

Noise Analysis of Photonic Modulator Neurons

Thomas Ferreira de Lima^{*}, Alexander N. Tait, Hooman Saeidi, Mitchell A. Nahmias, Hsuan-Tung Peng, Siamak Abbaslou, Bhavin J. Shastri, and Paul R. Prucnal

Abstract—Neuromorphic photonics relies on efficiently emulating analog neural networks at high speeds. Prior work showed that transducing signals from the optical to the electrical domain and back with transimpedance gain was an efficient approach to implementing analog photonic neurons and scalable networks. Here, we examine modulator-based photonic neuron circuits with passive and active transimpedance gains, with special attention to the sources of noise propagation. We find that a modulator nonlinear transfer function can suppress noise, which is necessary to avoid noise propagation in hardware neural networks. In addition, while efficient modulators can reduce power for an individual neuron, signal-to-noise ratios must be traded off with power consumption at a system level. Active transimpedance amplifiers may help relax this tradeoff for conventional p-n junction silicon photonic modulators, but a passive transimpedance circuit is sufficient when very efficient modulators (i.e. low C and low V-pi) are employed.

Index Terms—Neuromorphic Computing, Neuromorphic Photonics, Analog Links, Neural Networks

I. INTRODUCTION

THE GAP between current computing capabilities and current computing needs ushered research in the field of neuromorphic computing [1]–[5]. This new field aims to bridge the gap between the energy efficiency of von Neumann computers and the human brain [6], [7]. As a consequence, this thrust spawned research into novel brain-inspired algorithms and applications uniquely suited to neuromorphic processors. These algorithms attempt to solve artificial intelligence tasks in real-time while using less energy. We posit that we can make use of the high parallelism and speed of photonics to bring the same neuromorphic algorithms to applications requiring multiple channels of multi-gigahertz analog signals, which digital processing struggles to process in real-time.

By combining the high bandwidth and parallelism of photonic devices with the adaptability and complexity attained by methods similar to those seen in the brain, photonic processors have the potential to be at least ten thousand times faster than state-of-the-art electronic processors while consuming less energy per computation [8]. An example of such an application is nonlinear feedback control; a very challenging task that involves computing the solution of a constrained quadratic optimization problem in real time. Neuromorphic photonics can enable new applications because there is no general-purpose

hardware capable of dealing with microsecond environmental variations [9].

These benefits can be accomplished by use of wavelength-division multiplexing (WDM), which explores the enormous bandwidth of optical waveguides (\sim THz). Integrated neuromorphic circuits based on WDM can now be manufactured using silicon photonics platforms. Tait et al. recently demonstrated a way of performing neural computations on WDM signals via a photodetector directly driving a modulator [10], which is capable of integrating hundreds of wavelength channels carrying gigahertz signals. We call this an O/E/O-based photonic neuron (Fig. 1).

In this architecture, photonic neurons output optical signals with unique wavelengths. These are multiplexed into a single waveguide and broadcast to all others, weighted, and photodetected. Each connection between a pair of neurons is configured independently by one MRR weight [11], [12], and the WDM carriers do not mutually interfere when detected by a single photodetector. Consequently, the physics governing the neural computation is fully analog and does not require any logic operation or sampling, which would involve serialization and sampling. Thus they exhibit distinct, favorable trends in terms of energy dissipation, latency, cross-talk and bandwidth when compared to electronic neuromorphic circuits [8, Sec. 5].

But the same physics also introduce new challenges, especially reconfigurability, integration, and scalability. Information carried by photons is harder to manipulate compared to electronic signals, especially nonlinear operations and memory storage. Photonic neurons described here solve that problem by using optoelectronic components (O/E/O), which can be mated with standard electronics providing reconfigurability. However, neuromorphic photonic circuits are challenging to scale up because they do not benefit from digital information, memory units and a serial processor, and therefore requires a physical unit for each element in a neural network, increasing size, area and power consumption. Here, integration costs must also be considered, since the advantages of using analog photonics (high parallelism and high bandwidth) must outweigh the costs of interfacing it with digital electronics (requiring both O/E and A/D conversion). The optimal cost-benefit tradeoff can be computed for particular applications by engineers, but one factor that has not been addressed in prior literature is the accumulation of noise in O/E/O neuromorphic analog links. This paper will provide a quantitative study in this dimension of O/E/O neural networks.

^{*}Corresponding author: tlima@princeton.edu

T.F.L., A.N.T., H.S., M.A.N., H.-T.P., S.A., B.J.S. and P.R.P. are with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA.

B.J.S. is with the Department of Physics, Engineering Physics & Astronomy, Queen's University, Kingston, ON K7L 3N6, Canada.

A.N.T. is currently with the Physical Measurement Laboratory, National Institute of Standards and Technology, Boulder, CO 80305, USA.

“cascadability” condition can be written with the following inequality:

$$\underbrace{R_{\text{TIA}}}_{\text{electrical}} \cdot \underbrace{P_L}_{\text{optical}} \cdot \underbrace{\left(\frac{\text{MD}}{V_{\text{pp}}}\right)}_{\text{mod. sens.}} > \frac{N_{\text{FO}}}{2T_{1/2}R_d\eta_{\text{pp}}} \quad (1)$$




Fig. 1. Diagram of an O/E/O photonic neuron.

The fact that lightwaves are transduced into electrical currents and back is very advantageous for implementing nonlinear operations, compared to all-optical strategies. If all devices worked perfectly, this scheme would work even with extremely low power and voltage levels. At the low-energy limit, one photon transforms into a electron-hole pair, whose charge can be used to modulate the transmission of an electro-optic device, sourcing a number of photons at the output of the circuit (Fig. 1). This mode of operation would require lossless optics, coupled with very sensitive photodetectors and modulators.

Because of such loss and inefficiency, we need to operate the neuron with stronger signals. In addition, due to the analog nature of this communication scheme, it is not immune to noise accumulation. Ultimately it is the shot noise and the noise of circuit components that currently prevents driving these O/E/O photonic neurons with quantum-scale power levels at room temperature.

