Research Article | Vol. 34, No. 21 / 19 Oct 2026 / Optics Express 39389
Optics EXPRESS

Optically writable atomic vapor memory as a substrate for optical reservoir computing

ELIZABETH ROBERTSON,[1,2,3,†] MINGWEI YANG,[2,†] LINA JAURIGUE,[4] GUILLERMO GALLEGO,[3,5] KATHY LÜDGE,[4] AND JANIK WOLTERS[1,2,5,6,*]

[1]Institute of Space Research, German Aerospace Center (DLR), Berlin, Germany
[2]Institute of Physics, Technische Universität Berlin, Berlin, Germany
[3]Department of Electrical Engineering and Computer Science, Technical University of Berlin, Berlin, Germany
[4]Institut für Physik, Technische Universität Illmenau, Illmenau, Germany
[5]Einstein Center Digital Future (ECDF), Berlin, Germany
[6]AQLS UG Haftungsbeschränkt, Germany
[†]Both authors contributed equally.
[*]janik.wolters@dlr.de

Abstract: We present an optical random access memory (ORAM) based on warm cesium (Cs) atomic vapor and demonstrate its operation as the physical substrate of a reservoir computer. Information is stored in the hyperfine population distribution of a Cs ensemble via optical pumping and retrieved through differential probe absorption. Spatial multiplexing via acousto-optic deflection provides eight addressable memory rails able to store up to 3.9 bits of information per rail. Employing this platform as a temporally multiplexed reservoir, we achieve a kernel rank (KR $=8.8 \pm 0.4$), and a minimum bit error rate of 0.02 ± 0.01 on the Exclusive-or (XOR) benchmark. We find the limited memory lifetime constrains the achievable temporal depth, encouraging further research into fast addressable memories. This constitutes the first demonstration of a free-space, optically writable atomic RAM as a substrate in an optical reservoir computing system.

Published by Optica Publishing Group under the terms of the Creative Commons Attribution 4.0 License. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

1. Introduction

Memory is a fundamental element of information processing, as evidenced by the central role of the read-write tape in the Turing machine's definition of classical computability [1]. Historically, encoding information in classical optical signals has proven to be a highly viable storage method. Optical Read-Only Memories (OROMs)—spanning analog formats like paper and photographs to robust digital media such as CDs, DVDs, and Blu-ray discs—are mature, commercialized technologies. Moreover, the high degree of parallelism offered by optical memory modules [2] continues to drive research into high-volume archival storage for data centers [3,4]. In contrast, optical random access memories (ORAM) have not seen the same commercial success. This is primarily due to the overwhelming dominance of electronic counterparts (e.g., SRAM, DRAM, and flash memory), which provide versatile, multi-timescale storage natively integrated on silicon chips. However, as the computing industry increasingly pivots toward domain-specific hardware accelerators e.g. tensor processing units (TPUs) [5], in-memory compute architectures [6,7], and optical vector-matrix multipliers [8], compatible memory architectures must be researched in tandem. Optical neural network accelerators have been identified as a particularly promising system, due to their low latency, highly parallel, low-power operation [9–11]. However, realizing the full energy advantage of these architectures requires that both data processing *and* data storage

#612547 https://doi.org/10.1364/OE.612547
Journal © 2026 Received 30 Jul 2026; revised 10 Sep 2026; accepted 21 Sep 2026; published 9 Oct 2026

Optics EXPRESS

remain optical, since repeated conversion to and from electronic memory is power-expensive, limiting the power savings that optical hardware seeks to improve [12]. Optical random access memories (ORAMs) are therefore a necessary prerequisite for all-optical computing [13,14].

Within the broad field of optical computing, optical neural networks and specifically optical reservoir computing has been identified as a promising platform for high-throughput, low-energy machine learning [15]. A key characteristic of the reservoir computer (RC) and a distinguishing feature from other architectures, such as an extreme learning machine (ELM) [16], is its fading memory. Memory may be actively included into optical reservoir computers, explicitly as an element for introducing delay e.g. delay-line systems [17], or alternatively it may arise as an intrinsic feature of the hardware substrate used as a device, often seen in spatially distributed RCs [18]. Moreover, memories are required to store the information input into reservoir computers. Often, electro-optic devices are used to modulate input information, necessitating storage in the electrical domain; however, all optical architectures seek to remove this conversion. Here, optical memories which can be read or written to optically are required. In spatially distributed optical reservoir computers, for example, data are commonly injected via digital micromirror devices (DMDs) or spatial light modulators (SLMs), which are electronically writable and introduce conversion overhead [19,20]. Optically addressable SLMs (OASLMs) could address this limitation [21], but require a native optical memory to buffer and update weight data. In delay-line reservoir computers, the optical fiber loop that provides recurrence acts as a strict first-in first-out (FIFO) queue, fixing virtual node connectivity to a simple ring topology. A random-access optical memory replacing the delay-line could enable arbitrary inter-node coupling, enabling investigation into potentially more varied reservoir dynamics. A preliminary investigation into this system is the subject of this work.

Some exemplary physical platforms that have been studied as optically writable ORAMs include phase change materials (PCMs) [22–24], and different devices that form optical bistable memories [13,25] for integrated memory. For free-space systems, optically addressable liquid-crystal cells [21,26,27] and metasurfaces [28,29] have been investigated. Many of these devices have demonstrated successful operation in a research setting, but require further development to be used in real world applications.

In this work, we employ a room temperature cesium (Cs) atomic vapor as a free-space storage medium. Compared with fabricated storage platforms, atomic vapor offers high uniformity because all Cs atoms have identical properties, eliminating fabrication-induced variations and making peripheral devices the dominant source of non-uniformity. The inherently continuous three-dimensional spatial structure of the vapor cell also enables scalable multimode operation without the need for microfabrication. Importantly, the nonlinear saturation absorption of the atomic transition provides a native nonlinear response within the memory, enabling in-memory computation and furnishing the activation function for reservoir computing.

A coherent Cs vapor RAM has been demonstrated in Ref. [30]; this work presents an incoherent memory storage mechanism and its apparently novel application as the reservoir in reservoir computing. Importantly, coherent memories require phase coherence between the stored pulses, whereas incoherent memories do not, making the latter technically simpler and thus more suitable for practical applications.

In the following section, we present the design, characterization, and reservoir computing application of a free-space Cs vapor ORAM. After introducing the memory operating principle and experimental apparatus (Section 2), we describe the use of a Cs vapor ORAM as an element in a reservoir computing architecture and its performance metrics (Section 3), and evaluate the reservoir against the linear memory capacity, kernel rank, and exclusive-or (XOR) benchmarks (Section 4).

Optics EXPRESS

2. Memory concept and characterization

2.1. Operating principle

The proposed memory encodes a number to be stored in the hyperfine ground-state population distribution of a warm Cs ensemble. The writing protocol utilizes optical pumping [31], specifically depopulation pumping (or optical shelving), wherein atoms are optically transferred from the upper hyperfine ground state $|s\rangle = |F = 4\rangle$ into a dark state $|g\rangle = |F = 3\rangle$. The population pumping results in a polarized ensemble, where the degree population polarization P indicates the magnitude of the number stored. To read from the memory, the transmission behavior of the probe field is measured as a proxy for the population residing in the dark state. The protocol for reading and writing to the memory is illustrated in Fig. 1. Here, a numerical value is encoded in the atomic vapor using a strong laser pulse tuned to the Cs D1 $F = 4 \rightarrow F' = 4$ transition, with the pulse intensity modulated according to the number to be stored. To read from the memory cell, a subsequent weak probe pulse of the same frequency but below the saturation intensity (I_{sat}), is incident on the cell. The probe light is transmitted with an efficiency proportional to the population that remains in the $|s\rangle$ state. Consequently, higher pump intensities yield greater transmission, representing the storage of larger numbers. A key feature is the *accumulative* nature of writing in the linear regime: successive pump pulses pump more of the atomic population to the $|g\rangle$ state, provided the pump amplitudes remain below I_{sat}. Furthermore, when the effective pump intensity exceeds the saturation threshold, the ensemble enters a nonlinear regime characterized by its saturation absorption profile. Memory reset occurs passively as polarized atoms diffuse out of the laser interaction region and are replaced by thermally distributed atoms from the surrounding vapor. While the medium could be actively reset via a repumper laser on the $|F = 3\rangle \rightarrow |F' = 4\rangle$ transition, the necessary calibration and further investigation falls outside the scope of this work.

