



Numerical demonstration of neuromorphic computing with photonic crystal cavities

FLORIS LAPORTE,^{1,*} ANDREW KATUMBA,¹ JONI DAMBRE,² AND PETER BIENSTMAN¹

¹Photonics Research Group, UGent-imec, Technologiepark-Zwijnaarde 15, 9052 Ghent, Belgium

²IDLab, UGent-imec, Technologiepark-Zwijnaarde 15, 9052 Ghent, Belgium

*floris.laporte@ugent.be

Abstract: We propose a new design for a passive photonic reservoir computer on a silicon photonics chip which can be used in the context of optical communication applications, and study it through detailed numerical simulations. The design consists of a photonic crystal cavity with a quarter-stadium shape, which is known to foster interesting mixing dynamics. These mixing properties turn out to be very useful for memory-dependent optical signal processing tasks, such as header recognition. The proposed, ultra-compact photonic crystal cavity exhibits a memory of up to 6 bits, while simultaneously accepting bitrates in a wide region of operation. Moreover, because of the inherent low losses in a high-Q photonic crystal cavity, the proposed design is very power efficient.

© 2018 Optical Society of America under the terms of the [OSA Open Access Publishing Agreement](#)

OCIS codes: (200.3050) Information processing; (200.4560) Optical data processing; (200.4700) Optical neural systems; (200.4740) Optical processing; (060.4510) Optical communications; (050.5298) Photonic crystals;

References and links

1. H. Jaeger, "The echo state approach to analyzing and training recurrent neural networks," Bonn, Germany: German National Research Center for Information Technology GMD Technical Report **148**(34), 13 (2001).
2. W. Maass, T. Natschläger, and H. Markram, "Real-time computing without stable states: A new framework for neural computation based on perturbations," *Neural comput.* **14**(11), 2531–2560 (2002).
3. D. Verstraeten, B. Schrauwen, M. d'Haene, and D. Stroobandt, "An experimental unification of reservoir computing methods," *Neural networks* **20**(3), 391–403 (2007).
4. Q. Vinckier, F. Duport, A. Smerieri, K. Vandoorne, P. Bienstman, M. Haelterman, and S. Massar, "High-performance photonic reservoir computer based on a coherently driven passive cavity," *Optica* **2**(5), 438–446 (2015).
5. L. Appeltant, M. C. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, "Information processing using a single dynamical node as complex system," *Nat. Commun.* **2**, 468 (2011).
6. Y. Paquot, F. Duport, A. Smerieri, J. Dambre, B. Schrauwen, M. Haelterman, and S. Massar, "Optoelectronic reservoir computing," *Sci. Rep.* **2**, 287 (2012).
7. L. Larger, M. C. Soriano, D. Brunner, L. Appeltant, J. M. Gutiérrez, L. Pesquera, C. R. Mirasso, and I. Fischer, "Photonic information processing beyond turing: an optoelectronic implementation of reservoir computing," *Opt. Express* **20**(3), 3241–3249 (2012).
8. D. Brunner, M. C. Soriano, C. R. Mirasso, and I. Fischer, "Parallel photonic information processing at gigabyte per second data rates using transient states," *Nat. Commun.* **4**, 1364 (2013).
9. L. Larger, A. Baylón-Fuentes, R. Martinenghi, V. S. Udaltsov, Y. K. Chembo, and M. Jacquot, "High-speed photonic reservoir computing using a time-delay-based architecture: million words per second classification," *Phys. Rev. X* **7**, 011015 (2017).
10. K. Vandoorne, W. Dierckx, B. Schrauwen, D. Verstraeten, R. Baets, P. Bienstman, and J. V. Campenhout, "Toward optical signal processing using photonic reservoir computing," *Opt. Express* **16**(15), 11182–11192 (2008).
11. K. Vandoorne, P. Mechet, T. Van Vaerenbergh, M. Fiers, G. Morthier, D. Verstraeten, B. Schrauwen, J. Dambre, and P. Bienstman, "Experimental demonstration of reservoir computing on a silicon photonics chip," *Nat. Commun.* **5**, 3541 (2014).
12. H.-J. Stöckmann and J. Stein, "Quantum chaos in billiards studied by microwave absorption," *Phys. Rev. Lett.* **64**(19), 2215–2218 (1990).
13. M. Sieber, U. Smilansky, S. Creagh, and R. Littlejohn, "Non-generic spectral statistics in the quantized stadium billiard," *J. Phys. A: Mathematical and General* **26**(22), 6217 (1993).
14. C. Liu, R. E. Van Der Wel, N. Rotenberg, L. Kuipers, T. F. Krauss, A. Di Falco, and A. Fratalocchi, "Triggering extreme events at the nanoscale in photonic seas," *Nat. Phys.* **11**(4), 358–363 (2015).