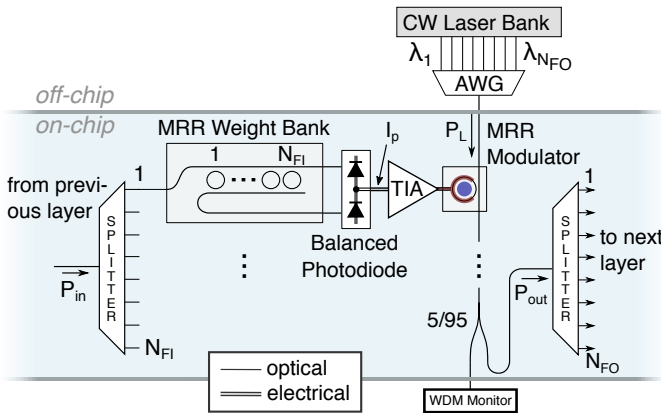


Fig. 2. Photonic neurons that receive inputs from one WDM broadcast medium and transmit into another, corresponding to the next layer. Distinct layers can thus reuse the same optical spectrum for broadcasting, much like a cellular telephone network reuses spectrum geographically. This strategy of using more waveguides for more layers can be extended, in principle, to an arbitrary depth.

In order for one layer to physically drive the next layer, enough laser power P_L must be provided to compensate for loss and power splitting due to fan-out to the next layer (N_{FO}). Alternatively, electrical gain via a high transimpedance (R_{TIA}) can provide the necessary amplification. This “gain

where MD is the modulation depth; $T_{1/2}$, mean transmission, R_d , photodiode’s responsivity, and η_{pp} is the optical point-to-point efficiency between connected neurons, representing excess optical loss between the output of one neuron and the input of the next, i.e. propagation loss and insertion loss of weighting devices and couplers.

Equation 1 describes the relationship between different kinds of gain: electrical, optical, modulator sensitivity. We draw attention to the $R_{\text{TIA}}/V_{\text{pp}}$ ratio (transimpedance over peak-to-peak voltage), which represents the *neuron’s sensitivity*, because it quantifies how much photocurrent is necessary to effect a full amplitude swing in the modulator. Since optical pump power (P_L) is an expensive resource, it is desirable to maximize this sensitivity in order to minimize overall power consumption. However, higher sensitivity comes at a price, as noise accumulation degrades the signal-to-noise ratio at the output (Sec. II).

B. Suppressing Noise Accumulation

Without careful design, analog circuits with long chains of cascaded neurons can accumulate and amplify noise, eventually burying signals under the noise floor. The optoelectronic devices shown in Fig. 2 are mostly linear and noisy, resulting in a signal-to-noise ratio (SNR) degradation, i.e. a noise factor greater than unity ($F > 1$). Fundamentally, this happens because not only they linearly amplify noise as well as signals, they generate noise on top of the output. To counter that effect, we need a device that can amplify signals while decreasing noise. This can be achieved with a nonlinear device – in our case, a modulator with a nonlinear transfer curve. The more nonlinear the modulator is, the more it can compensate for accumulating noise (Fig. 3).

C. Organization of the Paper

In this paper, we will numerically analyze the noise as it propagates across a neural network composed of O/E/O neurons (Sec. II). We will quantify the nonlinearity requirements for a neural circuit to suppress noise accumulation. We propose a simple experiment involving the cascability property of the neuron which verifies that noise is properly suppressed.

Following the noise analysis, we show two options of constructing the O/E/O circuit: one with a passive electronic link, and the other with an active transimpedance amplifier which provides energy gain in the electronic domain (Sec. III). The electronic gain can enhance the neuron’s response for weak input signals. The effect of this enhancement is a reduction of overall optical power levels in the circuit, leading to more energy efficiency.

II. NOISE PROPAGATION IN O/E/O NEURONS

Neural networks are known to be robust to noise [13]. In fact, noise can be exploited to train neural networks when other optimization algorithms might fail [14], [15]. Noise originating in hardware was used to implement on-line learning to a VLSI neural network [16].

There are two methods for avoiding noise propagation across a network. The first involves a collective approach, using redundant neurons encoding correlated information. This is called population coding in neuroscience, and is it required by physiological neural networks to overcome the noise generated by individual neurons [17], [18]. In essence, the estimation error for information carried by N neurons scales with $1/\sqrt{N}$ [18]. This concept has been adapted to machine learning and has been proven to mitigate noise in multilayer perceptron networks [13], the category into which photonic neural networks described in this paper falls. The second method, the focus of this paper, relies on every individual neuron to have noise suppressing circuit. This section studies the noise accumulation mechanisms within a single neuron and describes how a modulator's nonlinearity can suppress noise accumulation.

A. Modulator Nonlinearity as Noise Trimmer

Consider non-return-to-zero (NRZ) modulated signals at the input and output of each neuron. Assume that the neuron is biased so that zeros and ones fall on each side of its S-shaped transfer function (Fig. 3). Because the derivative of this function is relatively small in these regions, noise variance is reduced and the output looks "cleaner" than the input. This operating principle can be generalized to other modulation schemes and other transfer functions, but exploring all theoretical possibilities is beyond the scope of this paper.




Fig. 3. Modulator's transfer function showing the noise trimming principle. V_{pp} and $T(V)$ represent the peak-to-peak voltage and transfer function of this modulator, respectively. $T(V)$ is assumed to be a symmetric S-shape, so $T'(V_0) = T'(V_1)$. Here, an NRZ signal with probability distribution shown in the x-axis can be transduced into an optical signal with lower noise (y-axis). The inset illustrates how to extract the distribution empirically.

