100 Gbps. *J. Lightwave Technol.* **43**, 557–571 (2025).
- Marciano, P. R. N. et al. Chromatic distortion precompensation in OFDM-based optical systems through an integrated silicon photonic neural network. *J. Lightwave Technol.* **43**, 3034–3040 (2025).
- Van der Sande, G., Brunner, D. & Soriano, M. C. Advances in photonic reservoir computing. *Nanophotonics* **6**, 561–576 (2017).
- Argyris, A. et al. Comparison of photonic reservoir computing systems for fiber transmission equalization. *IEEE J. Sel. Topics Quantum Electron.* **26**, 5100309 (2020).
- Ranzini, S. M., Dischler, R., Da Ros, F., Bülow, H. & Zibar, D. Experimental investigation of optoelectronic receiver with reservoir computing in short reach optical fiber communications. *J. Lightwave Technol.* **39**, 2460–2467 (2021).
- Da Ros, F., Ranzini, S. M., Bülow, H. & Zibar, D. Reservoir-computing based equalization with optical pre-processing for short-reach optical transmission. *IEEE J. Sel. Topics Quantum Electron.* **26**, 7701912 (2020).

30. Dejonckheere, A. et al. All-optical reservoir computer based on saturation of absorption. *Opt. Express* **22**, 10868–10881 (2014).
31. Shen, Y.-W., Li, R.-Q., Liu, G.-T., Yu, J., He, X., Yi, L. & Wang, C. Deep photonic reservoir computing recurrent network. *Optica* **10**, 1745–1751 (2023).
32. Argyris, A., Bueno, J. & Fischer, I. Photonic machine learning implementation for signal recovery in optical communications. *Sci. Rep.* **8**, 8487 (2018).
33. Argyris, A., Bueno, J. & Fischer, I. PAM-4 transmission at 1,550 nm using photonic reservoir computing post-processing. *IEEE Access* **7**, 37017–37025 (2019).
34. Estébanez, I. et al. 56 Gbaud PAM-4 100 km transmission system with photonic processing schemes. *J. Lightwave Technol.* **40**, 55–62 (2022).
35. Sackesyn, S., Ma, C., Dambre, J. & Bienstman, P. Experimental realization of integrated photonic reservoir computing for nonlinear fiber distortion compensation. *Opt. Express* **29**, 30991–30997 (2021).
36. Gooskens, E., Sackesyn, S., Dambre, J. & Bienstman, P. Experimental results on nonlinear distortion compensation using photonic reservoir computing with a single set of weights for different wavelengths. *Sci. Rep.* **13**, 21399 (2023).
37. Ma, C. et al. Integrated photonic reservoir computing with an all-optical readout. *Opt. Express* **31**, 34843–34854 (2023).
38. Wang, B. et al. An all-optical signal processor enabling terabit-per-second real-time equalization. *Science* **392**, eady5344 (2026).
39. Juniper Networks. *Understanding Pre-FEC BER Monitoring and BER Thresholds* (Juniper Networks, 2023); <https://www.juniper.net/documentation/us/en/software/connectivity-services-director5.3/csd-user-guide/connectivity-services-director-user-guide/topics/concept/monitor-understand.html>
40. Vandoorne, K. et al. Experimental demonstration of reservoir computing on a silicon photonics chip. *Nat. Commun.* **5**, 3541 (2014).
41. Stark, P. et al. Non-volatile photonic weights and their impact on photonic reservoir computing systems. In *Proc. Conference on Lasers and Electro-Optics Europe and European Quantum Electronics Conference* jsi_1_3 (Optica Publishing Group, 2019).
42. Jaeger, H. & Haas, H. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* **304**, 78–80 (2004).
43. Maass, W., Natschläger, T. & Markram, H. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* **14**, 2531–2560 (2002).
44. Sackesyn, S., Ma, C., Dambre, J. & Bienstman, P. An enhanced architecture for silicon photonic reservoir computing. In *Proc. Cognitive Computing 2018—Merging Concepts with Hardware* 1–2 (2018).
45. Masaad, S., Sackesyn, S., Sygletos, S. & Bienstman, P. Experimental demonstration of 4-port photonic reservoir computing for equalization of 4 and 16 QAM signals. *J. Lightwave Technol.* **42**, 8555–8563 (2024).
46. Van Gasse, K., Wang, R. & Roelkens, G. 27 dB gain III–V-on-silicon semiconductor optical amplifier with >17 dBm output power. *Opt. Express* **27**, 293–302 (2019).
47. Gooskens, E. et al. Wavelength dimension in waveguide-based photonic reservoir computing. *Opt. Express* **30**, 15634–15647 (2022).
48. Argyris, A. Photonic neuromorphic technologies in optical communications. *Nanophotonics* **11**, 897–916 (2022).
49. Silva, L. C. B. et al. Integrated photonic neural networks for equalizing optical communication signals: a review. *Photonics* **12**, 39 (2025).
50. Li, S., Dev, S., Kühl, S., Jamshidi, K. & Pachnicke, S. Micro-ring resonator based photonic reservoir computing for pam equalization. *IEEE Photon. Technol. Lett.* **33**, 978–981 (2021).
51. Sarantoglou, G., Sozos, K., Kamalakis, T., Mesaritakis, C. & Bogris, A. Experimental demonstration of an extreme learning machine based on Fabry Perot lasers for parallel neuromorphic processing. In *Proc. Optical Fiber Communication Conference (OFC) 2022 M1G.3* (Optica Publishing Group, 2022).
52. Sozos, K. et al. Recurrent optical spectrum slicing receiver for power fading mitigation in highly dispersive links using programmable photonics. In *Proc. IEEE Photonics Conference (IPC)* 1–2 (IEEE, 2024).
53. Staffoli, E., Mancinelli, M., Bettotti, P. & Pavesi, L. Equalization of a 10 Gbps IMDD signal by a small silicon photonics time delayed neural network. *Photon. Res.* **11**, 878–886 (2023).
54. Tang, X. et al. C-band 56-Gb/s PAM4 transmission over 80-km SSF with electrical equalization at receiver. *Opt. Express* **27**, 25708–25717 (2019).
55. Yin, M., Zou, D., Wang, W., Li, F. & Li, Z. Transmission of a 56 Gbit/s PAM4 signal with low-resolution DAC and pre-equalization only over 80 km fiber in C-band IM/DD systems for optical interconnects. *Opt. Lett.* **46**, 5615–5618 (2021).
56. Wang, H. et al. Adaptive channel-matched detection for C-band 64-Gbit/s optical OOK system over 100-km dispersion-uncompensated link. *J. Lightwave Technol.* **38**, 5048–5055 (2020).