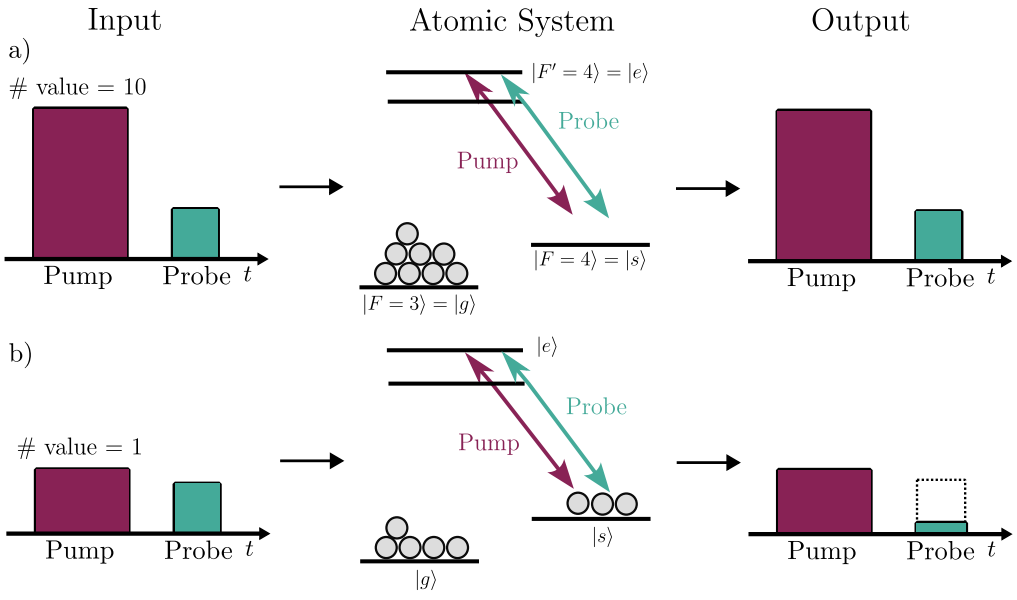


Fig. 1. Operating principle of the incoherent atomic ORAM. (a) A strong pump pulse transfers atoms to the dark state $|g\rangle$, rendering the medium transparent to a subsequent weak probe. (b) A weaker pump transfers fewer atoms, yielding lower probe transmission. The stored analog value is encoded in the measured probe intensity.

Optics EXPRESS

To store multiple values in a random access manner, we spatially multiplex the Cs vapor cell by deflecting the read and write pulses using an acousto-optic modulator (AOM) [30]; an example configuration with the two outermost memory rails being addressed simultaneously is shown in the experimental setup (Fig. 2(a)). By modulating the amplitude of the RF waves driving the AOMs, pulses can be generated. A typical modulation waveform for the input AOM (AOM1) and the collection AOM (AOM2) can be seen in Fig. 2(b). Different spatial locations can be accessed simultaneously by applying multi-frequency driving waves. Modulation waveforms corresponding to the different values to be written to each rail are applied to RF driving frequencies, which determine the spatial location that forms individual memory rails. These individual frequency components are summed to access multiple rails simultaneously. To avoid confusion, in this work we distinguish between the storage medium (denoted cesium "cell") and addressable memory location (memory "rail").

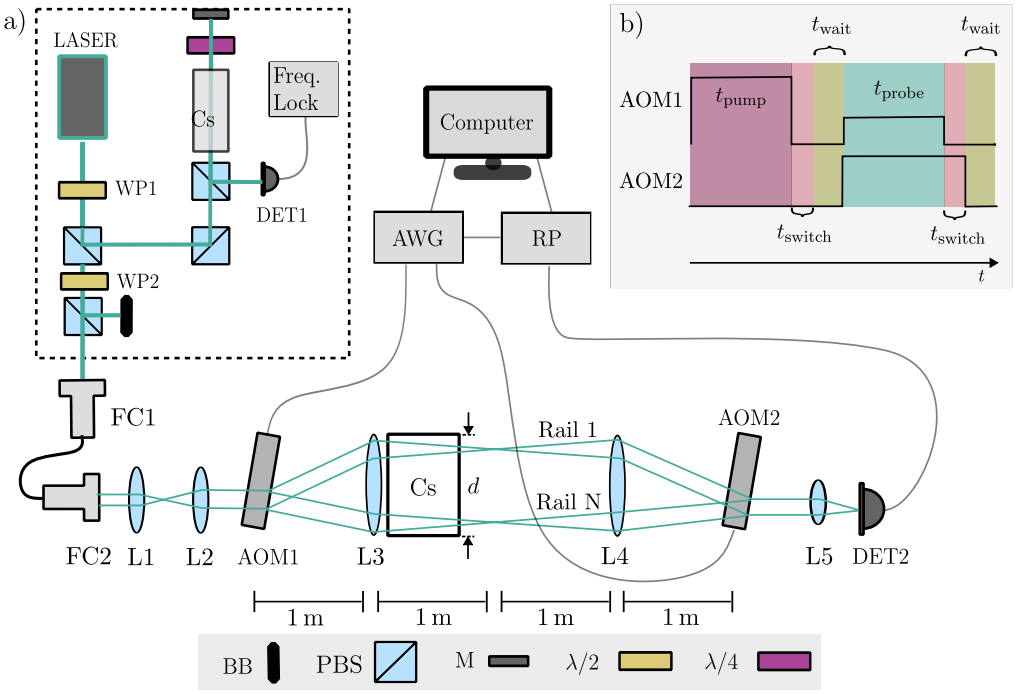


Fig. 2. a) Experimental setup of the Cs ORAM. The laser is frequency-locked via Doppler-free spectroscopy (dashed box). After magnification from lenses L1 and L2, the pump/probe beam is diffracted by AOM1 into N spatial rails that traverse the Cs cell with diameter $d = 5$ cm; AOM2 recombines the modes onto a single detector path. Figure not to scale. FC: Fiber coupler; L: Lens; DET: Detector; AWG: arbitrary wave generator; PBS: polarizing beamsplitter; BB: beam block; M: mirror; RP: Red Pitaya; WP: waveplate; $\lambda/2$: half-wave plate; $\lambda/4$: quarter-wave plate. b) Single-rail pump-probe timing sequence. AOM2 is turned off during pumping to prevent damage to the detector.

2.2. Experimental setup

The full experimental setup is shown in Fig. 2. An external-cavity diode laser (ECDL; Sacher Lasertechnik Micron) is locked to the $F = 4 \rightarrow F' = 4$ Lamb-dip of a Doppler-free spectroscopy using a Red Pitaya STEMlab 125-14 running the open-source *Linien* software [32]. The beam is then expanded by a telescope formed of L1 and L2 to a $1/e^2$ diameter of 3.1 mm and directed

Optics EXPRESS

to the first acousto-optic modulator (AOM1; AA Optoelectronics DTSX-400-900). The beam power is adjusted to $P_{\text{in}} = 10\,\text{mW}$ at the fiber coupler output (FC2, Schäfter + Kirchhoff 60FC-SP-4-M8-10) using a polarizing beamsplitter (PBS) and a half-wave plate (WP2). AOM1 is driven by a multi-tone RF signal from a Zurich Instruments HDAWG-4 arbitrary waveform generator, which simultaneously generates eight carrier frequencies ($f_{\text{RF}} = 75, 78, \ldots, 96\,\text{MHz}$, separated by 3 MHz to avoid crosstalk), each amplitude-modulated independently to produce pump and probe pulse pairs at eight distinct rails within the vapor cell. A 4f relay system (lenses L3, L4; $f = 1\,\text{m}$) images the spatially separated beams through the Cs cell (Precision Glassblowing; 10 cm length, 5 cm diameter, 10 Torr N_2 buffer gas at 30°). A second AOM (AOM2) recombines all spatial modes onto a single avalanche photodetector (DET2; Thorlabs APD430A/M) and acts as an optical gate to protect the detector during high-power pump pulses (Fig. 2(b)). Probe traces are acquired using the deep acquisition mode of a Red Pitaya STEMlab 124-14 board and transferred to a host computer for further application specific processing.

In this work, unless otherwise stated, the rails are pumped for $t_{\text{pump}} = 10\,\mu\text{s}$, and probed for $t_{\text{probe}} = 10\,\mu\text{s}$, with switching and waiting intervals of $t_{\text{switch}} = 2\,\mu\text{s}$ and $t_{\text{wait}} = 3\,\mu\text{s}$. This yields a memory operation time of $t_{\text{pp}} = 30\,\mu\text{s}$, defined as the time required to execute a single read and write operation. Notably, the architecture supports simultaneous write operations across multiple rails. The amplitude of the probe pulses, corresponding to the value read out of our memory is determined by averaging over a 2 μs subsection centered with the probe interval of the measured pulse. In this system, we deflect the beams along a single axis, with the driving electronics limiting the number of resolvable spots to 8 rails. In future experiments without RF-electronics limitations, a $1/e$ beam diameter of $d = 3.1\,\text{mm}$ would allow $N = 10$ rails within the cell aperture. Increasing the number of rails by reducing the beam diameter comes at the expense of the memory lifetime T_1. Replacing the single-axis AOMs and corresponding electronics with 2D AOMs would enable up to 64 memory modes, making the cell aperture the limiting factor.