B. Sources of Noise in the OEO Link

The accumulation of thermal noise, shot noise, amplifier noise, and relative intensity noise (RIN) must be counteracted




Fig. 4. Simplified neuron circuit showing the sources of noise in each step of the photonic link.

by the modulator nonlinearity in order to guarantee that the SNR at the output equals the SNR at the input. This condition leads to the following equation:

$$\frac{1 - T_n^2}{\text{SNR}} = \frac{4T_n^2}{V_{pp}^2} \left(\underbrace{\frac{4k_B T \Delta f R_{TIA}^2}{R_L}}_{\text{thermal noise}} + \underbrace{2q \Delta f V_{pp} R_{TIA}}_{\text{shot noise}} + \underbrace{R_{TIA}^2 \Delta f I_{TIA,n}^2}_{V_{TIA,n}^2 = \text{TIA noise}} \right) + \underbrace{\text{RIN}^2 \left(1 + \frac{1}{\text{MD}^2} \right)}_{\text{RIN}} \quad (2)$$

where four sources of voltage noise (depicted in Fig. 4) are balanced against a potentially very low term T_n . The expression derivation, assumptions, and precise meaning of each term are described in Appendix A.

Equation 2 shows that the SNR converges to a finite level which cannot be arbitrarily large, since the term on the right is always positive. That is expected since noise is generated at every stage. To increase the neuron's SNR, we need to decrease the terms on the right side of the equation, in particular T_n^2 (noise transmission) and (R_{TIA}/V_{pp}) (sensitivity). The modulator's nonlinearity is very important in guaranteeing a high SNR. A completely linear modulator (leading to a noise transmission factor $T_n \rightarrow 1$) results in a completely noisy signal ($\text{SNR} \rightarrow 0$). This happens because noise increases more strongly than the signal at every stage, eventually reducing signal integrity at the infinite cascability limit. However, a completely nonlinear modulator ($T_n \rightarrow 0$) results in a high quality signal ($\text{SNR} \rightarrow (2\text{RIN})^{-2}$), which can be over 100 dB (assuming $\text{RIN} = -160 \text{ dB/Hz}$ [19]).

C. Noise vs. Gain Tradeoff

Modulators in reality have intrinsic nonlinearities, but are often operated in their linear region, often limiting their extinction ratio. However, nonlinearities are not only necessary by the mathematics of the neural network but are also here exploited to suppress noise. The microring resonator modulator, in particular, is known to have a Lorentzian-shape transfer function, which provides an ideal high-sensitivity S-shape for this application. Another candidate is a Mach-Zehnder modulator, with a sinusoidal transfer function. In either case, we expect their noise multipliers to lie strictly between 0 and 1.

Assuming $0 < T_n < 1$ and a fixed V_{pp} , then SNR can be increased by reducing transimpedance R_{TIA} as much as

possible, at the expense of greater power consumption (Eq. 1). In other words, a higher quality signal requires a lower gain circuit, resulting in a higher power consumption. This tradeoff can be exploited to save energy in neural network applications for which SNR is not a critical factor. Section III shows a few realistic estimations for silicon photonic modulator neuron implementations.

D. Autapse Test as Cascadability Standard

We use the notion of cascadability to verify whether a particular photonic neuron design can be scaled up to form large networks. There are three kinds of cascadability: physical, gain and noise.

a) Physical: A photonic neuron device is physically cascable if the nature of its output can be directly connected to another’s input. For example, the neuron introduced in Fig. 1 outputs a lightwave of a single wavelength, while receiving a number of inputs in parallel at different wavelengths. The signals are carried in the amplitude envelope (not phase, or polarization) of both input and output lightwaves. Because of that, this neuron can be interconnected and form arbitrarily-large neural networks so long as the fan-in to each neuron is equal or less than the number of WDM channels it is designed to support.

b) Gain: Beyond being physically cascable, each neuron must be able to provide enough optical power to excite the next layer of neurons, if needed. Equation 1 provides an estimate of the gain cascadability condition for the worst case scenario: one neuron, alone, delivering enough optical energy to N_{FO} other neurons. In this subnetwork, N_{FO} neurons can “replicate” the output of the initial neuron, which allow for N_{FO} subnetworks to process the input in parallel. In practical deep neural networks, however, multiple neurons share the burden of providing enough optical energy for the next layer. This metric can only be quantified if the shape and weight configuration of the neural network is known in advance, i.e. in the presence of an application benchmark, which is out of the scope of this article.

c) Noise: The other potential scalability limitation is noise accumulation. This is particularly important for deep networks. In the worst case scenario, the information contained in a signal fed to the first layer’s input must survive uncorrupted as it goes through the remaining layers of the network, even in the presence of noise. The calculations leading to Eq. 2 show that at the limit of infinitely deep neural networks, the SNR stabilizes to a certain value (solution of the equation) by balancing noise generation by the electronic O/E/O link and the noise suppression by the modulator’s nonlinearity.

d) Autapse Test: A simple experiment can be constructed to demonstrate and quantify all three conditions: a self-connection, also referred to as an *autapse*. With an autaptic connection with unity weight, the neuron emulates an infinite chain of neurons, where each connection delay τ represents a *virtual* neuron. The autapse experiment thus allows to study infinite cascability without producing an infinite chain of cascaded neurons. In this experiment, an initial pulse perturbation is sent to the neuron at $t = 0$, triggering an

output pulse in response. This output pulse travels through the autapse and, provided the gain cascability condition is met, excites another perturbation at $t = \tau$. The evolution of the pulse amplitude and shape at times $t = n\tau$ will determine whether this neuron has met both gain and noise cascability conditions. This experiment emulates an infinite series of neurons connected on a one-to-one basis. However, it can also emulate a one-to-N connectivity pattern if the autapse weight is set to $1/N$, which represents a $1/N$ loss in optical power between consecutive layers. Many-to-one and many-to-many connectivity can be extrapolated from this test but not directly emulated. We also note that this experiment tests for indefinite cascability, which might not be required in small neural networks.

Autapse experiments as described here were conducted in both modulator-based [10, Sec. E] and laser-based [23] photonic neurons, but they focused on demonstrating physical and gain cascability. Observing noise accumulation in an autapse is the next logical step in testing these devices.