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) 2026

Methods

Setup

The experimental setup is depicted in Supplementary Fig. 1. A 1,550-nm continuous-wave laser is modulated using an OOK modulator to generate a 28-Gbps intensity-modulated optical bitstream. To generate this bitstream, the Mersenne Twister algorithm is used as a generator, as opposed to a pseudo-random binary sequence signal. This is done to avoid that the system learns the pseudo-random binary sequence rule being used instead of being able to equalize the actual signal. As such, the training and testing data are generated using different initialization seeds. To precisely control the input power without affecting the system's noise figure, the modulated signal first undergoes a fixed amplification to 20 dBm, followed by tuneable attenuation. The amplification is provided by an erbium-doped fibre amplifier. As the signal propagates through the optical fibre, dispersion introduces distortions that necessitate compensation. Because of design and fabrication imperfections, the photonic integrated circuit experiences an insertion loss of 30 dB, and therefore, the signal needs to be amplified accordingly. Performance is optimal when this is done through both pre- and post-amplification. Erbium-doped fibre amplifier 2 amplifies the signal up to 15 dBm, whereas erbium-doped fibre amplifier 3 amplifies the output of the reservoir to 10 dBm, an optimal power level for the high-speed photodetector. The photocurrent of this detector is captured by a real-time oscilloscope operating at 160 gigasamples s^{-1} . The acquired data are processed using an iterative algorithm (to be discussed later) that dynamically adjusts the driving currents for the photonic integrated circuit's weight elements, forming a closed-loop optimization cycle. It should be noted that the amplifiers add amplified stimulated emission noise to the signal, degrading the performance of the equalizer. A critical component of the setup is polarization control. The photonic circuit integrates both edge couplers and grating couplers, the latter being highly polarization sensitive. Additionally, on-chip waveguides exhibit polarization-dependent propagation characteristics, with transverse-magnetic modes suffering from substantially higher losses. To mitigate polarization-induced performance degradation, a polarization controller is used before the circuit, ensuring that the input signal is appropriately aligned for optimal transmission through the device.