15. A. J. Izenman, "Linear discriminant analysis," in *Modern multivariate statistical techniques* (Springer, 2013), pp. 237–280.
16. M. Jeruchim, "Techniques for estimating the bit error rate in the simulation of digital communication systems," J-SAC 2(1), 153–170 (1984).

1. Introduction

Reservoir Computing (RC) is a well-established machine-learning concept, first proposed in the early 2000s [1, 2] as a brain-inspired computing mechanism to process temporal signals in real time. In RC, a highly dynamical and nonlinear system is used as a signal-mixing *reservoir*, which essentially produces a number of non-linearly mixed versions of the input signal. Those resulting *reservoir states* are then *read out* by performing an application-dependent linear combination [3]. This linear combination is then used to perform a certain task, such as header recognition or bitwise XOR.

For a passive reservoir operating in discrete time and for a certain sequence of inputs $\mathbf{u}_n$, the resulting reservoir states $\mathbf{x}_n$ and the corresponding readout values $\mathbf{y}_n$ are given by the following formulas [4]:

$$\mathbf{x}_n = W_{\text{in}}\mathbf{u}_{n-1} + W_{\text{res}}\mathbf{x}_{n-1} \quad (1)$$

$$\mathbf{y}_n = W_{\text{out}}f(\mathbf{x}_n) \quad (2)$$

Here, W_{in} and W_{res} perform linear combinations on the inputs and the states respectively, while f is a nonlinear function. W_{out} performs the readout operation by doing an application-dependent linear combination on the reservoir states [3]. It is clear that, to have a reliable computation, the reservoir needs to possess a form of memory. This so-called *fading memory* [1, 2] needs to retain enough information of the history of the input signal to perform meaningful mixing, while not having too much memory such that it becomes very sensitive to initial conditions.

Backed by the promise of the ultra high-speed and high-bandwidth signal propagation of photonics, the reservoir computer has already found its way into several optical hardware implementations. These *photonic reservoir computers* roughly split into two kinds: the single-node reservoir, based on delayed feedback [4–9] and the photonic reservoir on chip [10, 11].

One of the advantages of these photonic reservoir computers is the possibility for removing the nonlinearity inside the reservoir in favor of a nonlinear measurement operation at the readout [4, 11]. These kind of reservoir computers are called *passive* photonic reservoir computers, which exhibit ultra high speed of operation because of the absence of internal nonlinearities. In fact, the operation speed is only limited by the operation speed of the final photodetector. Removing the nonlinearities inside the reservoir is of course only possible because we are working with coherent light, which possesses an amplitude and a phase, resulting in a nonlinearly mixed magnitude in the detector.

While the on-chip reservoir is already a lot more compact than the delayed feedback reservoir, it still suffers from some serious drawbacks, such as a very limited interconnection topology, a low density of nodes per chip surface area and high losses because of the many 3 dB-splitters. To address these issues, we propose a new passive design on a silicon photonics chip, consisting of a photonic crystal cavity with a quarter-stadium shape that seeks to improve on the reservoir-on-chip technology. This quarter-stadium resonator is known to foster interesting mixing dynamics [12–14], while the inherent properties of photonic crystals result in low losses, which makes this resonator a very good candidate for a reservoir.