2.3. Memory characterization

To characterize the memory operation, the pump pulse energy is varied while monitoring the response of the probe pulse. The experimental sequence for each rail consisted of an initial 10 μs probe pulse, a 5 μs delay, a 10 μs pump pulse, and a second 5 μs delay before a second 10 μs probe, for each rail. The transmitted intensity of this second probe pulse is shown in Fig. 3. To ensure consistency, the AOM2 RF driving power is normalized such that a single probe pulse yielded a uniform photodiode voltage. Throughout the measurements, the probe power is maintained at 50 μW, and the pump power is varied between $0 - 1\,\text{mW}$ in 50 μW increments, and scanned from $1.5 - 4\,\text{mW}$ in 500 μW steps. This increased step size at higher powers is adopted to optimize data acquisition time. The response of the probe transmission to increasing pump power follows a saturable absorber profile of the form $\theta(P_{\text{pump}}) = a\frac{P_{\text{pump}}}{1+P_{\text{pump}}/P_{\text{sat}}} + b$. In the linear regime ($P_{\text{pump}} \ll P_{\text{sat}}$), probe transmission increases approximately linearly with pump energy. Successive pump pulses result in the accumulation of the population, resulting in in-memory addition, provided the pulses are weak enough to not enter the saturated regime. At higher pump powers, the atomic population available for transfer to $|g\rangle$ saturates, producing a rectified response. We measure the average pump saturation power across all rails as $P_{\text{sat}} = 1.5\,\text{mW}$. We see that the normalization of the rails at the edges of the AOM operating range, and for the 87 MHz rail was not successful. We attribute this to large variation in the diffraction power of the outermost rails. The maximum pump power is set to $P_{\text{max}} = P_{\text{sat}}$ to access the onset of saturation while remaining within a bounded dynamic range. A tunable input scaling factor $g_1 \in [0, 2.4]$ controls the fraction of P_{max} applied to each pulse, allowing systematic variation of the effective nonlinearity.

To quantify the resolution of the system, we measure the difference in transmitted probe amplitude between the pumped and unpumped states, for a pump pulse of $P_{\text{pump}} = 1.5\,\text{mW}$.

Optics EXPRESS

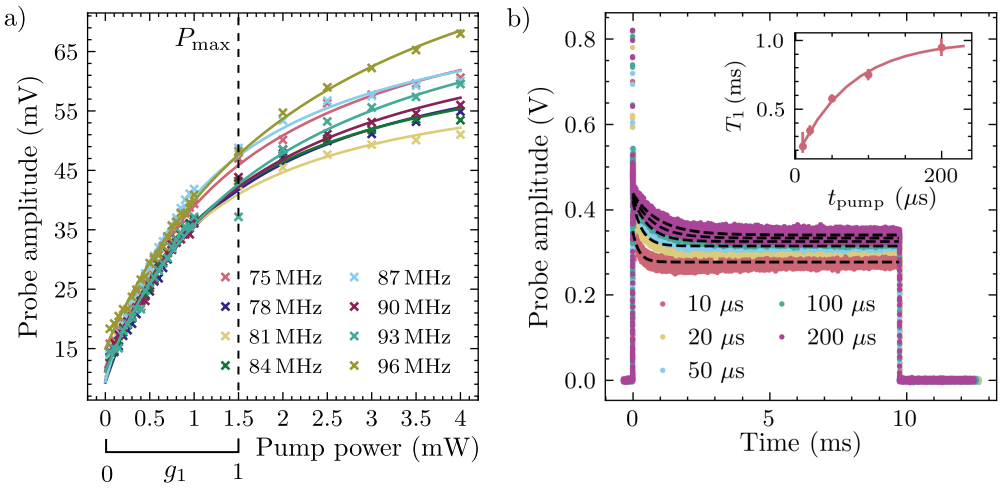


Fig. 3. Characteristic memory responses (a) Saturation absorption transfer function of the eight ORAM rails. Experimental data are fit with a saturable absorber profile, which defines the nonlinear function θ used in reservoir simulations. Mean saturation power: $P_{sat} = 1.5\,\text{mW}$. (b) Decay of atomic polarization for pump durations $t_{pump} \in \{10, 20, 50, 100, 200\}\,\mu\text{s}$ at $P_{pump} = 1.5\,\text{mW}$. The peak feature at $t = 0$ is leaked pump light indicating the switch from the pump to the probe pulse. **Inset:** Measured T_1 vs. t_{pump}, fit with $T_1 = a(1 - e^{-kx}) + b$; the curve saturates at approximately 1 ms for $t_{pump} \geq 200\,\mu\text{s}$.

We determined that a probe power of $P_{probe} = 50\,\mu\text{W}$ provides the largest dynamic range for a cell temperature $T = 30^\circ\,\text{C}$, which is the probe power used in the rest of this work unless otherwise specified. Using a pump pulse duration $t_{pump} = 10\,\mu\text{s}$, we measured a difference in detected photodiode voltage of between $U = 15\,\text{mV}$ and $U = 44\,\text{mV}$, with a standard deviation of $\sigma = 1\,\text{mV}$. The observed fluctuations arise from technical noise in the AOMs, where thermal variations in deflection efficiency are most pronounced at low RF driving powers. These effects may be further reduced through temperature stabilization of the AOMs or by decoupling the pulse generation process and the deflection to different rails. We define two signal values as resolvable if they differ by at least 2σ. The full range then contains $N = \Delta U/2\sigma = 14.5$ resolvable levels, or an effective resolution of $\log_2 N = 3.9$ bits. This resolution is time-dependent; the value reported here corresponds to a readout performed $3\mu\text{s}$ after pumping; we discuss the decoherence mechanism further in the next section.

2.4. *Memory lifetime and time-bandwidth product*

The stored population decays as polarized atoms diffuse out of the laser interaction volume and are replaced by atoms in thermal equilibrium. This effect occurs continuously from the moment the atoms are pumped, thus meaning the polarization state of the atoms is continuously depleting. Charge depletion is a known phenomenon in electronic dynamic random access memory (DRAM), overcome by periodic re-charging. The diffusion-limited lifetime is estimated as $\tau_{mem} = d_{beam}^2/4D \approx 1.3\,\text{ms}$ for our beam diameter $d_{beam} = 3.1\,\text{mm}$ and Cs-in-N_2 diffusion coefficient $D = 17(10)\,\text{cm}^2\,\text{s}^{-1}$. We measure the $1/e$ lifetime T_1 by pumping a single rail (the central rail at 81 MHz) for $10\,\mu\text{s}$, with a pump power of 1.5 mW. The transmission of a $\sim 10\,\text{ms}$ weak probe pulse ($P_{probe} = 9\,\mu\text{W}$) is measured for pump pulses of varying durations.

The measured $1/e$ lifetimes are $T_1 = \{230(50), 340(40), 570(30), 750(40), 950(60)\}\,\mu\text{s}$ for pump durations $t_{pump} = \{10, 20, 50, 100, 200\}\,\mu\text{s}$ respectively. We observe that the duration of

Optics EXPRESS

the pump saturates to approximately 1 ms with increasing pump pulse duration, as shown in the inset of Fig. 3.

To quantify the information storage capacity, we define the memory time-bandwidth product $\mathrm{TBP}_{\mathrm{mem}}$ as the number of pump-probe operations executable within the $1/e$ lifetime:

$$\mathrm{TBP}_{\mathrm{mem}} = \frac{T_1}{t_{\mathrm{pp}}}, \tag{1}$$

where t_{pp} is the duration of a single pump-probe cycle. For our system, $t_{\mathrm{pp}} = 30\,\mu$s (including a 10 μs pump, 10 μs probe, and switching/wait times), yielding $\mathrm{TBP}_{\mathrm{mem}} = 7.7$ at $T_1 = 230\,\mu$s. This presents a second limitation in the number of memory rails which are accessible within the storage time. To increase the number of accessible memory modes, one must either increase the storage time or decrease the time of a single read-write (pump-probe) operation. The storage time may be extended using an array of miniaturized Cs cells, where the pump/probe beam fully encompasses the atomic volume such that the atoms cannot diffuse out of the interaction region. These cells must be filled with an anti-relaxation coating to prevent collisions with the cell walls becoming a limiting mechanism and are a topic of current active research [33–35]. The duration of the pump and probe sequence could be reduced by utilizing AOMs with a higher acoustic velocity, and thus shorter rise times on the order of tens of nanoseconds, thereby enabling shorter pulses. However, the gain in addressing speed comes at the expense of the number of addressable rails. The Bragg deflection angle of the AOM scales inversely with the acoustic velocity, so that a shorter rise time also results in a reduced angular range for a given RF bandwidth. Alternatively, electro-optic deflectors, which have rise times on the order of 100 ns, may also be explored as a solution.

The key figures of merit are summarized in Table 1. Having established the characteristics of the medium as a memory element, we now describe the reservoir computing architecture implemented in this work.

Table 1. The key characteristics of the incoherent memory experiment.

Figure of merit	Value
Memory lifetime (T_1)	230 μs
Resolution	3.9 bits
Number of rails (N)	8
Read/write duration (t_{pp})	30 μs
Memory bandwidth ($\mathrm{TBP}_{\mathrm{mem}}$)	7.7

3. Reservoir computing architecture and metrics

In the previous section, we introduced the operating principle, experimental setup, and characterization of the incoherent optical RAM. We now turn to its application as a reservoir, and present the key metrics which allow us to evaluate its performance within a reservoir computing architecture.