III. TRANSIMPEDANCE AMPLIFIER IN SILICON PHOTONICS

Noise suppression relies on using the modulator’s intrinsic nonlinear transfer curve. But the nonlinearity is only observed if the voltage swing to the modulator is large enough (Fig. 3). The suite of standard silicon photonic components today are based on Ge photodetectors and p-n junction index modulators which possess a capacitance on the order of tens of femtofarads and require a voltage swing of a few volts. This voltage swing is provided by a transimpedance amplifier (TIA), which transduces photocurrent into voltage swing. Equation 1 already showed us that increasing the transimpedance reduces the optical pump power for the neuron. But increasing it too much limits the bandwidth of the circuit (inversely proportional to RC). In Sec. III-A we explore how the modulator’s parameters affect this power-bandwidth tradeoff. Sec. III-B discusses how active TIA can be used to mitigate some of that tradeoff. We show that an active TIA is no longer necessary to maintain 10 GHz bandwidth for sub-femtofarad nanophotonic devices.

A. Passive Transimpedance

An easy way to control the transimpedance in passive silicon photonic chips is to use a simple resistor in parallel with the modulator, whose value determines the transimpedance gain. This design is simple and works well, but the photodetector and modulator’s junction capacitance add in parallel with the transimpedance value. The achievable bandwidth is determined by the dominant pole of the circuit ($\Delta f = 1/2\pi R_{TIA}C$) (see Fig. 5). As a result, this limits how large the transimpedance can be, since the capacitances add up to 50 fF, which for a 10 GHz bandwidth corresponds to a maximum of $R_{TIA} = 320 \Omega$. This is significant because large networks will require on the order of 100 parallel wavelengths in a single waveguide. Assuming a maximum safe power of 100 mW per waveguide (avoiding nonlinear effects [24]), that gives us a maximum of 1 mW per wavelength, generating only ~ 0.3 V of swing at the modulator, far from the typical $V_{\pi} \sim 2$ V required

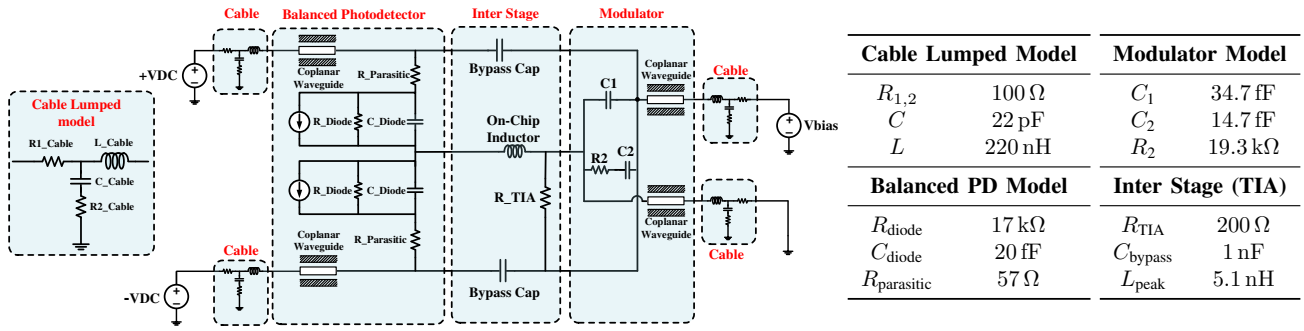


Fig. 5. Circuit schematic of a silicon photonic neuron with a passive transimpedance circuit whose bandwidth is enhanced via inductive gain peaking [20]. The circuit parameters are typical of recent literature and have been experimentally verified [20]–[22].

in silicon modulators [25]. Fig. 5 shows an implementation compatible with standard silicon photonic foundry chips.

In this case, gain cascading can only be achieved if (a) one uses a modulator with $V_{\pi} \sim 0.3$ V or (b) with a smaller capacitance, or (c) one uses an active TIA component which can decouple the transimpedance from the modulator’s capacitance, allowing for a higher gain with the same bandwidth.

The solution involving improving modulators is promising, as we are far from the fundamental limits of photonics [26]. Efficient modulators are key to cope with increasing demand in data communications, so research in this direction abounds. Exotic materials such as graphene are being used to reduce the switching energy of nanophotonic modulators toward sub-fJ [27]. Photonic crystals also offer an avenue for ultracompact O/E/O conversion. For example, Nozaki et al. [28] have demonstrated a nanophotonic (InP-based) O/E/O link with 1.6 fF capacitance and 25 k Ω transimpedance, with a voltage swing of 0.5 V. 40 μ W was sufficient to operate this O/E/O device.

B. Active Transimpedance Amplifier

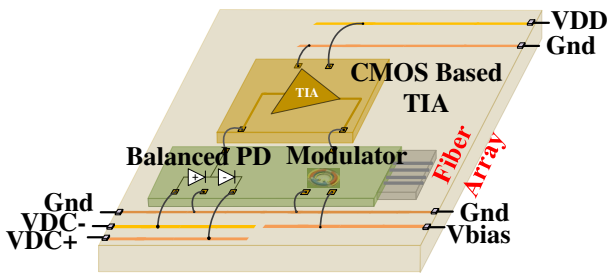


Fig. 6. Concept diagram of a silicon photonic integrated circuit packaged with a CMOS-based TIA. A flip-chip bonded alternative would yield similar electrical performance for the bandwidth of interest (10GHz).

Another way to provide gain without compromising bandwidth is to use an active TIA circuit [29], [30] instead of a RLC circuit. The TIA is designed to enhance the voltage swing of the modulator when the photocurrent is limited. In the short term, photonic integrated circuits can be coupled with CMOS-based TIAs via wirebonds or flip-chip bonding (Fig. 6). In the long run, however, these may be homogeneously integrated on the same chip, via a zero-change platforms [31].

