Integrated-Photonic Characterization of Single-Photon Detectors for Use in Neuromorphic Synapses

S.M. Buckley,[*] A.N. Tait, J. Chiles, A.N. McCaughan, S. Khan, R.P. Mirin, S.W. Nam, and J.M. Shainline

National Institute of Standards and Technology, 325 Broadway, Boulder, Colorado 80305, USA

Large-scale spiking neural networks have been designed based on superconducting-nanowire single-photon detectors (SNSPDs) as receiver elements for photonic communication between artificial neurons and synapses. Such large-scale artificial neural networks will require thousands of waveguide-integrated SNSPD devices on a single chip. Efficient waveguide-coupled SNSPDs and high-throughput methods for their characterization at this scale are very challenging. Here, we design, fabricate, and measure SNSPDs that are compatible with these large-scale networks. We demonstrate integrated-photonic circuits to simultaneously characterize many waveguide-coupled SNSPDs with a single-fiber input. We achieve up to 15 waveguide-coupled SNSPDs in a single integrated-photonic circuit and a total of 49 waveguide SNSPDs with a detection plateau out of 49 tested. This is perhaps the largest number of SNSPDs integrated in a photonic circuit to date. We employ several types of photonic circuits to enable rapid and reliable characterization of detector performance. We further demonstrate several important synaptic functions of these detectors. These include a binary response in the detectors with average incident photon numbers ranging from less than 10^{-3} to greater than 10, indicating that synaptic responses based on these detectors are independent of the number of incident photons in a communication pulse. Such a binary response is ideal for communication in neural systems. We further demonstrate that the response has a linear dependence of the output-current pulse height on the bias current with up to a factor of 1.7 tunability in pulse height, a function that cannot be obtained without a detection plateau.

I. INTRODUCTION

A defining feature of neural systems is the presence of highly interconnected neurons making thousands of synaptic connections that are often spatially distant. Many electronic circuits based on semiconductors [1–9] and superconductors [10–17] have been explored to achieve synaptic and neuronal computations but these circuits attempt to perform communication as well as computation in the electronic domain. It has been argued extensively that artificial neural systems will benefit from employing photonic rather than electronic communication to achieve this high fan-out [18–21]. A number of implementations of neural systems employing light for communication have been explored [22–26] but these typically do not include the same diversity of circuits for learning and computation and require much higher energy light sources and detectors. We have proposed an approach that combines photonic circuits for communication with superconducting electronics for computation [20,27–29]. In this approach, a light-emitting diode (LED) and transmitter circuit mimics

the firing operation of the neuron and the phase transition of a SNSPD converts the optical pulse to a binary electrical pulse, in analogy with the electrical-to-chemical conversion that occurs at the synaptic cleft. Synaptic, dendritic, and neuronal computations make use of the inherent thresholding nonlinearity of Josephson junctions.

While superconducting optoelectronic hardware has been proposed as a means to achieve large-scale neural systems, significant experimental progress is required to prove the practical viability. Before the full neuronal circuit can be realized, individual subcomponents must be fabricated and tested, so that each is performing as needed before integration. There are four main device subcomponents under investigation for future integration [30]: waveguide-coupled SNSPDs (this work), waveguide-coupled LEDs [31], amplifier devices [32], and SNSPDs integrated with Josephson junctions. While waveguide-coupled SNSPDs have been demonstrated previously, neural networks with modest numbers of neurons will require thousands of SNSPD devices integrated with waveguides, light sources, and superconducting electronics. Therefore one of the major contributions of this work is the high yield of the detectors on single-mode silicon waveguides,

_____
[*]sonia.buckley@nist.gov

allowing integration of large numbers (up to 15 in this work) of SNSPDs with saturating internal efficiency in a single photonic circuit.

There are many aspects of a SNSPD that must be considered when designing a synapse, including the critical current, inductance, device area, operating temperature, detection efficiency, detection plateau width, device yield, and superconductor material. All of these properties are interconnected and therefore they must be optimized self-consistently in the system context. Fabrication and testing of waveguide-SNSPDs with the most likely parameters to be used in a system integration is therefore a necessary precursor to building neuronal circuits. In Sec. II, we focus on experimental characterization and measurement techniques of waveguide-coupled SNSPDs with parameters similar to those expected to be utilized in the full neuronal-circuit demonstrations. Because synapses are the most numerous elements in neural systems, the ability to rapidly characterize many SNSPDs in the integrated-photonic environment in which they will be utilized is necessary both initially and for any further system optimization. Therefore, both the characterized parameters and the measurement techniques are highly relevant for the synaptic application.

The requirement for saturated internal efficiency in the definition of the yield of SNSPDs has often been overlooked in the literature, despite being of critical importance for most applications due to its direct impact on detection efficiency. It is of particular importance in this work for implementations of synaptic function. In Sec. III, we show how technological improvements that lead to high-yield detectors directly enable a technique for setting analog synaptic weights, an exciting space in applied physics. While synaptic weights may be either digital or analog, signaling between neurons is binary. Binary signaling is an important technique to reduce noise in communication from a neuron to its synaptic connections and the effectiveness of communication in neural systems is due in part to the binary nature of action potentials. Binary communication is therefore a key feature of the superconducting opto-electronic hardware platform and we demonstrate the conversion of photonic pulses with intensities varying by up to 4 orders of magnitude to binary electrical pulses. Finally, in Sec. V, we discuss the relevance of this work in a broader context.

II. INTEGRATED-PHOTONIC CHARACTERIZATION OF SINGLE-PHOTON DETECTORS

Waveguide-integrated SNSPDs are being pursued for a variety of applications in integrated quantum optics, metrology, and large-scale neural systems. Several review articles have been written on the subject [33–35], all of which present positive prospects for the future of systems integrating large numbers of waveguide-coupled SNSPDs. However, it is also clear that further work is needed to improve the yield of high-quality high-performance detectors. Refs. [35] and [36] suggest that a "pick-and-place" technique will allow prescreening of high-quality detectors before individual placement in photonic circuits, while Ref. [34] predicts that once SNSPDs are fabricated in commercial foundries, yield issues will be eliminated.

In this work, we follow most closely the approach suggested in Ref. [33] to realizing high-yield detectors, by using the amorphous material WSi. To date, this is the material that has been shown to yield the largest number of detectors with a broad plateau [37], although other amorphous materials with higher transition temperatures are strong candidates as well [38] and may ultimately prove more promising due to reduced cryogenic complexity. WSi nanowire detectors have been demonstrated in kilopixel arrays, with yields greater than 99% [39].

Recent work has shown that detection plateaus can be achieved in much wider wires than previously thought possible [40,41]. We use the silicon-rich WSi developed in Ref. [41], which allows straightforward fabrication and patterning with photolithography while giving a broad detection plateau. The ability to photolithographically pattern SNSPDs greatly reduces the fabrication time and complexity compared to electron-beam (e-beam) lithography, which has much longer write times and is a more manual process. Every e-beam step eliminated significantly decreases the fabrication turnaround time. For these reasons, we study WSi nanowires patterned with photolithography in this work. It is the use of this material that has enabled the integration of large numbers of detectors with saturated internal efficiency in a single device, a task that has previously proved challenging [36,42]. In this work, up to 15 working waveguide-coupled superconducting detectors are tested in a single integrated device [31,34,36,42–46], providing a scalable demonstration of highly complex integrated-photonic circuits utilizing SNSPDs.

The ability to fabricate devices with large numbers of waveguide-coupled SNSPDs leads to challenges in testing. One major challenge in the measurement of waveguide-integrated SNSPDs is to couple light from an optical fiber into an on-chip waveguide. Due to variations in fiber coupling efficiency and the complexity of the fiber coupling and/or packaging procedure, it is helpful to make as many measurements as possible using a single-fiber input, particularly when one wishes to compare the efficiency of different detector designs. We therefore use waveguides and beam splitters to route light to different devices from a single-fiber input. This is in contrast to the previous work on the characterization of waveguide-coupled SNSPDs, in which the characterization devices typically only include a single SNSPD per input grating coupler [47–55]. The integrated devices in this work therefore allow rapid

characterization without realignment and reliance on grating or coupler reproducibility.

In Sec. II A, we describe the integrated-photonic tree device, which evenly divides input light between seven detectors and can be used to compare different detector designs. In Sec. II B, we describe the high-dynamic-range detector array (HiDRA). The full integrated circuits consist of up to 15 detectors, integrated with up to 18 beam splitters and/or beam taps and several centimeters of waveguides. One promising light source used in the proposed superconducting spiking neural systems operates at wavelengths of around 1.22 μm [27,31]. Therefore, we design and demonstrate all devices for operation at both the standard 1.55 μm wavelength used for telecommunications and most quantum-optics experiments as well as 1.22 μm.