3.1. Memory-based reservoir computing

Physical reservoir computing (RC) exploits the nonlinear transient dynamics of a fixed physical system to map low-dimensional inputs into a high-dimensional state space, after which only a single linear readout layer requires training [36]. Optical delay-line reservoirs achieve recurrence via a fiber loop that functions as a FIFO queue [37]; virtual nodes are accessed sequentially, and connectivity is topologically fixed. By replacing the delay line with a multi-rail ORAM,

Optics EXPRESS

arbitrary read/write scheduling becomes possible, enabling flexible intra-node coupling. In temporally multiplexed systems such as delay-line reservoirs and the system presented in this work, information is injected into and read sequentially from the reservoir. Hence, we define the time scales relevant to the system in question; the input dataset $\mathbf{u}$ is injected into the reservoir as a piecewise function, with each element of $\mathbf{u}$ being injected with duration T, indexed by a discrete time t. Consequently, $\mathbf{u}(t = 0)$ corresponds to injection of the first logical data point. The number of data points injected into the reservoir is given by the training (test) set size $N_{\text{train}}(N_{\text{test}})$. We define virtual nodes by further subdividing the clock cycle into N_v distinct segments, where N_v is the length of a mask weight vector $\mathbf{w}_{\text{mask}}$. Each masked element is injected into the reservoir with duration T_v, such that $N_v = T/T_v$, as is typical in temporally multiplexed reservoir computing [38–40]. Experimentally, we set $T_v = t_{\text{pp}} = 30\,\mu$s, the time to execute a single read and write operation on a memory rail. The value of a single element injected into the reservoir is given by: $\mathbf{j}(k) = \mathbf{u}(t) \cdot \mathbf{w}_{\text{mask}}(k')$ where $k' \in (0, N_v]$ is the mask or virtual node index, and the system index $k = N_v t + k'$. The mask is drawn from a uniform distribution in the interval $[0, 1]$; differing from the convention for the mask interval of [-1,1] as our hardware cannot represent negative numbers. The input layer of the reservoir computer is depicted in Fig. 4

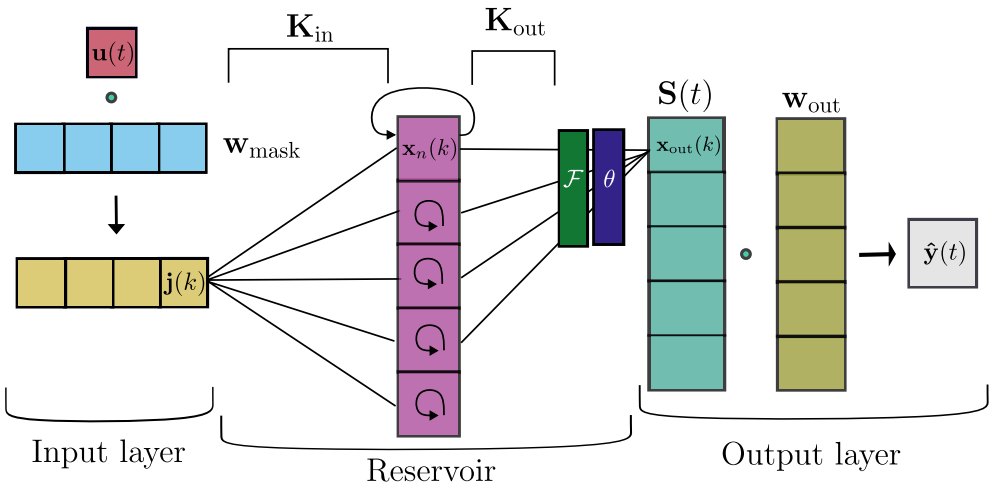


Fig. 4. Schematic of the memory-based reservoir computer. At each time step, the memory rails are first probed (read) before being pumped (written). The coupling matrices $\mathbf{K}_{\text{in}}$ and $\mathbf{K}_{\text{out}}$ determine which rails are addressed, establishing the effective network topology.

We inject the masked input sequentially into the reservoir and fill the state matrix $\mathbf{S}$ with responses. The state matrix is then normalized such that the entries are between $\mathbf{x}_{\text{out}} \in [0, 1]$. We include an additional bias vector, resulting in a $N_{\text{train}} \times (N_v + 1)$ dimensional state matrix. This is shown in Fig. 4 as the output layer. To train the reservoir an output weight vector $\mathbf{w}_{\text{out}}$ is found that minimizes the cost function:

$$\arg\min_{\mathbf{w}_{\text{out}}} \left(||\mathbf{S}\mathbf{w}_{\text{out}} + \mathbf{y}||_2^2 - \lambda_{\text{reg}} ||\mathbf{w}_{\text{out}}||_2^2 \right) \tag{2}$$

where $||\cdot||_2$ is the Euclidean norm, λ_{reg} is the regularization parameter and $\mathbf{y}$ is the true value of the task for inputs $\mathbf{u}$.

Having discussed information injection and population of the state matrix from the reservoir output, the next section details how to utilize the memory presented in Section 2 as a reservoir. The following sections then describe the training process and the metrics used to evaluate the reservoir's performance.

Optics EXPRESS

3.2. Reservoir model

In our reservoir architecture, the physical memory rails are temporally multiplexed to create a high-dimensional output state space. The masked input $\mathbf{j}(k)$ is written to a rail determined by the input coupling matrix $\mathbf{K}_{\text{in}}$. $\mathbf{K}_{\text{in}}$ contains values between $[0, 1]$, which represent the strength with which a specific memory rail should be written to at timestep m, where $m = k \mod M$. M is a coupling matrix repetition index, defining after how many time steps the coupling matrix should repeat. We set $M = N = 8$. This choice of M enables easier interpretation of the relation between the reservoir performance and virtual node dynamics arising from specific N_v. $\mathbf{K}_{\text{in}}$ is set to the identity matrix with dimension $(M \times N)$, such that a single rail was addressed at each timestep. A similar approach using coupling matrices to determine the input and output coupling was presented in Ref. [41]. We model the state of a specific rail x_n, as the effective pump energy of that rail; an overview of the architecture can be seen in Fig. 4. As the atomic population state decays with the storage time, the state of the memory rail from the previous timestep is multiplied by a decay factor $\mathcal{F}(\Delta k_n)$, modeling the exponential population decay between successive writes. This is given by the exponential decay as observed in Section 2.4, where $\mathcal{F}(\Delta k_n) = \exp(-\Delta k_n\, t_{\text{pp}}/T_1)$ and Δk_n is the number of steps elapsed since rail n was last written.

The state of rail n evolves as:

$$\mathbf{x}_n(k+1) = \left(\mathbf{K}_{\text{in}}\right)_n^m (g_1 \mathbf{j}(k+1) P_{\max}) + \mathcal{F}(\Delta k_n) \cdot \mathbf{x}_n(k), \tag{3}$$

At each timestep, we can also choose the memory rails which are read from by defining the $\mathbf{K}_{\text{out}}$ coupling matrix. The reservoir output is:

$$\mathbf{x}_{\text{out}}(k) = \sum_{n=1}^{N} \theta\left(\left(\mathbf{K}_{\text{out}}\right)_n^m \cdot \mathcal{F}(\Delta k) \cdot \mathbf{x}_n(k)\right), \tag{4}$$

where θ is the empirically measured characteristic saturation absorption transfer function (Fig. 3). The states $\mathbf{x}_{\text{out}}$ are then used to populate the state matrix $\mathbf{S}$. Note that the memory cells are first read from before they are written to. Consequently, we set $\mathbf{K}_{\text{out}}$ to the identity matrix, with dimension N, rolled over by one, such that the final memory cell is read first. This results in reading the cell that was just written out at each time step, increasing the effective memory capacity of the system. .

In this work, the network topology is controlled by varying (N_v) relative to (N). A mismatch such that $(N_v = N \pm 1)$ shifts the rail-to-virtual-node mapping at each clock cycle, thereby producing a ring-coupling structure between successive virtual nodes. In this preliminary investigation, only the number of virtual nodes is varied to modify the reservoir topology, while $\mathbf{K}_{\text{in}}$ and $\mathbf{K}_{\text{out}}$ are kept fixed to limit the scope of the study. Aside from the coupling, we may change the reservoir operational behavior by varying the linearity of the response of a single write action by changing the coefficient g_1. Furthermore, by changing the duration between the pump and probe t_{wait}, we can vary the memory capacity of the system. It is useful to define the logical memory capacity T_{logic}, which corresponds to how many logical inputs $\mathbf{u}$ can be input into the reservoir before the state of the first memory cell written to decays to its $1/e$ value. Thus we define:

$$\text{TBP}_{\text{logic}} = \frac{T_1}{N_v T_v} = \frac{\text{TBP}_{\text{mem}}}{N_v}. \tag{5}$$

The minimum number of virtual nodes considered is $N_v = 3$, yielding a maximum $\text{TBP}_{\text{logic}} = 2.6$. These variables, N_v, g_1 and T_{mem}, define the performance parameter space we explore in this work.

Optics EXPRESS

3.3. Training

The output weight vector $\mathbf{w}_{\text{out}}$ is trained using ridge regression on a state matrix $\mathbf{S}$ assembled from $N_{\text{train}} \approx 300$ input samples. Performance is evaluated on an unseen test set of $N_{\text{test}} = 300$ samples. The size of the train and test set varies by approximately 10% due to data set segmenting; a single experimental run of the reservoir required multiple experimental traces. We find this to have no appreciable effect on system performance.