To enable a direct performance comparison between the photonic system and an electrical FFE based on a tapped delay line, two different baselines are used. One baseline compares an FFE applied on the signal with an equivalent signal-to-noise ratio at the receiver as when the chip is used. To achieve this, a secondary signal path is introduced in which the photonic reservoir is replaced by a 25-dB optical attenuator. This configuration maintains comparable power levels and eliminates the impact of the photonic reservoir's transformation. The other baseline used is to apply an FFE on a signal without this extra insertion loss, to mimic what would be detected without an optical solution. This also allows us to skip the third amplifier in the setup, which increases the signal quality. Furthermore, all post-processing techniques applied to the reservoir output are identically applied to these baselines, ensuring an unbiased evaluation of system performance.

Post-processing

The data processing pipeline is different for the two successive key stages in the experiment: (1) training, which determines the optimal currents for driving the weight elements on the photonic chip, and subsequently, (2) inference, where the received signal is processed by the reservoir using the weights determined in the training step. It is important to emphasize that all processing steps used in this work beyond the reservoir are standard techniques in IM/DD systems and are not introduced specifically due to the presence of the photonic chip.

Training. During training, we send a known pattern of 20,000 bits through the system and adapt the weights of the system in such a way that the error rate is minimal. The following data processing steps

are also presented schematically in Supplementary Fig. 2. Fourier resampling, synchronization and thresholding are performed both during training and inference. Filtering and normalization take place only during inference. Due to the non-integer relationship between the real-time oscilloscope sampling rate (160 gigasamples s^{-1}) and the data rate (28 Gbps), the signal is first resampled via NumPy's Fourier interpolation to ensure a consistent sampling grid aligned with the bit period. Note that in an actual system with an integrated receiver, this step would not be necessary. The resampled signal is then synchronized with the known training pattern we sent through the link using a predefined header consisting of 1,000 consecutive zero bits at the start of the transmission. This is important, since we are constantly sending data to the oscilloscope, but without this synchronization, the machine learning algorithm would not be able to tell what the first bit of the signal is, and therefore, it would not have the correct labels to perform the training. Once synchronization is achieved, the signal is thresholded to convert it into a bitstring. Finally, the BER is computed through straightforward counting as the primary performance metric.

To optimize the read-out layer weights, we use the covariance matrix adaptation evolution strategy⁵⁷, a robust evolutionary optimization algorithm designed to efficiently refine high-dimensional parameter spaces. The optimization process spans only 750 iterations, structured into 15 generations with a population size of 50. To accelerate convergence, the standard deviation of the algorithm is manually tuned from an initial value of 3 to a final value of 0.3, progressively narrowing the search space. We verified that the 20,000 bits that are used for training in these experiments are sufficient to ensure a well-conditioned model for subsequent inference. For every iteration of the training process, the signal propagates through the complete system, where the chip processes the signal in real time. In each iteration, the weights are adjusted, and the signals are sent through the chip again. From the perspective of a telecom operator, this training period should be reduced as much as possible, since during training, no actual useful data can propagate through the system and services would have to be paused. From the perspective of performance, however, this training with hardware in the loop has the benefit that any variations and limitations of the setup will be taken into account during each iteration, even if they change slightly over time. In this way, a robust solution can be found.

Inference. During inference, additional processing steps are performed to ensure an accurate estimation of the system's BER. The signal first undergoes resampling using the same methodology as in the training phase, followed by low-pass filtering with a bandwidth equal to the data rate (28 GHz) to suppress high-frequency noise. After filtering, synchronization is carried out analogously to the training stage, to be able to determine the ground truth for the received bits. The signal is then normalized by subtracting the mean and dividing by the standard deviation, a process performed in packets of 1,000 received bits to account for system fluctuations. The final processing step is thresholding. Unlike in the training phase, where the training signal itself is used to determine the optimal decision boundary, inference relies on a separate validation dataset to establish a fixed threshold, which is then applied to all test data. This threshold is determined using a validation set of 1,000 bits. Once thresholding is completed, the BER is computed by comparing the predicted bitstream with the target sequence and counting the errors. Each experiment utilizes a test set comprising 18 million bits. Due to statistical variations, BER estimation is subject to a resolution limit, which, according to the rule of thumb⁵⁸, corresponds to a minimum of 100 erroneous bits, or 5.55×10^{-6} for these experiments.