Our extremely small design results in a reservoir with dimensions smaller than 0.01 mm^2 , which is at least two orders of magnitude smaller than the previous reservoir-on-chip design [11]. As we will see, this photonic crystal design promises very low loss combined with excellent performance on several benchmark telecom tasks proposed in [11], such as the highly nonlinear XOR task and header recognition tasks, while still accepting bitrates in a wide region of operation.

2. Design




Fig. 1. Waveguides connected to the photonic crystal cavity. (a) Snapshot of the field profile in the cavity. The mixing of the signal can clearly be witnessed by inspecting the field profiles. (b) Close up of the entrance waveguide. Some scattering losses can be observed at the waveguide-W1 transition.

Our system consists of an on-chip photonic crystal cavity in the shape of a quarter-stadium resonator, which is known to foster interesting mixing of the fields in an almost chaotic manner [12–14]. The photonic crystal cavity was designed for the 220 nm silicon photonics platform, consisting of holes etched from a 220 nm silicon slab with radius $r = 0.37a$, for which $a = 420$ nm. The light is sent through one of the seven standard 450 nm waveguides, which are connected to W1-defects in the wall of the photonic crystal cavity. The light inside the cavity subsequently leaks out of the cavity via all of those defects. The six other defects are used for readout. It is clear that this cavity is a great candidate for a reservoir, because alongside its *mixing property*, it also possesses a *fading memory*: the signal is bound to remain in this cavity for a certain amount of time directly proportional to the Q-factor and the dimensions of this cavity (Fig. 2). Moreover, light is trapped inside the cavity and can only leak out using the wall defects, where they contribute to the useful output signal. This results in low overall losses in the reservoir.

The size of the cavity will obviously have a profound effect on the region of operation of the reservoir in terms of the bitrate. The dimensions of the cavity with 7 W1-waveguides (Fig. 1(a)) are $30\ \mu\text{m} \times 60\ \mu\text{m}$ and were chosen to exhibit a memory of up to 5 bits at the highest attainable bitrate of our detector (around 25 Gbps). For these dimensions, the cavity exhibits a Q-factor of approximately 16400, which corresponds to a pulse half life of about 20 ps. This means that for the requested binary input of 25 Gbps, about 4 to 5 bits will be able to get mixed with each other before the intensity of the first bit drops below the noise level. As our simulations will confirm later in this paper, this results in sufficient memory for the XOR task and up to 5-bit header recognition at 25 Gbps. However this frequency cutoff is not steep, and the detector remains useful for bitrates up to about 50 Gbps, as can be seen in later results.




Fig. 2. Block diagram of the simulation.

3. Simulations

3.1. Detecting the bit stream response

To obtain the response of the photonic crystal reservoir to an arbitrary bit stream, the most time-consuming step of the calculation is simulating the propagation of the light through the cavity. However, we do not simulate the propagation of the complete bit stream, as this would result in enormous calculation times. Instead, we use an alternative approach illustrated in Fig. 2.

Suppose we have a bit stream $\mathbf{s}$ consisting of the bits $b_1, b_2, \dots, b_N \in \{0, 1\}$, for which the nonzero bits are specified by a pulse $\mathbf{u}(t)$:

$$\mathbf{u}(t) = \begin{pmatrix} \mathbf{E}_0(t) \\ \mathbf{H}_0(t) \end{pmatrix} \quad \text{with } \mathbf{u}(t) = \mathbf{0} \text{ if } t < 0 \text{ or } t > T, \quad (3)$$

with T the bit period of the signal. Then the value of the bit stream at time t is given by

$$\mathbf{x}(t) = \begin{pmatrix} \mathbf{E}_{\text{in}}(t) \\ \mathbf{H}_{\text{in}}(t) \end{pmatrix} = \sum_{n=1}^N b_n \mathbf{u}(t - nT) = b_k \mathbf{u}(t - kT) \quad \text{with } k = \left\lceil \frac{t}{T} \right\rceil \quad (4)$$

Each of the exit waveguides i will have an exponentially decaying response $\mathbf{U}_i(t)$ to the single bit pulse $\mathbf{u}(t)$, with $\mathbf{U}_i(t) > 0$ for all $t > 0$. This means that the response $\mathbf{X}_i(t)$ to the total bit stream at waveguide i can be described by

$$\mathbf{X}_i(t) = \begin{pmatrix} \mathbf{E}_{\text{out}}^i(t) \\ \mathbf{H}_{\text{out}}^i(t) \end{pmatrix} = \sum_{n=1}^N b_n \mathbf{U}_i(t - nT). \quad (5)$$

Note that in this case, the responses can not be decoupled anymore since $\mathbf{U}_i(t) \neq 0$ for $t > T$. This means that the state of the reservoir depends linearly on the previous input values, just like one would require from the passive reservoir described in (1).