To train the reservoir, we solve for $\mathbf{w}_{\text{out}}$ using ridge regression on the state matrix generated from the train set. This yields a closed-form solution to the minimization problem given by Eq. (3.1):

$$\mathbf{w}_{\text{out}} = (\mathbf{S}^T\mathbf{S} - \lambda_{\text{reg}}\,\mathbb{1})^{-1}\mathbf{S}^T\mathbf{y}, \tag{6}$$

where $\mathbf{S}^T$ is the transpose of the state matrix, $\mathbb{1}$ is the identity matrix and λ_{reg} is the regularization parameter, set to $\lambda_{\text{reg}} = 5 \times 10^{-6}$.

3.4. Performance metrics

The reservoir's performance is measured by its linear memory capacity (LMC), kernel rank (KR), and its ability to predict the XOR between current and previous inputs. The reported values represent an average over five different $\mathbf{w}_{\text{mask}}$ configurations, where the reported error bars indicate the standard deviation. We discuss these metrics in detail below, beginning with the LMC.

Linear memory capacity (LMC) The LMC quantifies how well the reservoir is able to recall past inputs [42]. To measure the LMC, we inject a random string of numbers $\mathbf{u}$ drawn from a uniform distribution $U[0, 1]$ into the reservoir. A reservoir is trained to recall the input $\mathbf{u}$ n time steps in the past $\mathbf{y}(t) = \mathbf{u}(t - n)$ and the performance of the reservoir in recalling that input is measured by the capacity C_n:

$$C_n = 1 - \text{NMSE}(\mathbf{y}_n, \hat{\mathbf{y}}_n) = 1 - \frac{\sum_{k=1}^{N_{\text{test}}} \|\hat{y}_n^k - y_n^k\|^2}{N_{\text{test}}\ \text{Var}(\mathbf{y}_n)}, \tag{7}$$

where $\mathbf{y}_n$ are the predictions for the input n steps ago and $\text{Var}(\mathbf{y}_n)$ is the variance of the target data set. The linear memory capacity is then the sum of all the time delays considered.

$$\text{MC} = \sum_{n=1}^{N_{\text{lim}}} C_n. \tag{8}$$

We choose the maximum evaluation time $N_{\text{lim}} = 5$. This is sufficient as it exceeds the maximum time-bandwidth product ($\text{TBP}_{\text{logic}} = 2.6$) of the system.

Kernel rank (KR) The KR quantifies the dimension of the feature space into which the reservoir projects the input information [42]. To measure KR, a sequence of random inputs $\mathbf{u}$ is randomly drawn from a uniform distribution $U[0, 1]$ and injected into the reservoir. The resulting reservoir states are collected into a state matrix $\mathbf{S}$, and by performing singular value decomposition, the kernel rank can be determined by taking the rank of the eigenvalue matrix Σ:

$$\mathbf{S} = \mathbf{U}\Sigma\mathbf{V}^{\dagger}, \tag{9}$$

$$r_{\mathbf{S}} = \text{rank}(\Sigma), \tag{10}$$

where $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices containing the left and right singular vectors of the state matrix. In practice, due to noise and experimental imperfections, singular values are rarely exactly zero. Therefore, when reporting KR, we provide a threshold below which SVD values are considered to be zero, in this case to 2% of the maximum singular value. The theoretical

Optics EXPRESS

upper bound of the kernel rank is determined by the number of virtual nodes in the system, here varied between $N_v \in [8, 14]$.

XOR bit error rate (BER) The XOR logic gate is a standard benchmark for reservoir computing systems because its truth table is nonlinearly separable. Consequently, a linear classifier cannot solve this task without the reservoir successfully mapping the input into a higher-dimensional feature space. The task consists of computing the XOR function of two consecutive binary inputs, the target function is thus:

$$\mathbf{y}(t) = \mathrm{XOR}(\mathbf{u}(t), \mathbf{u}(t-1)), \tag{11}$$

where the input sequence consists of binary values $\mathbf{u} \in \{0, 1\}$. The prediction performance is quantified by the bit error rate (BER), defined as the ratio of false classifications to the total number of samples in the test set:

$$\mathrm{BER} = \frac{N_{\mathrm{wrong}}}{N_{\mathrm{test}}}, \tag{12}$$

where N_{wrong} is the number of misclassified inputs and N_{test} is the size of the test set. Since the reservoir prediction $\hat{\mathbf{y}}(t)$ is a continuous variable, it is discretized into binary classes by applying a threshold. Specifically, predictions ≥ 0.5 are assigned to 1 and <0.5 are assigned to 0.

4. Results and discussion

Having motivated the need for a memory architecture with reconfigurable node couplings, presented a memory that offers such reconfigurability, and illustrated how we propose to use it as a reservoir in the previous sections, we now present the results. We evaluate the reservoir performance over three independent parameter axes: network topology (N_v), the decay timescale ($\mathrm{TBP}_{\mathrm{mem}} = T_1/t_{\mathrm{pp}}$, varied via t_{wait}), and the input nonlinearity (g_1).

4.1. Linear memory capacity

To understand the memory behavior of the reservoir, we first consider the linear memory capacity as a function of N_v and $\mathrm{TBP}_{\mathrm{mem}}$, as shown in Fig. 5. From simulation (Fig. 5(a)), we predict the highest achievable memory capacity to be LMC = 1.8(5) at $t_{\mathrm{wait}} = 3\,\mu$s ($\mathrm{TBP}_{\mathrm{mem}} = 7.7$), $N_v = 10$ corresponding to a $\mathrm{TBP}_{\mathrm{logic}}$ of 0.8. Experimentally, we find that the optimal operating point is $N_v = 14$, $t_{\mathrm{wait}} = 18\,\mu$s ($\mathrm{TBP}_{\mathrm{mem}} = 3.8$) yielding LMC = 1.0(2), which corresponds to a TBP of $\mathrm{TBP}_{\mathrm{logic}} = 0.29$. This indicates that the time bandwidth product alone is not the sole indicator for memory capacity. A sufficiently high dimensionality in the state space is required to separate the trajectories corresponding to different time scales; the dimension of the network (N_v) influences the linear memory capacity. Consequently, lower dimensional systems (N_v<N) perform poorly despite having a higher $\mathrm{TBP}_{\mathrm{logic}}$. Both simulation and experiment show improved memory capacity for virtual node numbers of $N_v = 10, 12, 14$. We also observe a performance reduction due to resonances within the reservoir; when the number of virtual nodes is an integer number, or a simple fraction of the number of memory rails, $N_v = N = 8$, or $N_v = 0.5N$, the reservoir performance drops significantly in both simulation and experiment. These resonance regimes are known to degrade performance in delay-based systems, due to a reduction in the effective richness in the reservoir dynamics [39,43,44].

The fundamental constraint is that the logical time-bandwidth product for N_v>8 at $\mathrm{TBP}_{\mathrm{logic}}$<1 under current operating conditions. This indicates that the state of the first virtual node decays to below $1/e$ before the complete set of virtual nodes comprising a single logical input has been injected—severely restricting the temporal depth of the reservoir. For comparison, delay-line reservoirs routinely achieve LMCs of order 10 or more [43]. Investigations into extending the memory capacity by increasing the pump time to $t_{\mathrm{pump}} = 200\,\mu$s ($T_1 = 950\,\mu$s) yielded no improvement; the maximum drops to LMC= 0.36 ± 0.18. Although a longer lifetime increases

Optics EXPRESS

$\mathrm{TBP}_{\mathrm{mem}}$, the accompanying increase in t_{pp} reduces $\mathrm{TBP}_{\mathrm{logic}}$ (Eq. (5)), negating the benefit. Thus, to achieve higher LMC, we must find alternate means to increase the $\mathrm{TBP}_{\mathrm{mem}}$ of the system, for example by reducing t_{pp} by using AOMs with a shorter rise time and with higher energy pump pulses. The limited LMC is characteristic of a weakly recurrent reservoir and thus close to the memory-less non-linear feature mapping observed in extreme learning machines (ELMs).