A. Branching tree structures

The "branching tree" structure provides a method for evenly distributing light to a number of different detectors. This allows comparison of the performance of several detectors under uniform waveguide-coupled illumination. A microscope image of such a device is shown in Fig. 1(a). Light from a single-fiber input is split evenly by a series of beam splitters and routed to seven detectors and an output grating. The output grating is used for fiber alignment during packaging at room temperature. The path lengths from the input grating to all detectors are equal. An enlarged view of a detector is shown in Fig. 1(b) and a SEM image of a SNSPD detector on a waveguide is shown in Fig. 1(c). We also implement a version of this device for room-temperature characterization. The room-temperature structures include an output grating after the hairpin device, as well as output reference gratings for normalization of the output power. Absorption of the hairpins at room temperature can then be compared to measurements of hairpins at cryogenic temperature. Figure 1(d) shows a microscope image of such a tree structure used for room-temperature characterization. The measurements described in this section are all performed at 1.55 μm on devices designed for operation at this wavelength. Once cooled, the current-voltage curves for each SNSPD detector are measured. From these data, the critical current of the standard detectors (200 μm long, 500 nm wide) is found to be 10 μA.

With the light evenly divided, comparisons can be made between detectors with different geometric parameters. A layout error causes one of the seven detectors on each tree to fail, so each tree in this work compares six detectors. For the first tree, the width of the SNSPDs is varied from 400 nm to 1.5 μm, keeping the length of the nanowire constant [see Fig. 7(c)]. The results are shown in Fig. 2. The detector count rate is plotted in Fig. 2(a) when the detectors are flood illuminated from above and in Fig. 2(b) when the detectors are illuminated through the waveguide tree structure. In both the uniform and waveguide-illuminated plots, the count rate is normalized to the count rate on the brightest detector. The detectors under uniform illumination require approximately 50 dB more light for the same count rate as the waveguide-coupled devices. A plateau region where the count rate on the detector is independent of bias current is observed for every width measured, as expected based on the results in Ref. [41]. This plateau region is the desired bias current range of operation. The narrower the wire, the steeper is the transition between the onset of detection and the plateau region. This has also been characterized for this and different film compositions for uniformly illuminated detectors in Ref. [41].

We further use the tree structure to investigate the height and width of the plateau region. The height of these curves quantifies the relative detection efficiency of wires with differing widths, while the width of the plateau quantifies the range of bias currents providing saturated internal quantum efficiency. We find very little dependence of the detection efficiency on the wire width (see Sec. IV F). Figure 2(c) shows the plateau width, P, versus the wire width, w. The value of P is calculated as the ratio of the width of the region between the value of I_{bias} at 0.9 of the maximum count rate and the switching current I_{sw} ($\Delta I_{plateau}$) to the switching current I_{sw}: $P \equiv \Delta I_{plateau}/I_{sw}$.
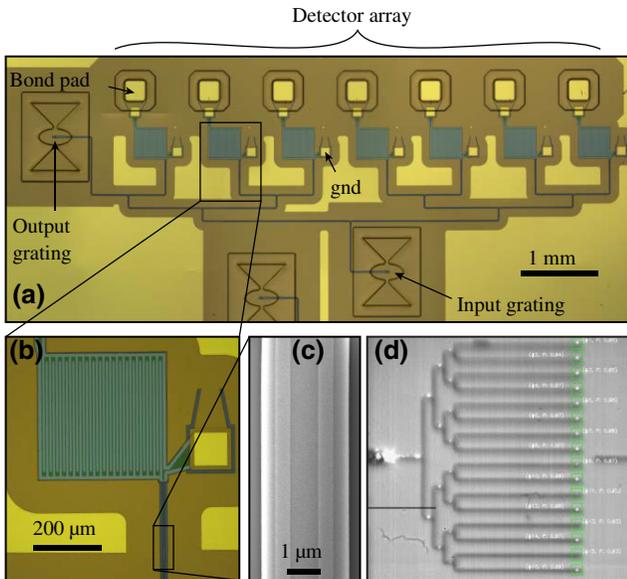

FIG. 1. (a) An optical-microscope image of the tree device. (b) An enlarged view of a waveguide-integrated detector with an associated meander inductor. (c) A SEM image of the nanowire on the silicon waveguide. (d) An infrared-camera image showing the passive tree device and the laser input. The green boxes are the automated detection of the output ports and the white text is automated output from the program, indicating the coordinates and intensity of the output.
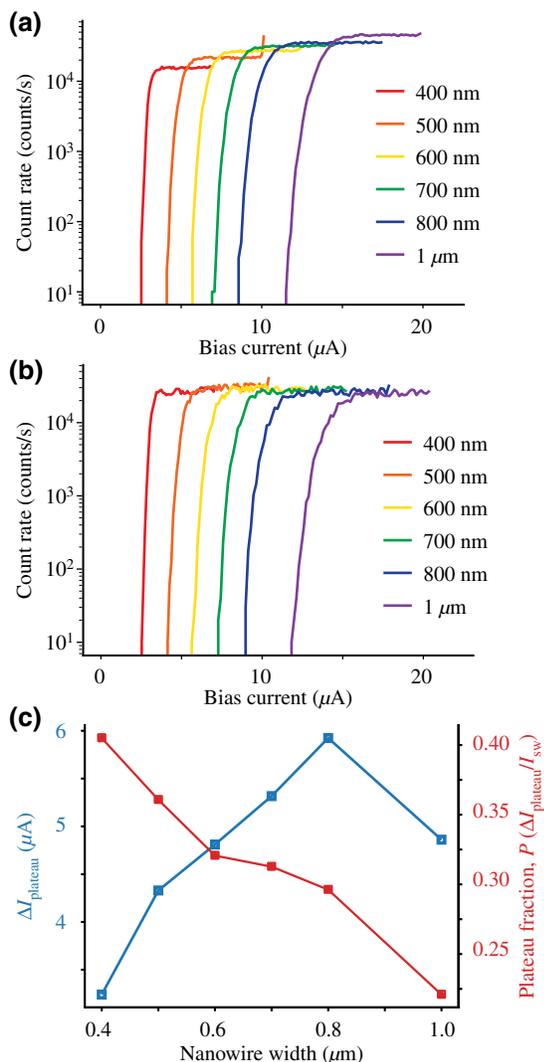
FIG. 2. (a) The count rate versus the bias current for six different SNSPD widths with uniform illumination of the detectors. (b) The count rate versus the bias current for six different SNSPD widths with waveguide-coupled light. (c) The width of the plateau versus the nanowire width.

For large-scale systems, the definition of yield should include this measure and therefore we do not include any SNSPDs with $P < 0.1$ in counting the number of working devices for the yield calculation. In fact, the 500-nm-wide standard detectors used in this work have a typical value of P much larger than this, with P between 0.3 and 0.4. We also include the length of the plateau region as an absolute current value on the right-hand axis of Fig. 2(c), as in experiments this may be more important than the fractional value P due to equipment and noise limitations.

The optimal performance depends on the application and the circuit used to read out the signal. Based on this analysis, any of the detector widths chosen would work for applications at 1.55 μm. The widest wires are the least sensitive to lithography but have the largest area and highest

energy consumption, while the reduced performance of the narrowest wires suggests that this is close to the resolution limit of our 365-nm *i*-line lithography process.

A second tree device with nanowires of different lengths is also fabricated and measured. The nanowire length is varied from 5 μm to 200 μm, based on simulations of the nanowire absorption length (see Sec. IV B). We find an absorption length of 160 dB/mm when measuring the detectors cryogenically at 1.55 μm (see Sec. IV F) and 175 dB/mm when measured at room temperature. The room-temperature measurement technique is automated and highly parallel and is described in Sec. IV E. A similar room-temperature versus cryogenic comparison has been performed previously for NbN SNSPDs on nanophotonic silicon and diamond waveguides [48,53] and room-temperature absorption measurements of similar devices have been performed for NbN on diamond and SiN waveguides [49,51]. In all previously published experiments, the detectors measured have each been on a separate integrated-photonic device, rather than being integrated in a single photonic circuit as in this work. The absorption values in this work are similar to those found in diamond (50–175 dB/mm for a single SNSPD hairpin) or silicon nitride waveguides. The absorption values are lower than the absorption coefficient (around 1000 dB/mm) found previously on silicon waveguides. This is due to (1) the thicker and wider silicon waveguides used in this work (220 nm thick instead of 110 nm thick) leading to a lower evanescent field, (2) the presence of the silicon nitride spacer between the waveguide and the nanowire and used as an etch stop, and (3) the thinner superconducting film used (2 nm thick WSi versus 3.5 nm thick NbN). These choices are made for ease of processing and would be straightforward to remedy in a more sophisticated foundry. The discussion of the room-temperature characterization of tree devices at 1.22 μm is in Sec. IV H.