In practice, the responses $P_i(t)$ at each of the waveguides to a single pulse are recorded from an FDTD simulation, performed by the *Lumerical FDTD* software. Then the response of a complete bit stream (typically a PRBS signal of 10^5 bits) is calculated by coherently adding together the individual bit responses for each channel in the above described way. Note that this described procedure is just a "bit-level" version of the impulse response method, where the response of an arbitrary system is found by convolving the function with the response of an ultrashort impulse. Here, we chose to work with the "bit-level" response instead of the true impulse response because of numerical rounding errors.

Finally, The photodetector performs the nonlinear operation described in (2) by detecting the light. This detector is simulated by a detector model with similar parameters as the detector in

our labs, which has a load resistance $R_L = 1 \text{ k}\Omega$, a bandwidth $f_c = 25 \text{ GHz}$ and a responsivity $\eta = 0.5 \text{ A/W}$. The detector noise introduced by the amplification of the photogenerated current is described by white thermal noise I_{tn} , modelled as a Nyquist process and shot noise I_{sn} , modelled as a Poisson process:

$$I_n = \sqrt{I_{\text{tn}}^2 + I_{\text{sn}}^2} = \sqrt{\left(\frac{4kTf_c}{R_L}\right)^2 + 2qIf_c}. \quad (6)$$




Fig. 3. (a) Waveforms detected at two of the exit waveguides as the result of a certain 50 Gbps bit sequence input. The outputs get sampled once per bit period. (b) After the readout, the prediction approximates the desired XOR target. The prediction and the target were aligned by shifting the prediction backwards in time according to the optimal latency of 0.8 bit periods.

3.2. Readout

In the readout, the reservoir output streams are sampled a fixed number of times per bit period, as is shown in Fig. 3(a). After this, a linear combination of the sampled values is made according to weights that are specifically trained for the intended application. Increasing the number of sampling points generally improves the performance of the reservoir. However, since we are already working at very high bitrates, we usually choose to sample only once per bit.

To obtain the weight matrix W_{out} for the linear combination acting on the output states of the reservoir, two different kinds of training algorithms are used: *ridge regression* (linear regression with a regularization parameter) for the binary classification tasks such as the XOR task, and *linear discriminant analysis* (LDA) [15] for multi-class classification tasks such as header recognition.

Linear Discriminant Analysis is a well-known linear machine learning technique, seeking to find the best linear combination of features (in our case output signals) that best separates the

different classes of the problem. It is often used for dimensionality reduction of the problem by discarding the least important singular values of the weight matrix.

The reservoir can be trained to reproduce the desired target signal with a certain delay, which we call the *latency* of the reservoir. We usually express the latency as a multiple of the bit period, i.e. the number (or fraction) of bit periods you have to wait after the last relevant bit has completely entered the cavity before you can reproduce the target signal. There usually exists only a narrow range of latencies for which the target can be approximated without error. The approximation for an XOR-target at optimal latency is shown in Fig. 3(b). Note that we aligned the target and the prediction by shifting the prediction backwards in time according to the latency.

3.3. Power budget of the reservoir

The simulated total energy measured at the exits of the cavity is about 83 % of the total energy inserted, as can be seen in Fig. 4. Looking deeper into the source of the losses, we find that there is a slight mode mismatch between the access waveguides and the W1 photonic crystal waveguides. When coupling out the light from the cavity, this causes scattering out of the exit waveguides, resulting in lost power. Further optimization to improve this transmission can still be performed, but is beyond the scope of this paper.