Figure 7 shows the bandwidth performance simulation of the O/E/O circuit in Fig. 5 using two active TIA designs (one




Fig. 7. Transimpedance gain characteristics of the O/E/O module assuming an AC current source at one of the photodetectors, while measuring AC voltage amplitude across the modulator. The passive transimpedance circuit parameters were introduced in Fig. 5. The commercial off-the-shelf TIA ONET8531T (Texas Instruments) replaces R_{TIA} and L_{peak} . The ideal TIA design is similar to that of ref. [30] but was optimized specifically for this bandwidth and gain range. Note: The f_{3dB} values are: 14.8 GHz (passive), 8.0 GHz (commercial), 21.9 GHz (ideal). The f_{peak} value for the passive circuit is 11.8 GHz.

commercial and one ideal), and how they compare against the passive transimpedance approach. They were all designed to a bandwidth greater than ~ 10 GHz. As expected, the use of active TIA allowed us to achieve a ~ 17 times higher gain-bandwidth product (21.9 GHz \times 2.4 k Ω vs. 14.8 GHz \times 0.2 k Ω).

We note that TIAs fabricated with modern CMOS nodes have their maximum output swing voltage limited by the maximum V_{DD} of the transistor gates, which, in turn, is limited by the breakdown voltage of the node. For example, the 0.18 μ m-CMOS node offer TIAs with $2.5 V_{DD}$ [30], thus limiting the maximum achievable V_{pp} to about 1.8 V. Since the modulation depth (MD) is proportional to V_{pp} , this limit does not impact the gain cascading condition (Eq. 1). But if $V_{pp} \leq 1.8$ V $<$ V_{π} , the modulator will operate in a more linear regime, which will then transmit more noise (higher transmission factor T_n), impacting the noise cascading condition (Eq. 2). As a result, neuromorphic photonics would benefit from photodiodes and modulators with driving voltages lower than the V_{DD} of modern technology nodes.

We study the effect of the ideal active TIA in the cascading condition of the O/E/O neuron. Equations 1 and 2 relate

TABLE I
COMPUTED TRADEOFFS OF VARIOUS O/E/O DESIGNS.

Modulator Class	TIA	V_{pp} V	R_{TIA} k Ω	P_L dBm	SNR (dB) ($T_n = 0.5$)
p-n junction [10], [22]	passive	4.8	0.4	26	64
	active	4.8	2.8	17	54
graphene [27] [32], [33, Fig.6]	passive	0.75	2.0	18	49
		0.1	1.2	11	41
	active	0.75	2.8	17	41
		0.1	2.8	8	24

This table was computed by using equations (1–2) with the following parameters: $T_n = 0.5$; $k_B T = 4.11 \times 10^{-21}$ J; $\Delta f = 10$ GHz; $q = 1.6 \times 10^{-19}$ C; $I_{TIA,n} = 20 \times 10^{-12}$ A/ $\sqrt{\text{Hz}}$ (active), 0 (passive); $RIN = 1 \times 10^{-6}$ [19]; MD = 0.61 (p-n junction), 0.33 (graphene); $N_{FO} = 10$; $T_{1/2} = 0.5$; $R_d = 1$ A/W (p-n junction), 0.35 A/W (graphene [32]); $\eta_{pp} = 0.5$; $R_L = R_{TIA}$ (passive), $R_L = \infty$ (active). R_{TIA} values were chosen to fulfill $2\pi RC = \Delta f^{-1}$, with capacitances listed in the references.

the autapse SNR to multiple design parameters, such as V_{pp} , P_L , R_{TIA} and bandwidth Δf . There are two main trends to bear in mind. With a passive transimpedance gain, there is a clear tradeoff between bandwidth and power consumption for a fixed modulator design. The larger the bandwidth Δf , the lower the load R_{TIA} has to be, and therefore the larger the optical power P_L to guarantee gain cascability (eq. 1). The autapse SNR is constant over values of Δf , since noise terms are proportional to $R_{TIA} \cdot \Delta f$ terms. On the other hand, an active TIA may provide a higher transimpedance-bandwidth product ($R_{TIA} \cdot \Delta f$), but introduces more noise and consumes more power than a passive transimpedance. However, a higher gain allows laser pump power to be lower. Since off-chip lasers tend to be inefficient, this can result in overall system-wide power savings. The TIA also allows for dynamic transimpedance tuning, which is useful to reduce power consumption at times when noise is not critical to the application.

The dependence between SNR and individual parameters shows very simple trends – linear or quadratic. Therefore, instead of displaying arrays of plots, we chose to use Table I to show a few specific examples tied to the devices reported in the literature. The table shows the effect of adding an active TIA to a p-n junction standard-platform modulator vs. a graphene-based ultra-sensitive modulator with V_π of 0.1 V or 0.75 V. A standard-platform silicon photonic neuron can benefit much more from the TIA; the laser power requirement decreases by 10-fold (from 26 dBm to 17 dBm), with the expense of a decrease in the SNR limit (from 64 dB to 54 dB). In contrast, almost no enhancement is seen by the graphene-based neuron, because it features low capacitance and therefore the passive R_{TIA} can be high without impacting bandwidth. With that said, TIAs will continue to be useful for applications requiring marrying conventional modulators, and low-power and lower bandwidth budgets, two factors that scale favorably to justify their use.

IV. CONCLUSION

We presented a quantitative overview of the interplay between electrical gain, optical gain, and modulator sensitivity

in O/E/O photonic neurons. Although these metrics have positive influences over the accuracy and stability of the neural network, they can impose difficult power and noise requirements.

We showed that current p-n junction-based silicon photonic platforms do not support highly-interconnected photonic neural networks unless they use (a) more sensitive modulators, (b) active transimpedance amplifiers (TIAs), or (c) operate at a sub-GHz bandwidth.

This occurs because modulators need a large voltage swing to reach the nonlinear threshold in their nonlinear transfer function, which suppresses noise directly between one neural layer and the next – a requirement for cascable analog links. This swing can either be achieved by increasing optical pump power at the modulator or by providing electric transimpedance gain. However, optical gain is limited by optical nonlinearities in waveguides (and potentially power budgets), whereas transimpedance gain is inversely proportional to the bandwidth of the circuit.