B. High-dynamic-range detector arrays

The next device examined in this study is a high-dynamic-range detector array (HiDRA). These have previously been proposed and demonstrated in Ref. [31] but the devices did not work as predicted due to the presence of multiple waveguide modes. In this work, all waveguides are single mode. An optical-microscope image of the device is shown in Fig. 3(a). Each on-chip system consists of three HiDRAs with a single input grating and a single output grating used for fiber alignment during packaging. Light from the input grating is split evenly to four waveguides by three 50:50 beam splitters [see Fig. 3(b)]. Three of these four branches are attached to a HiDRA, while the fourth is a straight waveguide connected to an output grating. The HiDRA device works as follows. Input light is incident on a beam tap [see Fig. 3(c)]. The beam tap drops
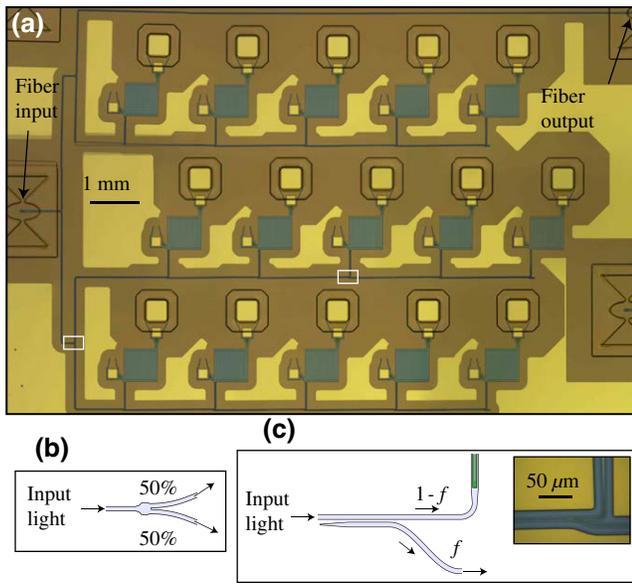
FIG. 3.   The HiDRA device used to test the dynamic range of the integrated SNSPD platform. (a) An optical-microscope image of the HiDRA device. (b) A schematic of a 50:50 beam splitter. (c) A schematic and an optical-microscope image of a beam tap.
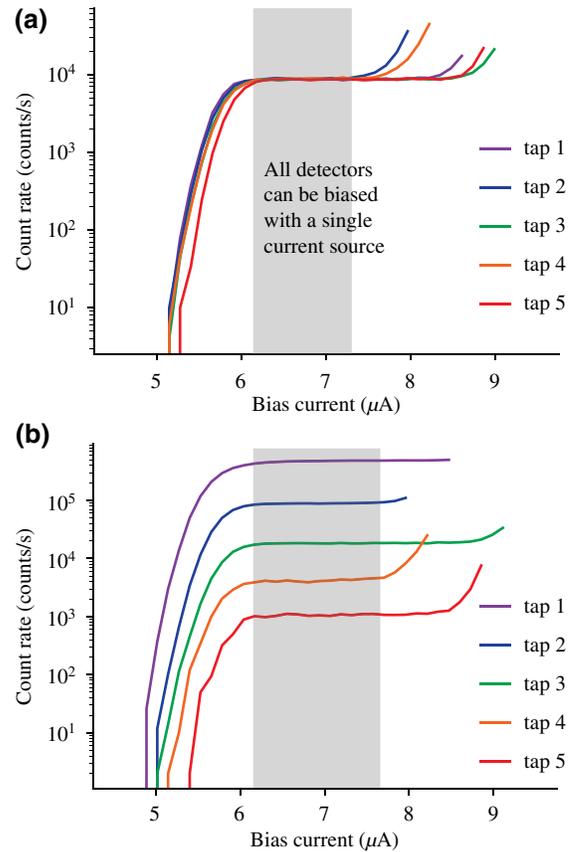

FIG. 4.   (a) The count rate versus the bias current for the five detectors in a HiDRA device when the chip is uniformly illuminated. (b) The count rate versus the bias current for the same device as in (a) with the input light waveguide coupled.

a fraction $1 - f$ of the light to a SNSPD, while a fraction f of the light is passed to the next beam tap.

Figure 4 shows data acquired from such HiDRAs. Figures 4(a) and 4(b) show the counts per second versus the bias current for the five detectors on a HiDRA designed to have a 10% beam tap ($f = 0.1$). Figure 4(a) shows the response of the detectors when the chip is uniformly illuminated. The detectors are designed to be identical and have very similar responses to this uniform illumination. Figure 4(b) shows the response when the light is waveguide-coupled, where successive detectors have significantly lower count rates. In each of Figs. 4(a) and 4(b), the gray-shaded region illustrates the range of bias currents over which the HiDRA can be biased with a single current source on all detectors (a guide to the eye). The range is truncated on the left by a decrease in detection efficiency and on the right by an increase in dark counts. The dark counts increase faster in the uniform-illumination case due to heating from the higher optical powers needed in this case. The ability to bias many SNSPDs with a single input is important when scaling to hundreds of SNSPDs on chip, as there will not be enough cryogenic I/O lines for a single individually tuned line per device, as is typically used in testing experiments. We extract tap fractions of 0.26 and 0.28 from the two HiDRAs designed to have taps with $f = 0.1$, with the mean value (solid black line) of 0.27. This deviation from the designed value is likely due to the gaps coming out narrower than designed. This, in combination with a slightly longer effective length due to the sine bends leading up to the beam tap, could explain

the discrepancy. A background light level around 30 dB below the level of the first waveguide-coupled detector on the connected HiDRA is also measured using disconnected devices. This indicates that with our current fabrication process and packaging scheme, there is a 30-dB noise floor. Previous measurements of LEDs coupled to waveguides have indicated a 40-dB noise floor [31]. Both the low fiber-coupling efficiency and the high waveguide-scattering loss contribute to this high background light level (for a precise characterization of these losses, see Sec. IV). As in the case of trees, "passive HiDRA" structures are fabricated for characterization at room temperature. Unlike the trees, in the case of the HiDRA structures measured at room temperature, the SNSPD is omitted and the structure is intended to characterize the photonic beam-tap ratios. The measurement technique used for characterization of the structures is discussed in Sec. IV E. The beam-tap ratio at room temperature is determined to be 0.29 for these devices, which agrees well with the cryogenic measurements described in the preceding paragraph. We also design the HiDRAs to have 1% and 0.1% beam taps; however, in

the fabricated devices the scattered background light prevents such low tap ratios from being useful. In Sec. IV H, we make use of the good agreement between cryogenic and room-temperature measurements for characterization at 1.22 μm. We test the detectors under uniform illumination only and then perform measurements of the passive integrated-photonic structures at 1.22 μm to avoid the lengthy packaging and cool-down processes. For the HiDRA, the noise floor is much higher for the 1.22 μm devices, at around 15 dB, likely due to a combination of a lower grating coupling efficiency (around 3 dB lower per coupler) and a higher propagation loss (14 dB/cm higher loss) (for details, see Sec. IV).