It is important to emphasize that clock recovery, normalization and thresholding are essential steps that are part of any commercially deployed system. Any practical equalization scheme—whether optical or electronic—requires these operations to recover symbol timing, set the signal swing and make binary decisions. In a deployed system, these

functions are implemented in the analogue front-end (for example, clock and data recovery circuits, transimpedance amplifiers or limiting amplifiers, and comparators) rather than as digital routines as we did in this laboratory setup. This analogue implementation ensures that they incur negligible additional latency and do not undermine the real-time nature of the equalizer. As discussed in the main text, the photonic RC system is benchmarked against a digital FFE. The FFE is a linear filter that constructs a weighted sum of the signal and its delayed copies, with its memory depth determined by the number of taps. Often in the filter design, an odd number of taps is chosen, since for an FFE, this means that the final delay is an integer number of samples, and all frequencies are delayed in the same way. Therefore, in this study, an FFE with 17 taps is used, ensuring a fair comparison, as the photonic reservoir system utilizes only 16 trainable weights. It is also worth pointing out that even an FFE with 16 taps would still exhibit a greater memory capacity than the photonic reservoir because of optical losses, making sure that we do not give preferential treatment to the reservoir in the comparison.

To further maximize the performance of the baseline, the latency between the FFE output and the target sequence is separately optimized. Note that this delay was not optimized for the reservoir, where we match the headers of the measured output signal and the target without adding extra delay. To find the optimal weights for the digital filter and threshold, a validation set is used. Given that the FFE relies on linear regression for weight optimization, a larger validation dataset is required to mitigate underfitting compared with the reservoir case in which the validation set is only used for thresholding. Consequently, the validation set size for the FFE is increased to 4,000 bits, compared with 1,000 bits for the photonic reservoir system.

Note that the resulting BER of the linear equalizer is heavily dependent on the launch power and, therefore, on the nonlinearities of the task. Our goal in these experiments is to compare with a linear equalizer and to compare the performance purely for the considered experiments. It is not our goal to imply superiority in general over FFE. To claim this, these experiments should be performed over a vast range of launch powers and received power at the detector, where only the minimal values of BER should be considered.

Data availability

The data used in this study are available from the corresponding author (R.V.A.) upon reasonable request.

Code availability

The code used in this study are available from the corresponding author (R.V.A.) upon reasonable request.

References

57. Nomura, M., Shibata, M. & Hamano, R. cmaes: A simple yet practical Python library for CMA-ES. Preprint at <https://arxiv.org/abs/2402.01373> (2024).
58. Jeruchim, M. Techniques for estimating the bit error rate in the simulation of digital communication systems. *IEEE J. Sel. Areas Commun.* **2**, 153–170 (1984).

Acknowledgements

We acknowledge useful discussions with E. Vlieg and A. Lupo.

Author contributions

S.S. designed the chip. R.V.A., S.M. and E.G. planned the experiments and worked on optimizing the setup. J.V.K. wire bonded the chip and helped with the equipment. X.Y. and P.B. supervised the project. R.V.A. and S.M. performed the experiments. R.V.A. analysed the results and performed the simulations. R.V.A. and P.B. wrote the paper.

Funding

This work received funding from the European Commission under the Horizon 2020 Framework Programme for Research and Innovation through grant agreement numbers 871658, 860360, 101070195, 101046329, 101098717 and 101194363. E.G. acknowledges support from the Fonds Wetenschappelijk Onderzoek (Research Foundation Flanders).

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/s41566-026-01968-2>.

Correspondence and requests for materials should be addressed to Ruben Van Assche.

Peer review information *Nature Photonics* thanks Antonio Hurtado, Miguel Soriano and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.