The analysis shown here also applies to other kinds of electro-optic transducers, such as lasers. In particular, excitable lasers [34] have a zero-or-one thresholding response that serves as both an amplifier and a noise suppression device. Cascadability properties have recently been demonstrated in one such devices [23]. Future work will involve testing and quantifying noise cascability by means of an autapse test experiment, setting another useful benchmark in the field of analog neuromorphic photonics.

APPENDIX A

DERIVATION OF CASCADABILITY EQUATIONS

In this section, we will derive the expressions contained in Eqs. 1 and 2. We will be using the following notations:

$$\bar{x}(t) \equiv \mathbb{E}[x(t)]; \quad \langle x(t) \rangle \equiv \int_0^T x(t) dt.$$

We start by computing the expressions of the photocurrent in Fig. 4. The photodetector is often modelled as a linear, added with shot noise and thermal noise terms following normal distributions. This approximation is valid if the optical power amplitude is much slower than the photodiode’s response time.

$$I = R_d \cdot (P_{in} + P_{in,n}) + I_{shot} + I_{thermal} \quad (3)$$

where

$$P_{in,n} \sim \mathcal{N}(0, P_n^2) \quad (4)$$

$$I_{shot} \sim \mathcal{N}(0, 2q_e \Delta f \bar{I}) \\ \sim \mathcal{N}(0, 2q_e \Delta f R_d \bar{P}_{in}) \quad (5)$$

$$I_{thermal} \sim \mathcal{N}\left(0, \frac{4k_B T}{R_L} \Delta f\right) \quad (6)$$

are random variables sampled in time. Note: I_{shot} should technically not be a normal random variable, but in many textbooks it is modelled that way. It is a good approximation for large values of $\bar{I}/\Delta f q_e$ (average event count).

Similarly, the transimpedance amplifier is modelled as a linear component with input referred noise.

$$V = R_{\text{TIA}} \cdot I + V_{\text{bias}} + V_{\text{TIA},n} \quad (7)$$

$$V_{\text{TIA},n} \sim \mathcal{N}(0, R_{\text{TIA}}^2 I_{\text{TIA},n}^2 \Delta f). \quad (8)$$

The $I_{\text{TIA},n}$ is often presented in the literature in units of pA/ $\sqrt{\text{Hz}}$ and is assumed invariant for TIAs of the same design on the same technology node. V_{bias} is assumed to be noise-free.

So far, it can be easily shown that V is a random variable with Gaussian noise, resulting in

$$V = \bar{V} + V_{\text{noise}} \quad (9)$$

$$\bar{V} = R_{\text{TIA}} R_d \bar{P}_{\text{in}} + V_{\text{bias}} \quad (10)$$

$$V_{\text{noise}} \sim \mathcal{N}(0, V_n^2). \quad (11)$$

So far all noise terms have been additive. The next expression for the power envelope of the output optical signal P_{out} .

$$P_{\text{out}} = P_L (1 + n_{\text{RIN}}) \cdot T(V) \quad (12)$$

$$n_{\text{RIN}} \sim \mathcal{N}(0, \text{RIN}^2) \quad (13)$$

where $T(V)$ is the transfer function of the modulator (Fig. 3) and RIN is the relative intensity noise – r.m.s. integrated from 0 to Δf , $\int |\text{RIN}(\text{dBc}/\text{Hz})| df$ – at the pump. There are two nonlinear transformations in this expression that render analytical treatment difficult: RIN is a multiplicative noise term, and $T(\cdot)$ is a nonlinear function. That means that the probability distribution function of the output noise does not have an analytical form. but we can proceed to compute the signal-to-noise ratios at the input and output.

$$\text{SNR}_{\text{in}} = \frac{\langle \overline{P_{\text{in}}^2} \rangle - \langle \bar{P}_{\text{in}} \rangle^2}{\langle (P_{\text{in}} - \bar{P}_{\text{in}})^2 \rangle} = \frac{\langle \overline{P_{\text{in}}^2} \rangle - \langle \bar{P}_{\text{in}} \rangle^2}{P_n^2} \quad (14)$$

$$\text{SNR}_{\text{out}} = \frac{\langle \overline{P_{\text{out}}^2} \rangle - \langle \bar{P}_{\text{out}} \rangle^2}{\langle (P_{\text{out}} - \bar{P}_{\text{out}})^2 \rangle} \quad (15)$$

$$F = \frac{\text{SNR}_{\text{in}}}{\text{SNR}_{\text{out}}} = 1 \quad (16)$$

The noise cascability condition is expressed in Eq. 16. The idea is to compute Eq. 15 and then express the resulting P_n^2 term in terms of SNR_{in} per Eq. 14.

Equation 16 is general; it should work for arbitrary amplitude modulation schemes with known probability distribution. From here onwards, we need to choose a convenient input signal waveform that will allow us to compute an approximate expression. We choose an unbiased non-return-to-zero (NRZ) signal as modulation scheme, which is very common in optical links and makes it easy to extract values from experimental papers and plug them into these equations.

$$P_{\text{in}}(t) = P_{\text{in},n} + \begin{cases} P_1 & 0 \leq t < T/2 \\ P_0 & T/2 \leq t < T \end{cases} \quad (17)$$

We also make a few assumptions about how the modulator is biased, as hinted in Fig. 3. First, we assume that V_{bias} is picked so that

$$T \left(R_{\text{TIA}} R_d \cdot \frac{P_1 + P_0}{2} + V_{\text{bias}} \right) = \frac{1}{2} \equiv T(V_{1/2}).$$

Second, we denote $T_{0,1} = T(V_{0,1}) = T(V(P_{0,1}))$, $\Delta T \equiv T_1 - T_0$, $\text{MD} \equiv \Delta T / (T_1 + T_0)$, $\Delta P \equiv P_1 - P_0$, and $V_{\text{pp}} = R_{\text{TIA}} R_d \Delta P$.