III. CONSIDERATIONS FOR SYNAPTIC OPERATION

In the superconducting optoelectronic neuromorphic platform proposed in Refs. [20] and [28], communication is photonic but computation is electronic. The signals communicated from each neuron to its synaptic connections are binary few-photon pulses, meaning that the information communicated is independent of the number of photons in the pulse. Only the timing of spikes is relevant. Synaptic weighting of the pulses is done entirely electronically. This communication scheme requires the detectors to be insensitive to the number of incident photons. We use the tree and HiDRA structures to test whether or not SNSPDs demonstrate this behavior by measuring their output electrical pulses when they are illuminated over a light-intensity range of close to 50 dB. At the low end of this range, multiphoton absorption events are very improbable, with less than 10^{-3} photons per pulse on average, while at the high end, multiphoton events are likely, with greater than 10 photons per pulse on average. The measurement is performed with an input laser with 50-ps pulses, operating at a pulse rate of 200 kHz and a wavelength of 1.57 μm. The photon number is changed via a calibrated variable-fiber attenuator. For coherent light, the photon number should follow a Poisson distribution given by $p(n > 0) = 1 - e^{-\langle n_{\mathrm{ph}} \rangle}$. Since the laser pulse length is much shorter than the response time of the detector, at most one count is observed every time the photon number in a pulse incident on the detector is greater than zero. Therefore, $p(n > 0)$ can be approximated by the count rate on the detector R_{det} divided by the repetition rate of the laser, R_{laser}. The average photon number in the pulse received on the detector should therefore be related to these quantities by $\langle n_{\mathrm{ph}} \rangle = -\log(1 - R_{\mathrm{det}}/R_{\mathrm{laser}})$. Since this is the mean photon number received on the detector, no assumptions are being made so far about the absolute value of the laser power. Next, we assume that the laser power is proportional to the mean photon number and therefore the plot of the count rate versus the attenuation can be used to extract the mean photon number incident upon the detector at each
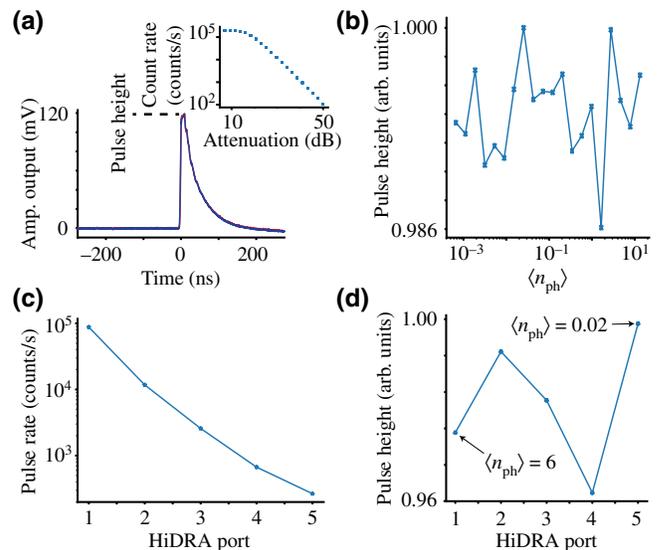

FIG. 5. (a) Pulses on a SNSPD with incident light with different average numbers of photons. The pulse height is indicated. The count rate versus the laser-attenuation plot used to calibrate the photon number is shown in the inset. (b) The pulse height versus the photon number for the detector shown in (a). (c) The count rate for each of the five HiDRA SNSPDs used in this experiment. (d) The pulse height for each of the five HiDRA SNSPDs with count rates indicated in (c). The average photon numbers for the first and last SNSPD are indicated.

set laser power. Figure 5(a) shows pulses from a SNSPD with waveguide-coupled light with mean photon numbers of 0.003 and 12. The difference is indistinguishable on this scale. The inset shows the count rate versus the laser attenuation for that SNSPD, from which the mean photon numbers are calculated. Figure 5(b) shows the pulse height versus the mean photon number for the detector in Figure 5(a). One hundred pulses are averaged to obtain the pulse height at each photon number. The pulse height varies by less than 2% over almost 5 orders of magnitude in incident photon flux, with no discernible dependence on the number of incident photons. Figures 5(c) and 5(d) show the results using the same laser configuration for detectors on a HiDRA device. Figure 5(c) shows the count rate versus the port number on the HiDRA, while Figure 5(d) shows the pulse height for each of the five detectors for the count rates shown in Figure 5(c). The pulse height varies by less than 5% over the HiDRA, despite using five different detectors with mean photon numbers varying from 0.02 to 6. The mean photon numbers are calculated for each detector in the HiDRA using the procedure described for the data in Figs. 5(a) and 5(b).

In the synaptic weighting scheme described in Ref. [56], the magnitudes of the SNSPD current pulses are used as the weights. This pulse amplitude can be controlled using the bias current applied to the SNSPD. In this case, it is very simple to generate synapses with a bit depth of 1

by simply biasing the SNSPD. However, further weighting can also be achieved by varying the bias current within the detection range of the synapse. It has also been shown that neural networks that perform simple tasks can be designed with these simple integrate-and-fire neurons [56]. In Fig. 6(a), pulses on a SNSPD for three different bias currents are shown. Figure 6(b) plots the pulse height versus the bias current for the detector over the available range of bias currents, indicating a range of 6–10 μA. The black line shows a linear fit to the curve. In Fig. 6(c), pulses from seven detectors on a symmetrical tree are shown. This is a tree device with all seven detectors fabricated with the same geometric parameters. The pulses are separated on the x axis (in time) for clarity—this separation has no physical meaning. Figure 6(d) shows the pulses on the same seven detectors with bias currents changed to give different pulse heights. In Ref. [56], the spiking neural network model assumes that bias currents can be varied between 5 μA and 15 μA (or set to zero). This is around twice the range that is observed here. It would be necessary to resimulate networks with these values to ensure that this is still a viable method for controlling the weights. Calculation of the bit depth in a quasianalog system is challenging, especially before the system is built and it is determined what constitutes the minimum reasonable change in weight

FIG. 7. A schematic of the fabrication process used in the paper, showing (a) a cross section and (b) a top view. (c) The geometry of the waveguide nanowire.

value. This is a challenge for many emerging technologies, such as magnetic Josephson junctions, which have > 20 000 internal states per μm^2 and are therefore quasianalog [57], and spin-torque nano-oscillators [58], where the frequencies are tuned using an applied current. Memristor technologies have similar issues, with typically reported bit depths of 4–6 [59] and with bit depths of up to 8 reported for phase-change memory [60]. However, even in these technologies, arguments have also been made for operating in a binary mode [61], for which these synapses are well suited. Superconducting-loop neurons [20] use SNSPDs to generate the electrical pulses and then Josephson junctions add current to an integrating loop in integer quantities of fluxons (although thermal noise will play a role and further investigation is required). The maximum number of fluxons that can be added determines the bit depth, with an upper limit of around 10 bits.

IV. METHODS

A. Fabrication

Waveguide-coupled SNSPDs are fabricated using a layer structure that builds toward the full superconducting-optoelectronic neuromorphic hardware fabrication process [30]. Certain layers, while not strictly necessary for this demonstration, are included for compatibility purposes. These include the small-pad metal layer, the waveguide hard mask (similar to the dielectric spacer used in the amplifier fabrication process) and its associated via layer
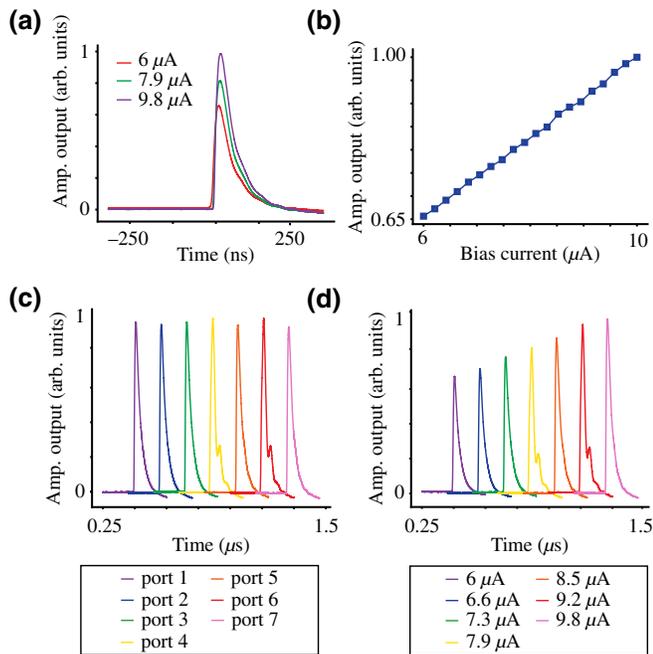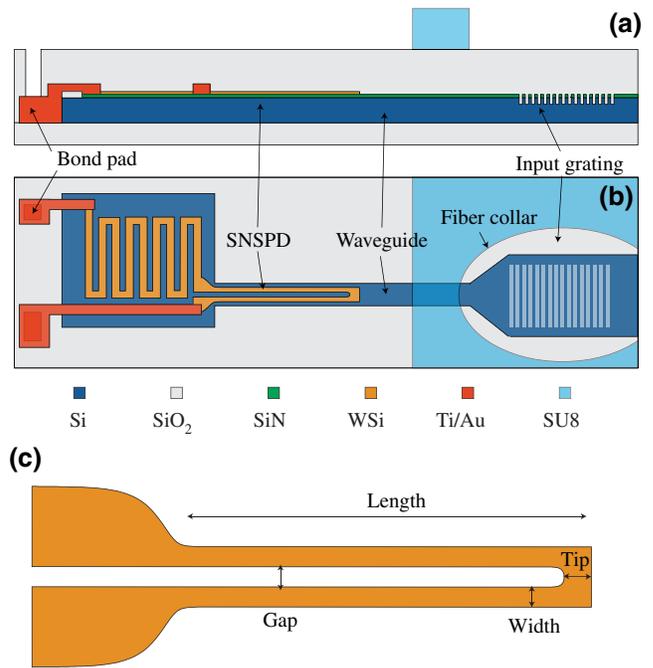

FIG. 6. (a) The pulses for three different values of the bias current on a single detector. (b) The pulse height versus the bias current for the detector in (a). (c) The pulses on each of seven detectors on an even tree device with a fixed bias current of 7.7 μA. (d) The pulses on the same seven detectors as in (c) for seven different values of the bias current. The x-axis separation of the pulses in (c) and (d) is for visibility purposes only.