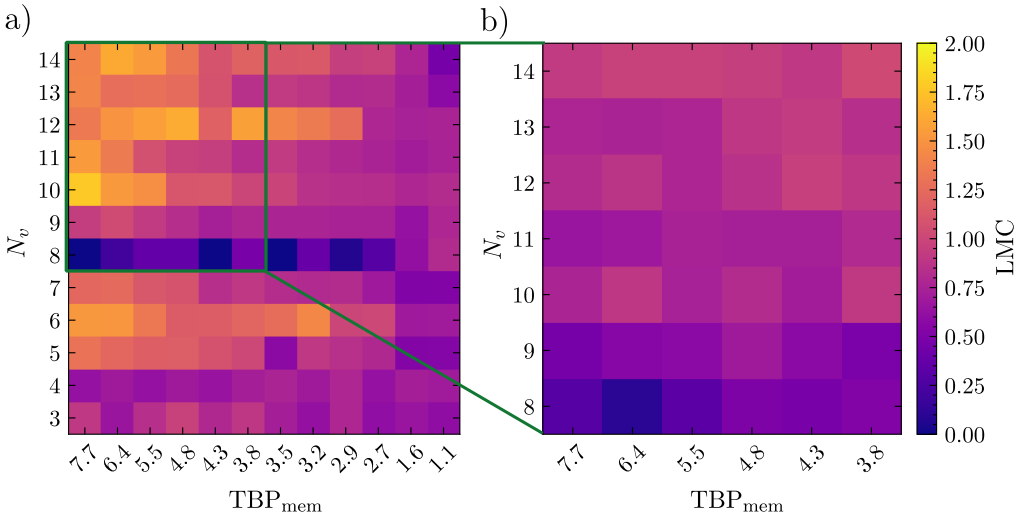


Fig. 5. Linear memory capacity as a function of virtual node number N_v and $\mathrm{TBP}_{\mathrm{mem}}$, for $T_1 = 230\,\mu$s ($t_{\mathrm{pump}} = 10\,\mu$s), $g_1 = 1$. (a) Simulation; (b) experiment. The experimentally scanned region is indicated in green. The maximum experimental MC= 1.0 ± 0.2 is achieved at $N_v = 14$, $\mathrm{TBP}_{\mathrm{mem}} = 3.8$.

4.2. *Kernel rank*

The dependence of the kernel rank on the reservoir parameters N_v and $\mathrm{TBP}_{\mathrm{mem}}$, for a short memory lifetime $T_1 = 230\,\mu$s is shown in Fig. 6. We observe that the KR is primarily determined by the number of virtual nodes; reducing $\mathrm{TBP}_{\mathrm{mem}}$ of the system has a negligible effect on the dimensionality of the projection in both simulation and experiment. Moreover, we find that higher KRs are reached for odd N_v, and conversely for even N_v we achieve high LMCs. We attribute the dependence of KR and LMC on the parity of N_v to the effective network topology imposed by the fixed number of memory cells. For even N_v, not all virtual nodes interact within the memory, and the network separates into subgraphs. Their number is given by $d = \gcd(N_v, 8)$. For odd N_v, $d = 1$, so that over successive cycles every virtual node is written to every memory cell. The formation of subgraphs removes connections between nodes and lowers the projection dimension of the reservoir, which we observe as a reduction in kernel rank. Since the total computational capacity is conserved [45], this loss of nonlinear projection is compensated by an increase in linear memory capacity, accounting for the higher MC at even N_v.

Experimentally, a maximum of KR $= 8.8 \pm 0.4$ is reached, approaching the theoretical upper bound of N_v. However, we do not anticipate the higher experimental KR to be indicative of higher-dimensionality performance compared to the simulated reservoir; rather, we expect that further noise correlations exist, which indicate that a threshold of 2% may be too small and thus overestimate the experimental kernel rank. Further investigation similar to the analysis presented in Ref. [46] is required to confirm this but was deemed beyond the scope of this work.

Optics EXPRESS

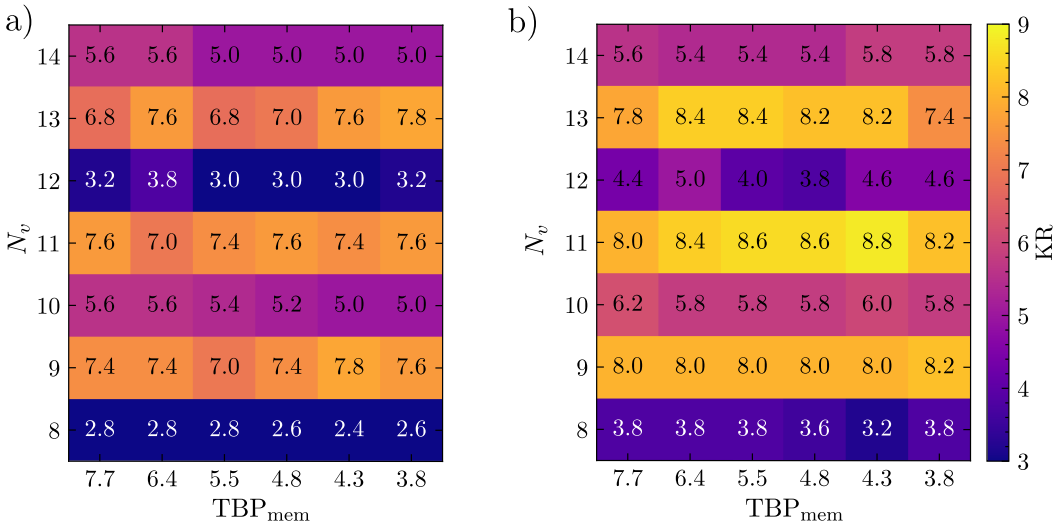


Fig. 6. Kernel rank as a function of N_v and $\mathrm{TBP}_{\mathrm{mem}}$ for $T_1 = 230\,\mu$s. (a) Simulation, $\mathrm{KR}_{\mathrm{sim}}^{\mathrm{max}} = 7.8 \pm 0.4$; (b) experiment, $\mathrm{KR}_{\mathrm{exp}}^{\mathrm{max}} = 8.8 \pm 0.4$.

We investigate the influence that varying the operation regime between linear and nonlinear systems has on the reservoir by varying g_1, revealing that varying the pump power—and thus the degree of saturation nonlinearity—has a weak effect on KR. The results can be seen in Fig. 7, where we observe the KR values ranging only between 6.8 and 8.0 across the full range of g_1. This indicates that the dimensional expansion of the reservoir state is driven primarily by the coupling topology (the choice of N_v relative to N), rather than by the atomic saturation nonlinearity. These results suggest that the saturation absorption profile provides insufficient nonlinear transformation for robust dimensionality expansion independent of topology. An investigation into whether the incorporation of an additional nonlinear element—such as a semiconductor optical amplifier—could further increase the nonlinearity remains an extension to this work.

4.3. *XOR classification*

We first investigate the effect of N_v and $\mathrm{TBP}_{\mathrm{mem}}$ on the performance of the XOR task, as can be seen in Fig. 8 . The test size of 300 samples imposes a quantization floor on the error metric; the minimum resolvable BER is $1/300 \approx 0.003$. This small test set and averaging over 5 masks presents limited statistical validation, however extending the size of the data set was not feasible in this work due to the experimental limitations from device data transfer. Importantly, this means that a BER = 0, as we achieve in the simulation, requires reevaluation with longer test sets to get an accurate measure of the performance. Optimal performance is obtained for even values of N_v, consistent with the correlation between even-N_v topologies and higher LMC observed in Section 4.1. This confirms that memory depth, rather than projection dimensionality, is the limiting factor for the XOR task, which requires an output representation containing the previous time step. The best performing configuration while scanning N_v and $\mathrm{TBP}_{\mathrm{mem}}$ achieved a BER of 0.03(3) at $N_v = 12$, $\mathrm{TBP}_{\mathrm{mem}} = 6.4$.

Figure 9 presents the XOR BER as a function of g_1 at $N_v = 12$. The range of g_1 is reduced to $[0.2, 1]$, because simulations such as those shown in Fig. 7 did not suggest any significant gain in the $g_1 > 1$ range. Here we achieved the minimum experimental BER of 0.02(1) at $g_1 = 0.6$, $\mathrm{TBP}_{\mathrm{mem}} = 6.4$. Crucially, varying g_1 across the full range from linear ($g_1 = 0.2$) to saturated

Optics EXPRESS

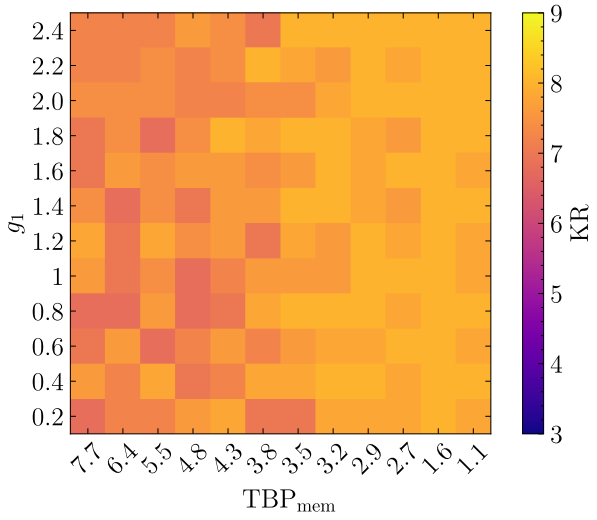


Fig. 7. Simulated kernel rank as a function of input scaling g_1 and TBP$_{\mathrm{mem}}$ for $N_v = 11$, $T_1 = 230\,\mu$s. g_1<1 restricts pumping to the linear regime. The KR range spans 6.8–8.0, with no strong dependence on g_1.