Third, we assume that $T(V)$ is symmetric around $V_{1/2}$ and thus $T'(V_1) = T'(V_0) \equiv T'$ as hinted in Fig. 3. It is useful to define the quantity we named the noise transmission T_n as (cf. Fig. 3):

$$T_n \equiv \frac{T' \cdot V_{\text{pp}}}{\Delta T}. \quad (18)$$

Here, we are able to define the gain cascability condition. We simply state that the output optical amplitude generated by an input amplitude ΔP should be able to provide at least an amplitude equivalent to ΔP to each of the fan-out neurons (N_{FO}). In other words:

$$\frac{P_L \Delta T \eta_{\text{pp}}}{N_{\text{FO}}} > \Delta P, \quad (19)$$

where η_{pp} is efficiency term from point to point, accounting for insertion losses of the WDM networking grid. Eq. 19 should lead directly to Eq. 1.

With these assumptions, Eq. 14 and the numerator of Eq. 15 become, respectively:

$$\text{SNR}_{\text{in}} = \frac{\Delta P^2}{4P_n^2} \quad (20)$$

$$\langle \overline{P_{\text{out}}^2} \rangle - \langle \bar{P}_{\text{out}} \rangle^2 = \frac{P_L^2 \Delta T^2}{4} \quad (21)$$

In order to complete the derivation, we must know how to approximate the propagation of P_n^2 in Eq. 12. As shown in Fig. 3, if the variance of the input signal around their means is small, we can linearize $T(V)$ such that this approximation holds:

$$T(V) \approx T(\bar{V}) + T'(\bar{V}) \cdot V_{\text{noise}} \quad (22)$$

Plugging Eq. 22 into Eq. 12 and observing that $n_{\text{RIN}} \cdot V_{\text{noise}} = 0$ because they are independent, then Eq. 15 should become equal to the desired Eq. 2.

REFERENCES

- [1] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, B. Brezko, I. Vo, S. K. Esser, R. Appuswamy, B. Taba, A. Amir, M. D. Flickner, W. P. Risk, R. Manohar, and D. S. Modha, "A million spiking-neuron integrated circuit with a scalable communication network and interface," *Science*, vol. 345, no. 6197, pp. 668–673, 2014.
- [2] S. B. Furber, F. Galluppi, S. Temple, and L. A. Plana, "The SpiNNaker project," *Proceedings of the IEEE*, vol. 102, no. 5, pp. 652–665, May 2014.
- [3] B. Benjamin, P. Gao, E. McQuinn, S. Choudhary, A. Chandrasekaran, J.-M. Bussat, R. Alvarez-Icaza, J. Arthur, P. Merolla, and K. Boahen, "Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations," *Proceedings of the IEEE*, vol. 102, no. 5, pp. 699–716, May 2014.