(all described below). The SNSPDs are made from WSi [shown in orange in Fig. 7(a)], while the waveguides are silicon (blue). The substrate is a silicon-on-insulator (SOI) wafer with a 220-nm device layer and a 3-μm buried-oxide layer. A 40-nm SiN layer (green) is sputtered on the SOI to act as an etch stop for the WSi etch. Electron-beam and photolithography alignment marks are then etched into the wafer. A 2-nm/30-nm/2-nm Ti/Au/Ti layer is patterned via liftoff (red). This metal layer is used to make electrical contact to the SNSPDs.

A 2.1-nm $W_{0.64}Si_{0.36}$ film is then cosputtered as described in Ref. [41]. An amorphous silicon capping layer is then deposited without breaking vacuum to protect the thin WSi film. This WSi recipe has been demonstrated to allow saturated internal quantum efficiency for nanowires as wide as 1.2 μm at 1.55 μm [41]. The nanowires are patterned with photolithography using a 365-nm *i*-line stepper and etched using Ar and SF_6 chemistry.

A partial silicon etch is performed to create the grating used for coupling light from an optical fiber to the waveguides. This patterning is performed using e-beam lithography. The etch depth is designed to be 50 nm and is measured at 65–70 nm. A 100-nm SiO_2 hard mask is then deposited using plasma-enhanced chemical-vapor deposition (PECVD) and the waveguides are patterned and etched with e-beam lithography, using a positive-tone resist with a 2-μm clearance around the photonic devices. The remainder of the Si is then cleared out with photolithography. All of the silicon etches use SF_6 and C_4F_8 chemistry.

Following the waveguide etch, vias are etched through the waveguide oxide hard mask to make electrical contact with the Au layer, using CHF_3/O_2 chemistry. A wiring layer of 2 nm/440 nm/2 nm Ti/Au/Ti is then deposited and patterned with liftoff. A 1.8-μm cladding oxide (white) is then deposited using PECVD. Openings to the wirebond pads and the ground plane are then etched using the same CHF_3/O_2 dry etch. Finally, a 50-μm-thick SU8 packaging layer (light blue) is spun on and patterned using photolithography [62].

Figure 7 shows (a) cross-section and (b) top-view diagrams of a fabricated device, indicating some of the main features. Optical-microscope images and SEM images of fabricated devices are shown throughout this paper alongside measured data. The geometry of a typical SNSPD is shown in Fig. 7 (c). The geometric parameters are indicated in the diagram. In this work, the length and width are varied to determine the optimal device performance.

B. Photonic device design

The devices described in this paper require the design of a number of components. These devices include waveguides, beam splitters, beam taps, and SNSPD hairpins. The photonic devices are designed for operation at both

1.55 μm (standard telecommunications wavelength) and at 1.22 μm, where the proposed spiking neural networks [20, 27–29] are designed to operate. Simulations of the photonic waveguides are performed using a finite-difference frequency-domain eigenmode solver. The effective index of the waveguide modes in 220-nm-thick waveguides of different widths for wavelengths of 1.55 μm and 1.22 μm are shown in Fig. 8(a). Based on these simulations, the single-mode waveguide widths are chosen to be 450 nm (350 nm) for waveguides designed for operation at 1.55 μm (1.22 μm). For branching-tree and high-dynamic-range array structures, the beam taps are simulated using the eigenmode method described in Ref. [63]. The results of the simulations are shown in Fig. 8(b). For the 1.55 μm beam taps, the waveguide width is brought down to 400 nm at the beam-tap region. The beam-splitter gap widths and interaction lengths are chosen based on the plots in Fig. 8(b). The 1.55 μm beam splitters are chosen based on the optimized design in Ref. [64]. The 1.22 μm beam splitters are multimode interferometers that are optimized in separate fabrication runs to minimize loss.
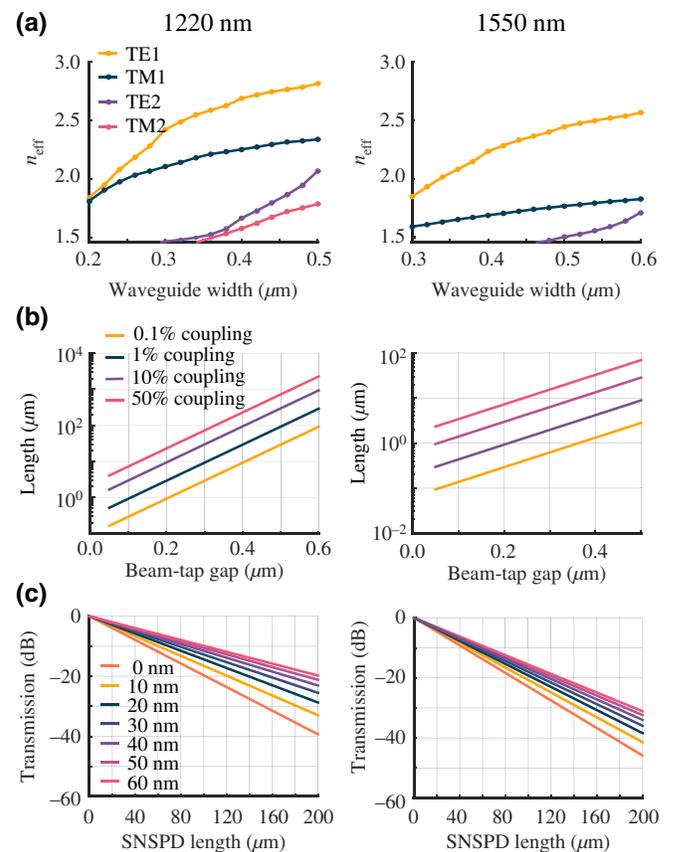

FIG. 8. (a) Simulated waveguide modes at 1.22 μm and 1.55 μm. (b) The simulated beam-tap length versus the gap for different values of the splitting ratio at 1.22 μm and 1.55 μm. (c) The simulated transmission versus the nanowire length for different values of the SiN-spacer thickness.

The waveguide-integrated nanowire geometry is shown in Fig. 7(c). We refer to this structure as a hairpin. Unless otherwise noted, the measured width of the nanowire is 500 nm and the length is 200 μm. There is an 800-nm gap between the wires in the hairpin, with a 400-nm gap on either side. The waveguides are adiabatically tapered from their single-mode widths up to 2.6 μm over a 100-μm length to accommodate the SNSPD. The detectors also have a meandering section of 3-μm-wide wire for extra inductance, to achieve a total value of 1.25 μH.

The nanowire absorption is calculated for a 500-nm-wide nanowire with the geometry shown in Fig. 7(c). The real and imaginary parts of the refractive index of the WSi are measured and used to simulate the real and imaginary propagation constants in the nanowire on the waveguide structure. The resulting absorption versus the length at 1.22 μm and 1.55 μm are shown in Fig. 8(c) for different thicknesses of SiN spacer. Based on the results of this plot in combination with concerns about ensuring electrical isolation of the superconducting layer, the SiN spacer is designed to have a 40-nm thickness.

C. Component characterization

For every aspect of the integrated devices described in Sec. II, test structures are also fabricated to characterize the performance and losses. These test structures are described below and in Fig. 9. All of these test devices are fabricated on the same wafer as the devices with integrated nanowires. The insertion losses of the grating couplers vary from 5 to 7 dB at 1.55 μm and from 8 to 10 dB at 1.22 μm.