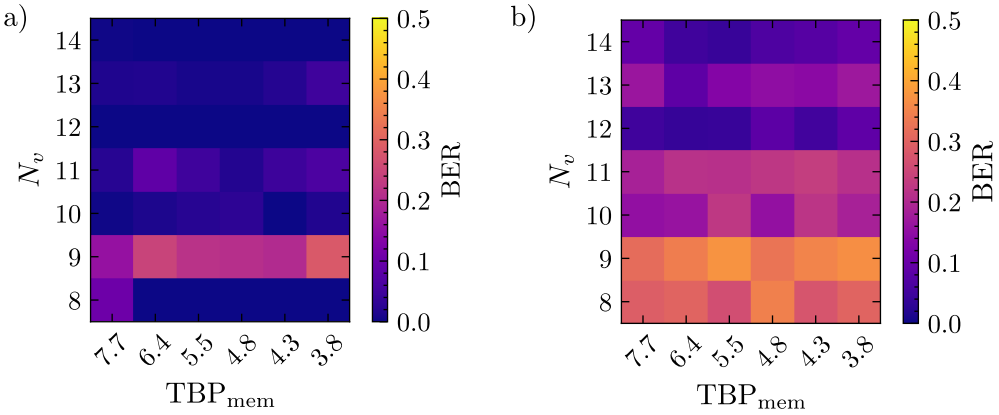


Fig. 8. (a) Simulated and (b) experimental XOR performance (BER) as a function of N_v and TBP$_{\mathrm{mem}}$ for $T_1 = 230\,\mu$s, $g_1 = 1$. In simulation $\mathrm{BER}_{\mathrm{sim}}^{(\mathrm{min})} = 0$, for $N_v \in \{8, 12, 14\}$, for all values of TBP$_{\mathrm{mem}}$, except TBP$_{\mathrm{mem}} = 7.7$. The lowest experimental accuracy $\mathrm{BER}_{\mathrm{exp}}^{(\mathrm{min})} = 0.03(3)$ is obtained at $N_v = 12$, TBP$_{\mathrm{mem}} = 6.4$.

($g_1 = 1.0$) operation produces negligible change in performance, corroborating the finding that the atomic saturation nonlinearity contributes weakly to computation in this system. Simulation predicts zero BER for all values of g_1 when TBP$_{\mathrm{mem}} \geq 4.8$, indicating that the nonlinearity provided by the network topology is sufficient to solve XOR once adequate memory depth is established. The discrepancy between simulated and experimental BER is attributed to measurement noise.

Optics EXPRESS

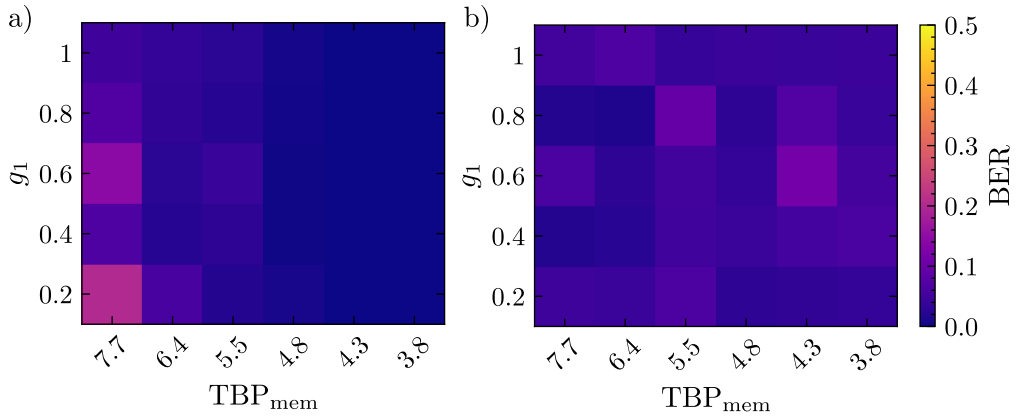


Fig. 9. XOR performance (BER) as a function of input scaling g_1 and $\mathrm{TBP}_{\mathrm{mem}}$ for $N_v = 12$, $T_1 = 230\,\mu$s. (a) Simulation: BER= 0 across all g_1 for $\mathrm{TBP}_{\mathrm{mem}} \geq 4.8$. (b) Experiment: minimum BER= 0.02(1) at $g_1 = 0.6$, $\mathrm{TBP}_{\mathrm{mem}} = 6.4$.

4.4. Discussion

The results reveal a principal limitation of the current implementation: the low memory lifetime is insufficient for deep temporal processing. The diffusion-limited lifetime $\tau_{\mathrm{mem}} \approx 1.3$ ms is not reached in practice; the measured $T_1 = 230\,\mu$s for standard pump parameters limits the logical TBP to approximately 2.6. Increasing N_v to improve network complexity proportionally consumes the available memory budget (Eq. (5)), creating a fundamental complexity-memory trade-off. This trade-off is structurally distinct from delay-line reservoirs, where the virtual node count can be increased by faster sampling without extending the total injection time. Moreover, while the system's short temporal dynamics and nonlinear projection suggest potential for use in low-recurrence tasks which are often used to benchmark extreme learning machines, the single-input reservoir architecture prevents the use of multi-feature inputs and therefore limits its applicability in this case.

Several strategies may mitigate the limited relaxation lifetime of the system. These include implementing active re-pumping to refresh atomic polarization or employing electronic pre-processing to augment the input with delayed copies, thereby shifting the memory requirements to the pre-processing layer [41,47]. Moreover, the ultimate limit on rail addressing is set by the atomic excited-state lifetime. Future work could therefore explore high-speed addressing via AOMs or EODs together with fast coherent pulses to accelerate read/write cycles [30], or parallelized operations updating multiple virtual nodes within a single time step. Memory access times of ~ 100 ns, achievable with faster deflection hardware, would place the incoherent architecture within the $\sim$ 10–1000 ns range of access times typically seen in delay-line memories [48–50]. In addition, we have demonstrated in this work storage times of order $\sim 100\,\mu$s, two orders of magnitude beyond the $\sim$ 100–1000 ns typically available from delay-line systems. Alternatively, a transition to telecom-compatible hardware, such as erbium-doped fiber memories, offers a promising approach, given the demonstrated capacity for highly multiplexed operation [51,52].

Despite the memory lifetime constraints, this work establishes that warm Cs atomic vapor is a viable substrate for free-space optical RAM and demonstrates its integration into a reservoir computer. The system successfully performs XOR with BER = 0.02(1), and the ORAM supports accumulative in-memory writing in the linear regime.

Optics EXPRESS

5. Conclusion

We have demonstrated an optically writable spatially multiplexed ORAM based on warm Cs atomic vapor and its application as a physical substrate for a reservoir computer. The memory operates through optical pumping of the D_1 hyperfine transition, with eight addressable spatial rails defined by acousto-optic deflection. Characterization yielded a $1/e$ lifetime of $T_1 = 230\,\mu$s and a maximum logical time-bandwidth product of $\mathrm{TBP}_{\mathrm{logic}} = 2.6$.

The reservoir benchmarking showed a maximum linear memory capacity of LMC$= 1.0 \pm 0.2$ and a kernel rank of KR$= 8.8 \pm 0.4$, both primarily governed by the network topology rather than the nonlinearity arising from atomic saturation. The system successfully solved the XOR classification task with a minimum BER of 0.02 ± 0.01. Performance is fundamentally limited by the interplay between memory lifetime and the temporal multiplexing overhead, which restricts the logical TBP and confines the reservoir to a low recurrence operating regime for tasks requiring multi-step memory. Future research will focus on optimizing read/write protocols for storage rails and evaluating the suitability of high-speed memory architectures, such as those based on electromagnetically induced transparency (EIT), to profit from the fast modulation possible in atomic systems.

Funding. Deutsche Forschungsgemeinschaft (445183921); Deutsches Zentrum für Luft- und Raumfahrt e. V. (DLR) (50RP2410).

Acknowledgments. E.R. acknowledges funding through the Helmholtz Einstein International Berlin Research School in Data Science (HEIBRiDS).

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper may be obtained from the authors upon reasonable request.