- [4] K. Meier, "A mixed-signal universal neuromorphic computing system," in *2015 IEEE International Electron Devices Meeting (IEDM)*, 2015, pp. 4.6.1–4.6.4.
- [5] M. Davies, N. Srinivasa, T. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, Y. Liao, C. Lin, A. Lines, R. Liu, D. Mathaikutty, S. McCoy, A. Paul, J. Tse, G. Venkataraman, Y. Weng, A. Wild, Y. Yang, and H. Wang, "Loihi: A neuromorphic manycore processor with on-chip learning," *IEEE Micro*, vol. 38, no. 1, pp. 82–99, 2018.
- [6] B. Marr, B. Degnan, P. Hasler, and D. Anderson, "Scaling energy per operation via an asynchronous pipeline," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 21, no. 1, pp. 147–151, Jan 2013.
- [7] J. Hasler and H. B. Marr, "Finding a roadmap to achieve large neuromorphic hardware systems," *Front. Neurosci.*, vol. 7, no. 118, 2013.
- [8] T. Ferreira de Lima, B. J. Shastri, A. N. Tait, M. A. Nahmias, and P. R. Prucnal, "Progress in neuromorphic photonics," *Nanophotonics*, vol. 6, no. 3, pp. 577–599, jan 2017.
- [9] T. Ferreira de Lima, H.-T. Peng, A. Tait, M. A. Nahmias, H. Miller, B. J. Shastri, and P. Prucnal, "Machine Learning with Neuromorphic Photonics," *Journal of Lightwave Technology*, p. 1, 2019.
- [10] A. N. Tait, T. Ferreira de Lima, M. A. Nahmias, H. Miller, H.-T. Peng, B. J. Shastri, and P. R. Prucnal, "A silicon photonic modulator neuron," *arXiv preprint arXiv:1812.11898*, pp. 8–9, Dec. 2018.
- [11] A. N. Tait, A. X. Wu, T. Ferreira de Lima, E. Zhou, B. J. Shastri, M. A. Nahmias, and P. R. Prucnal, "Microring weight banks," *IEEE J. Sel. Top. Quantum Electron.*, vol. 22, no. 6, 2016.
- [12] A. N. Tait, T. Ferreira de Lima, M. A. Nahmias, B. J. Shastri, and P. R. Prucnal, "Multi-channel control for microring weight banks," *Opt. Express*, vol. 24, no. 8, pp. 8895–8906, Apr 2016.
- [13] Y. L. Lee, "Input Noise Immunity of Multilayer Perceptrons," *ETRI Journal*, vol. 16, no. 1, pp. 35–43, apr 1994.
- [14] A. H. Abdelaziz, S. Watanabe, J. R. Hershey, E. Vincent, and D. Kolossa, "Uncertainty propagation through deep neural networks," *Interspeech*, sep 2015.
- [15] Y. Bengio, N. Léonard, and A. Courville, "Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation," pp. 1–12, 2013.
- [16] A. Jayakumar and J. Alspector, "A Cascadable Neural Network Chip Set With On-chip Learning Using Noise And Gain Annealing," in *Proceedings of the IEEE Custom Integrated Circuits Conference*. IEEE, 1992, pp. 19.5.1–19.5.4.
- [17] H. P. Snippe and J. J. Koenderink, "Discrimination thresholds for channel-coded systems," *Biological Cybernetics*, vol. 66, no. 6, pp. 543–551, apr 1992.
- [18] H. S. Seung and H. Sompolinsky, "Simple models for reading neuronal population codes." *Proceedings of the National Academy of Sciences*, vol. 90, no. 22, pp. 10749–10753, nov 1993.
- [19] P. Resneau, M. Calligaro, M. Krakowski, H. Liu, M. Hopkinson, A. Somers, A. Forchel, and J. P. Reithmaier, "High power and very low noise operation at 1.3 and 1.5 μm with quantum dot and quantum dash Fabry-Perot lasers for microwave links," in *Advanced Free-Space Optical Communication Techniques/Applications II and Photonic Components/Architectures for Microwave Systems and Displays*, L. J. Sjöqvist, R. A. Wilson, and T. J. Merlet, Eds., vol. 6399, no. October 2006, sep 2006, p. 63990K.
- [20] A. Novack, M. Gould, Y. Yang, Z. Xuan, M. Streshinsky, Y. Liu, G. Capellini, A. E.-J. Lim, G.-Q. Lo, T. Baehr-Jones, and M. Hochberg, "Germanium photodetector with 60 GHz bandwidth using inductive gain peaking," *Optics Express*, vol. 21, no. 23, p. 28387, nov 2013.
- [21] J. M. Lee and W. Y. Choi, "An equivalent circuit model for germanium waveguide vertical photodetectors on Si," *2014 International Topical Meeting on Microwave Photonics / the 9th Asia-Pacific Microwave Photonics Conference, MWP/APMP 2014 - Proceedings*, pp. 139–141, 2014.
- [22] Z. Xuan, Y. Ma, Y. Liu, R. Ding, Y. Li, N. Ophir, A. E.-J. Lim, G.-Q. Lo, P. Magill, K. Bergman, T. Baehr-Jones, and M. Hochberg, "Silicon microring modulator for 40 Gb/s NRZ-OOK metro networks in O-band," *Optics Express*, vol. 22, no. 23, p. 28284, nov 2014.
- [23] H.-T. Peng, T. Ferreira de Lima, M. A. Nahmias, A. N. Tait, B. J. Shastri, and P. R. Prucnal, "Autaptic circuits of integrated laser neurons," in *Conference on Lasers and Electro-Optics*, May 2019, pp. 1–2.
- [24] C. Koos, L. Jacome, C. Poulton, J. Leuthold, and W. Freude, "Nonlinear silicon-on-insulator waveguides for all-optical signal processing," *Optics Express*, vol. 15, no. 10, p. 5976, may 2007.
- [25] P. Dong, S. Liao, D. Feng, H. Liang, D. Zheng, R. Shafiqhi, C.-C. Kung, W. Qian, G. Li, X. Zheng, A. V. Krishnamoorthy, and M. Asghari, "Low V_{pp}, ultralow-energy, compact, high-speed silicon electro-optic modulator." *Optics express*, vol. 17, no. 25, pp. 22484–22490, 2009.
- [26] V. J. Sorger, R. Amin, J. B. Khurgin, Z. Ma, H. Dalir, and S. Khan, "Scaling vectors of attoJoule per bit modulators," *Journal of Optics*, vol. 20, no. 1, p. 014012, jan 2018.
- [27] R. Amin, Z. Ma, R. Maiti, S. Khan, J. B. Khurgin, H. Dalir, and V. J. Sorger, "Attojoule-efficient graphene optical modulators," *Applied Optics*, vol. 57, no. 18, pp. D130–D140, 2018.
- [28] K. Nozaki, S. Matsuo, T. Fujii, K. Takeda, E. Kuramochi, A. Shinya, and M. Notomi, "Ultracompact O-E-O converter based on fF-capacitance nanophotonic integration," in *Conference on Lasers and Electro-Optics*. Optical Society of America, 2018, p. SF3A.3.
- [29] J. Kim and J. F. Buckwalter, "A 40-Gb/s Optical Transceiver Front-End in 45 nm SOI CMOS," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 3, pp. 615–626, mar 2012.
- [30] Jun-De Jin and Shawn Hsu, "A 75-dB $\cdot \Omega$ 10-Gb/s Transimpedance Amplifier in 0.18- μm CMOS Technology," *IEEE Photonics Technology Letters*, vol. 20, no. 24, pp. 2177–2179, dec 2008.
- [31] V. Stojanović, R. J. Ram, M. Popović, S. Lin, S. Moazeni, M. Wade, C. Sun, L. Alloati, A. Atabaki, F. Pavanello, N. Mehta, and P. Bhargava, "Monolithic silicon-photonics platforms in state-of-the-art CMOS SOI processes [Invited]," *Optics Express*, vol. 26, no. 10, p. 13106, may 2018.
- [32] I. Goykhman, U. Sassi, B. Desiatov, N. Mazurski, S. Milana, D. de Fazio, A. Eiden, J. Khurgin, J. Shappir, U. Levy, and A. C. Ferrari, "On-Chip Integrated, SiliconGraphene Plasmonic Schottky Photodetector with High Responsivity and Avalanche Photogain," *Nano Letters*, vol. 16, no. 5, pp. 3005–3013, may 2016.
- [33] K. Nozaki, S. Matsuo, A. Shinya, and M. Notomi, "Amplifier-Free Bias-Free Receiver Based on Low-Capacitance Nanophotodetector," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 24, no. 2, pp. 1–11, mar 2018.
- [34] P. R. Prucnal, B. J. Shastri, T. Ferreira de Lima, M. A. Nahmias, and A. N. Tait, "Recent progress in semiconductor excitable lasers for photonic spike processing," *Advances in Optics and Photonics*, vol. 8, no. 2, p. 228, jun 2016.