The transmission of meandering waveguides of various lengths is used to measure the propagation loss of the waveguides (the cutback method). The transmission as a function of the waveguide length is shown in Fig. 9(a). Each device has an input, a reference port, and an output grating, as shown in the optical-microscope image in the inset. The intensity of the light from the reference and output gratings is measured using an infrared camera and compared for devices of different lengths to obtain the data in Fig. 9(a). The measured losses are 24 dB/cm at 1.22 μm and 10.2 dB/cm at 1.55 μm. These losses can be improved using well-known processes [65].

The loss per beam splitter is measured using a similar technique to the waveguide propagation loss. The transmission through sequences of various numbers of beam splitters is measured. The data are shown in Fig. 9(b) and the insets show a SEM image of a single beam splitter as well as a microscope image of a series of four beam splitters. For this measurement, light is collected through a second fiber and detected with a photodiode. The loss of 0.7 dB per beam splitter is around double the insertion loss simulated [63]. To quantify the power splitting ratio of the beam splitters, they are embedded in unbalanced
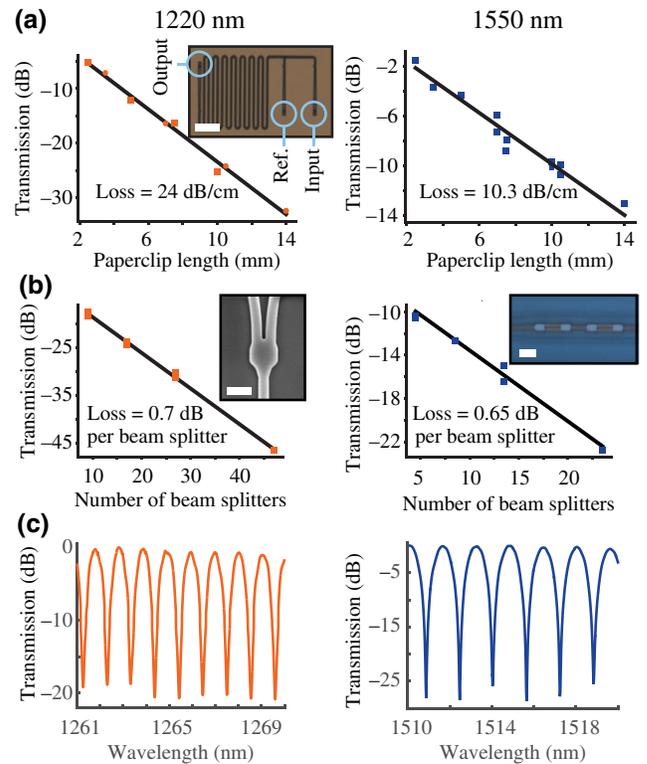

FIG. 9. (a) The measured waveguide loss at 1.22 μm and 1.55 μm with an inset optical-microscope image of the structure used for loss measurement (scale bar = 200 μm). (b) The measured beam-splitter loss at 1.22 μm and 1.55 μm, with the left inset showing a SEM image of a beam splitter (scale bar = 1 μm) and the right inset showing an optical-microscope image of the structure used for measuring the beam-splitter loss (scale bar = 5 μm). (c) The measured transmission spectrum of the Mach-Zehnder interferometer at 1.22 μm and 1.55 μm.

Mach-Zehnder interferometers (MZIs). The extinction of such an interferometer can be used to bound the splitting ratio [63]. Transmission spectra of the MZIs at 1.22 μm and 1.55 μm are shown in Fig. 9(c). The extinction at 1.55 μm is nearly 30 dB, while at 1260 nm (measured at this wavelength due to the availability of an O-band laser), it is near 20 dB. From these values, we obtain approximate splitting ratios of 48:52 at 1.55 μm and 45%:55% at 1260 nm. The extinction values measured with the MZIs are similar to the noise floor observed with SNSPDs at those wavelengths (see Secs. II A and II B). This result may indicate that the beam-splitting ratios are closer to 50%:50% than calculated from these measurements.

D. Cryogenic measurements

Room-temperature and cryogenic measurements are carried out on devices comprising both integrated-photonic and superconducting-nanowire components. For cryogenic measurements, detector samples are wire bonded and fiber

packaged by the method described in Ref. [62]. The packaged samples are then cooled to 800 mK in a closed-cycle sorption pump ^{4}He cryostat. The detectors are dc biased and are measured through ac-coupled amplifiers. Light can either be coupled onto the chip via the packaged fibers or via a separate fiber that floods the entire chip uniformly with light.

For each cryogenic measurement, a corresponding room-temperature measurement is also performed on a similar device, with output gratings in place of single-photon detectors. For these room-temperature measurements, we use a technique that has previously been described for the characterization of photonic structures in Ref. [66]. With this technique, light is coupled from an optical fiber to an input grating and the device is characterized by measuring the intensity emitted from output gratings at different points in the device on an infrared camera. The room-temperature measurements are directly comparable to the cryogenic measurements.

E. High-dynamic-range room-temperature camera measurements

Parallel automated measurements of the room-temperature properties of passive photonic devices are carried out using a procedure similar to that described in Ref. [66]. In brief, light is input through a single input grating. The light from the output gratings following the nanowires and reference paths is captured on an infrared-sensitive camera. Automated image processing is used to find the output grating locations and to extract the relative intensities of light from these outputs. The green boxes in Fig. 1(d) indicate the locations that have been automatically assigned as the outputs. The intensity of the light after each nanowire is compared to the intensity from the associated reference grating to normalize out the effect of the grating response as well as the propagation losses. The structures measured on this wafer have a high dynamic range, with up to 30 dB attenuation measured in the devices themselves. This poses an issue for the 8-bit camera used for characterization of the structures. We therefore use a calibrated attenuator to vary the input light and capture the light from the reference and measurement gratings over 50 dB of attenuation.

F. Tree analysis

In Fig. 10(a), the plateau height (the averaged value of the data points within the plateau region) is plotted versus the wire width for the tree with nanowires of different widths. The blue squares show the result when the chip is uniformly illuminated. The black line is a linear fit, which is expected due to the fact that the detector area is proportional to the wire width. The red circles show the same analysis for the waveguide-coupled data. Since all of the nanowires absorb more than 99% of the light, there should
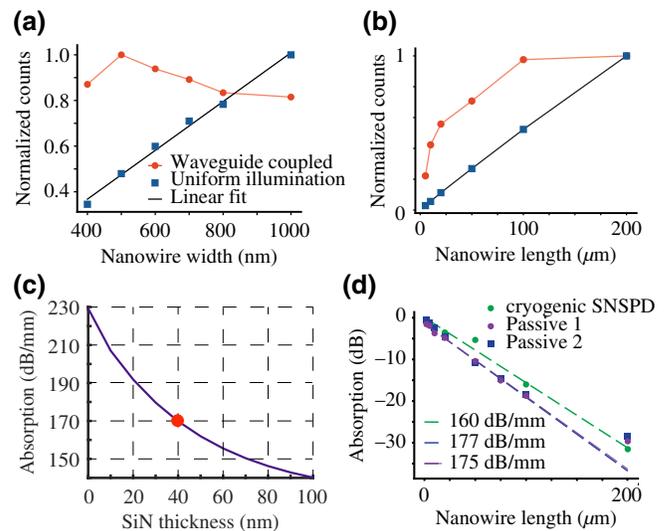

FIG. 10. (a) The height of the plateau versus the nanowire width for uniform illumination (blue squares) and waveguide-coupled light (red circles). The count rate is linearly proportional to the nanowire width when the chip is uniformly illuminated, as shown by the linear fit (black line). (b) The count rate versus the nanowire length for six different SNSPD lengths for waveguide-coupled and not-waveguide-coupled light. The count rate is linearly proportional to the SNSPD length when the detectors are uniformly illuminated, while the waveguide-coupled data saturates. (c) The simulation data for the absorption coefficient versus the SiN-spacer thickness, extracted from the data in Fig. 8(c). The red dot indicates the fabricated value. (d) The absorption versus the nanowire length for cryogenic and passive devices measured at room temperature, showing absorption between 160 dB/mm and 175 dB/mm.

be no discernible difference between the values for the different wire widths. However, the data indicate a slight drop in efficiency for wider wires. The cause of this drop in efficiency may be the tip of the wire, which is approximately 4 times wider than the rest, absorbing some of the light without leading to detection events. This design is intended to avoid current crowding [67]. For reference, a 4-μm length corresponds to 14% absorption, according to the measured absorption values [simulated in Fig. 8 and measured in Fig. 10(d)]. It may be possible to narrow this tip length, although the cost may be a lower switching current and a narrower plateau region. Interestingly, the narrowest wire is also less efficient, possibly due to edge roughness resulting from pushing the limits of the photolithography tool used in fabrication.