References

1. A. M. Turing, "On Computable Numbers, with an Application to the Entscheidungsproblem," Proc. London Math. Soc. **s2-42**(1), 230–265 (1937).
2. J. Mumbru, G. Zhou, X. An, *et al.*, *Optical memory for computing and information processing*, (Denver, CO, 1999), pp. 0, 14–24.
3. R. Chen, J. Wang, H. Wu, *et al.*, "Encoding and decoding of multidimensional optical field modulation in holographic data storage," Optica **13**(4), 591 (2026).
4. N. Cheriere, J. Chu, G. Brennan, *et al.*, "Holographic Storage for the Cloud: advances and challenges," ACM Trans. Storage **21**(1), 1–31 (2025).
5. N. P. Jouppi, C. Young, N. Patil, *et al.*, "In-Datacenter Performance Analysis of a Tensor Processing Unit," in *Proceedings of the 44th Annual International Symposium on Computer Architecture*, (ACM, Toronto ON Canada, 2017), pp. 1–12.
6. A. Sebastian, M. Le Gallo, R. Khaddam-Aljameh, *et al.*, "Memory devices and applications for in-memory computing," Nat. Nanotechnol. **15**(7), 529–544 (2020).
7. C. Ríos, N. Youngblood, Z. Cheng, *et al.*, "In-memory computing on a photonic platform," Sci. Adv. **5**(2), eaau5759 (2019).
8. K. P. Kalinin, J. Gladrow, J. Chu, *et al.*, "Analog optical computer for AI inference and combinatorial optimization," Nature **645**(8080), 354–361 (2025).
9. P. L. McMahon, "The physics of optical computing," Nat. Rev. Phys. **5**(12), 717–734 (2023).
10. Y. Liang, J. Wang, K. Xue, *et al.*, "High-clockrate free-space optical in-memory computing," Light: Sci. Appl. **15**(1), 115 (2026).
11. X. Xu, M. Tan, B. Corcoran, *et al.*, "11 TOPS photonic convolutional accelerator for optical neural networks," Nature **589**(7840), 44–51 (2021).
12. M. Aifer, Z. Belateche, S. Bramhavar, *et al.*, "Solving the compute crisis with physics-based ASICs," arXiv (2025).
13. T. Alexoudi, G. T. Kanellos, and N. Pleros, "Optical RAM and integrated optical memories: a survey," Light: Sci. Appl. **9**(1), 91 (2020).
14. N. L. Kazanskiy, M. A. Butt, and S. N. Khonina, "Optical Computing: Status and Perspectives," Nanomaterials **12**(13), 2171 (2022).
15. H. Wang, J. Hu, Y. Baek, *et al.*, "Optical next generation reservoir computing," Light: Sci. Appl. **14**(1), 245 (2025).
16. J. R. Rausell-Campo, A. Hurtado, D. Pérez-López, *et al.*, "Programmable Photonic Extreme Learning Machines," Laser Photonics Rev. **19**(9), 2400870 (2025).
17. M. Abdalla, G. Van Der Sande, A. Argyris, *et al.*, "Photonic reservoir computing: A thematic review," Journal of Physics: Photonics **8**, 022003 (2025).

Optics EXPRESS

18. J. Mulet and S. Balle, "Spatio-temporal modeling of the optical properties of VCSELs in the presence of polarization effects," IEEE J. Quantum Electron. **38**(3), 291–305 (2002).
19. J. Dong, M. Rafayelyan, F. Krzakala, *et al.*, "Optical Reservoir Computing Using Multiple Light Scattering for Chaotic Systems Prediction," IEEE J. Sel. Top. Quantum Electron. **26**(1), 1–12 (2020).
20. M. Rafayelyan, J. Dong, Y. Tan, *et al.*, "Large-Scale Optical Reservoir Computing for Spatiotemporal Chaotic Systems Prediction," Phys. Rev. X **10**(4), 041037 (2020).
21. V. V. Semenov, X. Porte, I. Abdulhalim, *et al.*, "Two-color optically addressed spatial light modulator as a generic spatiotemporal system," Chaos: An Interdiscip. J. Nonlinear Sci. **31**(12), 121104 (2021).
22. E. Gemo, S. G.-C. Carrillo, C. R. De Galarreta, *et al.*, "Plasmonically-enhanced all-optical integrated phase-change memory," Opt. Express **27**(17), 24724 (2019).
23. X. Li, N. Youngblood, C. Ríos, *et al.*, "Fast and reliable storage using a 5 bit, nonvolatile photonic memory cell," Optica **6**(1), 1 (2019).
24. J. Song, J. Pady, E. Gemo, *et al.*, "Fabrication-friendly all-optical plasmonically-enhanced integrated phase-change photonic memory device," Opt. Express **33**(25), 52965 (2025).
25. F. Ashtiani, "Programmable photonic latch memory," Opt. Express **33**(2), 3501 (2025).
26. M. G. Kirzhner, M. Klebanov, V. Lyubin, *et al.*, "Liquid crystal high-resolution optically addressed spatial light modulator using a nanodimensional chalcogenide photosensor," Opt. Lett. **39**(7), 2048 (2014).
27. P. K. Shrestha, Y. T. Chun, and D. Chu, "A high-resolution optically addressed spatial light modulator based on ZnO nanoparticles," Light: Sci. Appl. **4**(3), e259 (2015).
28. S. Gong, M. Ren, W. Wu, *et al.*, "Optically addressed spatial light modulator based on nonlinear metasurface," Photonics Res. **9**(4), 610 (2021).
29. X. Fan, W. Xiong, K. Xu, *et al.*, "Spatial light modulator via optically addressed metasurface," Nat. Nanotechnol. **21**(4), 561–570 (2026).
30. L. Messner, E. Robertson, L. Esguerra, *et al.*, "Multiplexed random-access optical memory in warm cesium vapor," Opt. Express **31**(6), 10150–10158 (2023).
31. W. Happer, "Optical Pumping," Rev. Mod. Phys. **44**(2), 169–249 (1972).
32. B. Wiegand, B. Leykauf, R. Jördens, *et al.*, "Linien: A versatile, user-friendly, open-source FPGA-based tool for frequency stabilization and spectroscopy parameter optimization," Rev. Sci. Instrum. **93**(6), 063001 (2022).
33. R. Mottola, G. Buser, and P. Treutlein, "Optical Memory in a Microfabricated Rubidium Vapor Cell," Phys. Rev. Lett. **131**(26), 260801 (2023).
34. D. Nodop, T. Kügler, J. Rücker, *et al.*, "Atomic vapor cells fabricated by femtosecond laser welding of standard-optical-quality glass," Opt. Express **34**(13), 24924 (2026).
35. Y. Li, M. R. Slot, M. T. Hummon, *et al.*, "Wafer-scale fabrication of temperature-compensated alkali vapor cells," Appl. Phys. Lett. **126**(22), 224001 (2025).
36. H. Jaeger, "The "echo state" approach to analysing and training recurrent neural networks," Tech. rep., German National Research Center for Information Technology GMD (2001).
37. L. Larger, M. C. Soriano, D. Brunner, *et al.*, "Photonic information processing beyond Turing: an optoelectronic implementation of reservoir computing," Opt. Express **20**(3), 3241 (2012).
38. L. Appeltant, M. Soriano, G. Van Der Sande, *et al.*, "Information processing using a single dynamical node as complex system," Nat. Commun. **2**(1), 468 (2011).
39. T. Hülser, F. Köster, L. Jaurigue, *et al.*, "Role of delay-times in delay-based photonic reservoir computing [Invited]," Opt. Mater. Express **12**(3), 1214 (2022).
40. S. Ortín, L. Pesquera, G. Van Der Sande, *et al.*, "5. Time delay systems for reservoir computing," in *Photonic Reservoir Computing*, D. Brunner, M. C. Soriano, and G. Van Der Sande, eds. (De Gruyter, 2019), pp. 117–152.
41. L. Jaurigue, E. Robertson, J. Wolters, *et al.*, "Reservoir Computing with Delayed Input for Fast and Easy Optimisation," Entropy **23**(12), 1560 (2021).
42. C. Wringe, M. Trefzer, and S. Stepney, "Reservoir computing benchmarks: a tutorial review and critique," Int. J. Parallel, Emergent Distributed Syst. **40**(4), 313–351 (2025).
43. F. Köster, D. Ehlert, and K. Lüdge, "Limitations of the Recall Capabilities in Delay-Based Reservoir Computing Systems," Cogn. Comput. **15**(5), 1419–1426 (2020).
44. A. Röhm, L. Jaurigue, and K. Lüdge, "Reservoir Computing Using Laser Networks," IEEE J. Sel. Top. Quantum Electron. **26**(1), 1–8 (2020).
45. J. Dambre, D. Verstraeten, B. Schrauwen, *et al.*, "Information Processing Capacity of Dynamical Systems," Sci. Rep. **2**(1), 514 (2012).
46. A. Skalli, X. Porte, N. Haghighi, *et al.*, "Computational metrics and parameters of an injection-locked large area semiconductor laser for neural network computing [Invited]," Opt. Mater. Express **12**(7), 2793 (2022).
47. E. Picco, L. Jaurigue, K. Lüdge, *et al.*, "Efficient optimisation of physical reservoir computers using only a delayed input," Commun. Eng. **4**(1), 3 (2025).
48. D. Brunner, M. C. Soriano, C. R. Mirasso, *et al.*, "Parallel photonic information processing at gigabyte per second data rates using transient states," Nat. Commun. **4**(1), 1364 (2013).
49. G. Donati, A. Argyris, M. Mancinelli, *et al.*, "Time delay reservoir computing with a silicon microring resonator and a fiber-based optical feedback loop," Opt. Express **32**(8), 13419 (2024).

Optics EXPRESS

50. Q. Vinckier, F. Duport, A. Smerieri, *et al.*, "High-performance photonic reservoir computer based on a coherently driven passive cavity," Optica **2**(5), 438 (2015).
51. N. G. Kamel, A. Mansourzadeh, U. Gautam, *et al.*, "Erbium-Doped Fibre Quantum Memory for Chip-Integrated Quantum-Dot Single Photons at 980 nm," arXiv (2025).
52. E. Saglamyurek, J. Jin, V. B. Verma, *et al.*, "Quantum storage of entangled telecom-wavelength photons in an erbium-doped optical fibre," Nat. Photonics **9**(2), 83–87 (2015).