Fig. 4. The accumulated power exiting the cavity via the output waveguides adds up to about 83 % of the total energy introduced in the system. This corresponds to a 0.8 dB loss.

3.4. Boolean tasks

In our simulations, a PRBS of 10^5 bits with an input power of 1 mW is sent through one of the photonic crystal W1-waveguides (the top waveguide on the left in Fig. 1(a)). The responses of the six other waveguides are then recorded. Finally, on the recorded output stream, the readout weights are trained to follow the intended target function with an as low as possible mean squared error, as is shown in Fig. 3(b). After performing a threshold, the bit error rate (BER) is calculated. Since we use 10^5 bits in our simulations, the general guideline is to crop the BER at 10^{-3} , i.e. 2 orders of magnitude higher than the lowest BER one can find in the simulation [16].

As one of the binary target functions, the XOR of two consecutive bits was chosen, which is known in machine learning to be a hard, nonlinear task due to the fact that the output cannot be found by just performing a linear classification algorithm such as linear regression on the inputs. However, as can be seen in Fig. 3(a) and Fig. 5(a), sending the input first through the photonic crystal cavity and training a readout on the nonlinearly detected outputs of the reservoir, clearly helps to perform the XOR.

During the optimization of the readout, the latency is swept. Fig. 5 illustrates again that the nonlinear XOR is a quite difficult task, as the maximal latency with errorless performance is about half of the maximal latency of the AND task.




Fig. 5. The XOR and AND of two neighboring bits at 50 Gbps can easily be performed by the reservoir. The latency represents the time shift with respect to the input, after which the output should give the desired value. The correct result of $b_n \text{ XOR } b_{n-1}$ can be retained until a new bit is sent in, while the AND of two bits can be retained until 2 new bits are sent in. Since the BER is cropped at 10^{-3} , we often look how well the output *continuously* approximates the target function by looking at the Mean Squared Error (MSE).




Fig. 6. By performing the performance vs latency sweep for every bitrate and looking at the minimal error for each sweep, the operating range of the reservoir can be determined. (a) The optimal bitrate - looking at the MSE - for the XOR task lies at 50 Gbps. However, there is a full band of frequencies the reservoir can work with between 25 Gbps and 67 Gbps. (b) The fact that the linear AND task is easier is again reflected in the the region of operation for the AND, which starts at lower bitrates.

By performing the performance vs latency sweep for every bitrate and looking at the minimal error for each sweep, we can infer the operation range of the reservoir. Looking at Fig. 6, we see that this is where this reservoir really excels: by just adjusting the readout weights, this simple cavity can perform the XOR task reliably between 25 Gbps and 67 Gbps. This region of operation is much wider than previously reported values for the XOR task [11].

At higher bitrates, both the AND and the XOR stop working. This can be explained by the fact that the reservoir remembers *too many* previous bits, which increases the signal to noise ratio on the most recent bits, which are relevant for the operation.

The interesting result here is that we succeeded to compute a highly nonlinear function such as XOR by using a completely passive device. This is of course only possible because of the

nonlinearity of the photodetector, which takes the magnitude of the complex-valued field at the exits of the reservoir.

3.5. Header recognition

For applications in telecom, recognizing headers in the bit stream is often a more useful task. Interestingly, the exact same simple cavity design can perform this task as well.

Concretely, all different headers were searched for simultaneously in a completely random bit stream. For each bit in the bit stream, a class label was given corresponding to the header of length L made by the current bit and the $L - 1$ previous bits. Then, LDA was used to find a different weight matrix for each of the different classes. At each different bitrate, the error rate for the worst performing header (class) is reported.




Fig. 7. Error Rate (ER) for the worst performing header at each latency. The reservoir can distinguish headers of up to $L=6$ bits without error at the optimal bitrate of 50 Gbps. To reduce simulation times, the sweep over the latencies was stopped when the ER became higher than 10^{-1} .

As can be seen in Fig. 7, errorless recognition of up to 6-bit headers is possible at 50 Gbps. Apart from providing us with readout matrix, LDA also allows us to project the 6-dimensional (number of exit waveguides) reservoir output state onto a lower dimensional state, in which we can see the separation of the headers visually [Fig. 8].