Figure 10(b) shows the height of the nanowire plateau versus the nanowire length for the different-length nanowires. Normalization is again performed against the highest-count-rate detector on that particular tree device. A linear dependence versus the length (proportional to area) is again observed for the case of uniform illumination. The same nanowires show a saturating behavior with

increasing length when the light is waveguide coupled, indicating that all light is absorbed in the longest wires. We use this tree to quantify the absorption of the nanowires versus the length. The expected absorption coefficient can be extracted from the simulations shown in Fig. 8(c). A plot of the simulated absorption coefficient versus the SiN-spacer thickness is shown in Fig. 10(c), with the value expected for the fabricated SiN thickness indicated by the red dot. Figure 10(d) shows $10 \cdot \log(A - R_{\text{det}})$, where R_{det} is the SNSPD count rate, versus the length for the value of A that gives the best linear fit (the green dashed line). The nanowire absorption per unit length calculated from the slope of this line is 160 dB/mm.

We also measure the waveguide-integrated nanowire absorption at room temperature. Separate photonic trees are fabricated for this purpose. Such a tree is shown in Fig. 1(d). In these structures, output grating couplers are fabricated after each nanowire. Each waveguide-integrated nanowire is also accompanied by a reference waveguide with the same geometry but without a nanowire. Light is coupled through a fiber to the input grating. A branching tree splits the light evenly along 16 paths in this structure, with eight nanowires each accompanied by a reference waveguide, for a total of 16 output ports. The automated measurement technique is described in Sec. IV E.

The results of this measurement for the tree structure with varying nanowire length are compared to cryogenic measurements of counts per second in Fig. 10(d) (purple circles and blue squares, indicating measurements of two equivalent devices). The absorption in this case is directly measured (unlike the case of the cryogenic measurement, where the total light to be absorbed is a fit parameter). We find 175 dB/mm attenuation (purple and blue lines) from the SNSPD hairpin at room temperature, compared with the cryogenic measurement of 160 dB/mm. This level of agreement may be satisfactory for some applications, in which case it would be sufficient to measure the absorption of prototype nanowires with room-temperature measurements. The simulations shown in Sec. IV give a value of 165 dB/mm for the structures studied in the experiment, which is in good agreement with the experimental results.

G. HiDRA analysis

In the HiDRA device, after n beam taps, a SNSPD should receive $(1 - f)f^{(n-1)}$ of the light. The waveguide loss must also be included and is approximated by $e^{2\alpha \cdot l(n-1)}$, where $\alpha = (\ln 10/20)L$, L is the waveguide loss in dB/m, and l is the length of each segment of the HiDRA. The total light intensity at the nth detector is therefore $(1 - f)f^{(n-1)}e^{2\alpha \cdot l(n-1)}$.

Figure 11(a) shows further analysis of the devices designed to have a 10% beam tap. The log of the detector count rate versus the detector tap number n is a straight line, indicating that the tap fraction is constant. The tap
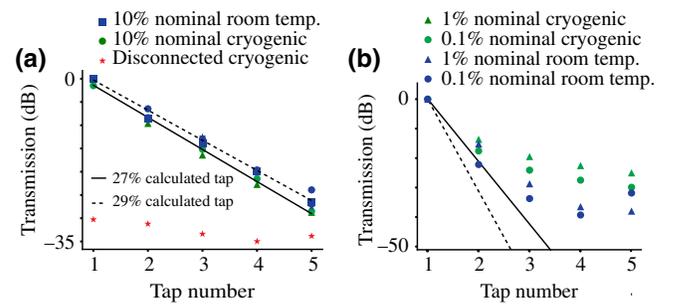

FIG. 11. (a) The count rate versus the port number for a HiDRA designed to have a 10% beam tap, including measurements of two cryogenic detector devices (green markers) and two room-temperature devices (blue markers). The black lines show linear fits to the cryogenic (solid) and room-temperature (dashed) data. The red stars show the result for a cryogenic device with intentionally disconnected detectors. (b) The count rate versus the port number for a HiDRA with differently designed beam-tap ratios for both cryogenic (green makers) and room-temperature (blue markers) operation. The scattered light prevents the dynamic range from exceeding 30 dB. The black lines show how a 1% (solid) and 0.1% (dashed) measured device should perform, indicating that the nominal values are far from the realized values.

fraction can be calculated from the slope m of this plot, as well as the known segment length l and waveguide loss L in dB/m. The tap fraction f is then given by $f = 10^{(m-Ll)/10}$. We extract tap fractions of 0.26 and 0.28 from the two HiDRAs designed to have taps with $f = 0.1$, with a mean value (the solid black line) of 0.27. This deviation from the designed value is likely due to the gaps coming out narrower than designed. This, in combination with a slightly longer effective length due to the sine bends leading up to the beam tap, could explain the discrepancy. A third device with intentionally disconnected waveguides is also measured (the red stars). The disconnected HiDRAs are identical aside from a 180° waveguide turn and a taper that couples light into free space, away from the SNSPD. The disconnected device shows a relatively constant level at around 30 dB below the level of the first waveguide-coupled detector on the connected HiDRA. This indicates that with our current fabrication process and packaging scheme, there is a 30-dB noise floor. Previous measurements of LEDs coupled to waveguides have indicated a 40-dB noise floor [31]. Both the low fiber-coupling efficiency and the high waveguide-scattering loss contribute to this high background light level.

Figure 11(a) also shows the transmission for each of the output ports of a passive room-temperature HiDRA. From the plot of the transmission versus the attenuation, the beam-tap ratio at room temperature is determined to be 0.29 for these devices, calculated from the slope of the dashed black line in the figure. This value agrees well with the cryogenic measurements described in the preceding paragraph.
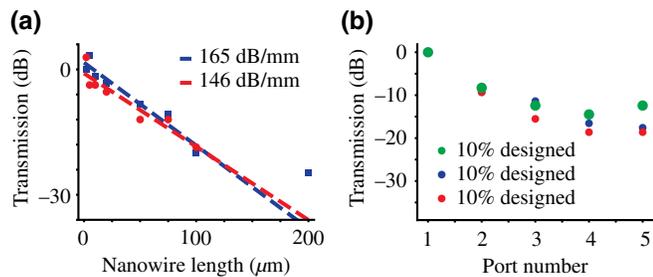
(a)

(b)

FIG. 12. Measurements of (a) length trees and (b) HiDRA at room temperature at 1.22 μm.

The results of the cryogenic and room-temperature measurements of HiDRAs designed to have 1% and 0.1% beam taps are shown in Fig. 4(d). The solid black line indicates where the data for a cryogenic HiDRA with a 1% tap ratio should fall, while the dashed black line indicates the same for a 0.1%-tap-ratio device. The data clearly do not follow this trend. The scattered background light prevents such low tap ratios from being useful. The passive HiDRAs with beam-tap ratios designed to be 1% (blue triangles) and 0.1% (blue circles) indicate similar trends. We again find very good agreement with the cryogenic measurement, even observing the same roll-off at a 30-dB noise floor.

H. Waveguide-integrated nanowire measurements at 1.22 μm

The values of the absorption and the dynamic range measured with the passive tree and HiDRA devices measured at 1.55 μm match those of the active devices and therefore we measure the 1.22-μm devices at room temperature without also performing cryogenic measurements. The measurement technique is the same as that described in the room-temperature-measurements portions of Secs. II A and II B of the main text, but with a 1.22-μm input laser. We find that the measured absorption of the nanowires at 1.22 μm is similar to that measured at 1.55 μm, giving 146 dB/mm and 165 dB/mm for two different (nominally identical) devices, as shown in Fig. 12(a). The large variation may be due to the fact that there is more scattering from the 1.22-μm waveguides due to edge roughness, which affects the accuracy of the camera intensity measurements.

The HiDRA devices designed to have 10% beam taps suffer from a very high background, with the intensity flattening out at around 20 dB, as shown in Fig. 12(b). This reduced extinction ratio is likely due to the very high waveguide scattering loss.

V. DISCUSSION

Realization of the hardware necessary for several forms of advanced cryogenic computing involves integrating devices in a scalable fabrication process. One of the main elements required for quantum and neural computing

hardware involving photonic communication is high-yield high-efficiency detectors. The work presented here introduces several tools leveraging integrated-photonic circuits for the characterization of waveguide-integrated SNSPDs and makes several measurements that are important for synaptic operation.