Fig. 8. We can visualize the separation of three bit headers by projecting on the two primary LDA axes. We see nice separation for all different headers, while similar headers are located closer together.

Finally, in terms of bitrate, our simulations show that longer headers work better at lower bitrates, as can be seen in Fig. 9. This is unsurprising, as for longer headers, one needs to keep more bits in memory. Therefore, the bitrate needs to be higher to accommodate this.




Fig. 9. By sweeping over the bitrate to find the operation range, we find that the reservoir can distinguish headers up to a header length of $L = 6$ bits without error at a bitrate of up to 100Gbps.

3.6. Q-factor and time scale of the reservoir

In the previous simulations, we have always chosen for 7 connected waveguides. This choice was rather arbitrary and one could argue that changing the number of connected waveguides will have a profound effect on the quality of the reservoir, as decreasing the number of exit waveguides will inevitably increase the Q-factor and thus the memory of the reservoir.

This form of optimization is far from trivial, as removing exits will likely also decrease the complexity of the tasks the reservoir can solve. To quantify this effect, we look at the performance of the reservoir on the XOR task for a range of number of exits.

When the light source is turned off, the power in the cavity with 7 connected waveguides decays exponentially with a slope $m = -0.070 \text{ ns}^{-1}$, as can be seen in Fig. 10. This yields for the Q-factor at $\lambda = 1550 \text{ nm}$:

$$Q = -\frac{2\pi c}{\lambda m} = 16400. \quad (7)$$

A more useful value for the time scale of the reservoir is however the *half-life* $T_{1/2}$ of the pulse, as it provides a more tangible metric for the memory of the reservoir.

$$T_{1/2} = -\frac{\log(2)}{m} = 10 \text{ ps}. \quad (8)$$




Fig. 10. Decay of the power inside the cavity. The envelope decays with a half-life of 10 ps, yielding a Q factor of 16400.




Fig. 11. The Q-factor decays harmonically with the number of connected waveguides

Looking at the Q-factor and half life for several variations of the cavity with less connected waveguides, we see in Fig. 11, as we would expect, that the Q factor decays harmonically with the number of connected waveguides. We also clearly see that the threshold for the number of connected waveguides for the XOR task lies at 6 waveguides (1 input and 5 outputs).

Although certainly possible, increasing the number of connected waveguides even more will not yield better results for the benchmark tasks presented here, however it might prove to be a useful tactic for tackling some more challenging problems in the future. One should keep into account, however, that introducing extra waveguides will also have an adverse effect on the Q-factor of the cavity and thus also on the memory of the reservoir.

4. Conclusions

We have proposed a completely new photonic reservoir computer on a silicon photonic chip. The proposed tiny photonic crystal cavities seem to be very good candidates for optical neuromorphic computing. They show excellent performance on Boolean tasks such as the nonlinear XOR and the AND task. Additionally, the exact same reservoir can also perform header recognition for up to 6-bit headers.

Our simulations also show that this reservoir accepts a very wide range of bitrates, with good performance between 25 Gbps to 67 Gbps for the XOR-task, while the header recognition task can be performed up to 100 Gbps without having to change anything to the design. For both applications, the region of operation can probably still be extended to lower bitrates by increasing the size of the cavity. Conversely, decreasing the size of the cavity should yield better performance at higher bitrates, allowing to trivially upscale the working range of the reservoir to bitrates that are presently unattainable.

One of the main challenges still to overcome is fabricating the reservoir. Photonic crystal cavities are notoriously well-known for their difficulty to fabricate with the correct properties. What's more, implementing the readout (weight-matrix) on chip could have an important effect on the total size of the system (reservoir+readout).

However, the relatively low power loss of the chip, together with having several outputs per cavity allows for a very modular structure, in which multiple cavities can be cascaded in a hierarchy to each solve a part of the problem at hand.

Funding

EU Horizon 2020 PHRESCO Grant (688579); BELSPO IAP P7-35 program Photonics@be; Research Foundation Flanders (FWO)(G024715N).