The integration of the detectors with photonic structures provides a powerful toolbox to rapidly assess large numbers of detectors with various design parameters in a context similar to that in which the detectors will be utilized. We describe the fabrication of waveguide-coupled SNSPDs that can be used as components in synapses in superconducting optoelectronic networks. We demonstrate several integrated-photonic structures that can be used to assess important properties of waveguide-integrated SNSPDs using a single-fiber input and we compare room-temperature measurements with cryogenic measurements. We demonstrate the effects of the photon number and the bias current on the SNSPD pulse height, these being relevant to synaptic operation. Such measurement methods leveraging integrated-photonic circuits are likely to be indispensable in the maturation of advanced computing technologies using photons for communication.

A primary focus of this work is on realizing SNSPDs that have high yield, enabling many such detectors to be integrated in large computational systems. To enable this, we choose to work with wide wires fabricated with photolithography rather than narrower wires commonly fabricated with e-beam lithography. In total throughout this paper, we test 49 waveguide-coupled SNSPDs, of which all 49 yield plateaus with $P > 0.1$. These wider wires consume more area, which is problematic in large integrated systems. However, we find that elimination of even a single e-beam lithography step greatly simplifies fabrication. This compromise will enable near-term demonstrations in academic and government-laboratory cleanrooms, which is important for the scientific goal of building and testing superconducting optoelectronic neural networks. By simplifying the fabrication of the detectors to be compatible with our 365-nm photolithography equipment, we are able to fabricate large numbers of detectors with saturated internal efficiencies in a single device, something that has proved challenging in other materials [33,36,42], in particular when integrated in complex photonic circuits [35]. Ultimately, if such detectors prove valuable to advanced neural or quantum computing, the improved fabrication and readily achievable feature sizes available at a commercial foundry will only improve the performance demonstrated here, while reducing the area of each detector considerably. When fabricated in a commercial foundry, it may be beneficial to optimize processing for an entirely different material with a much higher T_c. This would allow the cryogenic complexity to be reduced by operating the entire platform above 4 K. Nevertheless, the techniques described in this work will be necessary

for characterization of large numbers of SNSPDs in any material.

Since the inductance is set so that the L/R time constant is sufficiently long to avoid latching, in this work large inductance meanders are added in series with the detection hairpins to bring the total value of the inductances L above 1 μH and achieve the correct L/R. This leads to a large increase of the area of the devices. However, in a practical neuronal circuit, the inductance may be very different. In these measurements, SNSPD pulses are read out across a 50-Ω transmission line. In the context of synapses, the detectors will be in parallel with a much smaller resistance and the inductance (and therefore the size) can be reduced while maintaining the same L/R time constant for resetting. If the synapses of Ref. [56] are employed, the resistance affecting the SNSPD L/R time constant is related to the integration time of the neuron. If the synapses of Refs. [20,28] are employed, the normal-state resistance of the Josephson junctions (on the order of 1 Ω) establishes a minimum functional value of SNSPD inductance, on the order of 100 nH, or around a factor of 10 smaller than in the present work.

The experiments presented here point to important challenges for the scaling of future technologies. The HiDRA device shows that in the present devices, scattered light is significant, at only 30 dB below the measured signal. This finding is particularly relevant for photonic quantum computing, in which relatively strong pump lasers are used to generate entangled-photon pairs on chip, as well as for superconducting optoelectronic networks used for spiking neural systems, in which many light sources must be integrated on a chip with many single-photon detectors. The scattered light from a neuron must be less than the waveguide-coupled light pulses. Scattered light therefore puts an upper limit on the number of connections possible per neuron. This demonstrates that for most large-scale photonic systems incorporating on-chip detectors, some form of detector shielding will likely be necessary.

Finally, the synaptic measurements made in this paper confirm the utility of SNSPDs for the synaptic application, where the physics of the superconducting-to-normal phase transition is used as the binary signaling mechanism for neuromorphic computing. Future work will build on this to demonstrate complex synaptic circuits integrated with transmitters for integrate-and-fire functionality.

ACKNOWLEDGMENTS

[1] Carver Mead, *Analog VLSI and Neural Systems* (Addison-Wesley, Boston, MA, USA, 1989).

[2] R. J. Vogelstein, U. Mallik, J. T. Vogelstein, and G. Cauwenberghs, Dynamically reconfigurable silicon array of spiking neurons with conductance-based synapses, IEEE Trans. Neural Networks **18**, 253 (2007).

[3] G. Indiveri, B. Linares-Barranco, T. J. Hamilton, A. van Schaik, R. Etienne-Cummings, T. Delbruck, S.-C. Liu, P. Dudek, P. Häfliger, S. Renaud, J. Schemmel, G. Cauwenberghs, J. Arthur, K. Hynna, F. Folowosele, S. Saighi, T. Serrano-Gotarredona, J. Wijekoon, Y. Wang, and K. Boahen, Neuromorphic silicon neuron circuits, Front. Neurosci. **73**, 73 (2011).

[4] J. M. Cruz-Albrecht, M. W. Yung, and N. Srinivasa, Energy-efficient neuron, synapse, and STDP integrated circuits, IEEE Trans. Biomed. Cir. Sys. **6**, 246 (2012).

[5] T. Pfeil, A. Grubl, and K. Meier, Six networks on a universal neuromorphic computing substrate, Front. Neurosci. **7**, 11 (2013).

[6] A. S. Cassidy, J. Georgiou, and A. G. Andreou, Design of silicon brains in the nano-CMOS era: Spiking neurons, learning synapses and neural architecture optimization, Neural Networks **45**, 4 (2013).

[7] B. V. Benjamin, P. Gao, E. McQuinn, S. Choudhary, A. R. Chandresekaran, J.-M. Bussat, R. Alvarez-Icaza, J. V. Arthur, P. A. Merolla, and K. Boahen, Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations, Proc. IEEE **102**, 699 (2014).

[8] Elisabetta Chicca, Fabio Stefanini, Chiara Bartolozzi, and Giacomo Indiveri, Neuromorphic electronic circuits for building autonomous cognitive systems, Proc. IEEE **102**, 1367 (2014).

[9] Steve Furber, Large-scale neuromorphic computing systems, J. Neural Eng. **13**, 051001 (2016).

[10] Y. Harada and E. Goto, Artificial neural network circuits with Josephson devices, IEEE Trans. Magn. **27**, 2863 (1991).

[11] M. Hidaka and L. A. Akers, An artificial neural cell implemented with superconducting circuits, Supercond. Sci. Technol. **4**, 654 (1991).

[12] T. Kondo, M. Kobori, T. Onomi, and K. Nakajima, Design and implementation of stochastic neurosystem using SFQ logic circuits, IEEE Trans. Appl. Supercond. **15**, 320 (2005).

[13] T. Hirose, T. Asai, and Y. Amemiya, Spiking neuron devices consisting of single-flux-quantum circuits, Phys. C **445**, 1020 (2006).

[14] T. Hirose, T. Asai, and Y. Amemiya, Pulsed neural networks consisting of single-flux-quantum spiking neurons, Phys. C **463**, 1072 (2007).

[15] P. Crotty, D. Schult, and K. Segall, Josephson junction simulation of neurons, Phys. Rev. E **82**, 011914 (2010).

[16] K. Segall, M. LeGro, S. Kaplan, O. Svitelskiy, S. Khadka, P. Crotty, and D. Schult, Synchronization dynamics on the picosecond time scale in coupled Josephson junction networks, Phys. Rev. E **95**, 032220 (2017).

[17] M. L. Schneider, C. A. Donnelly, S. E. Russek, B. Baek, M. R. Pufall, P. F. Hopkins, P. Dresselhaus, S. P. Benz, and W. H. Rippard, Ultralow power artificial synapses using nanotextured magnetic Josephson junctions, Sci. Adv. **4**, 1701329 (2018).

[18] Y. S. Abu-Mostafa and D. Psaltis, Optical neural computers, Sci. Am. **256**, 88 (1987).

[19] Thomas Ferreira De Lima, Bhavin J. Shastri, Alexander N. Tait, Mitchell A. Nahmias, and Paul R. Prucnal, Progress in neuromorphic photonics, Nanophotonics **6,** 577 (2017).

[20] J. M. Shainline, S. M. Buckley, A. N. McCaughan, J. Chiles, M. Castellanos-Beltran, C. A. Donnelly, M. L. Schneider, A. Jafari-Salim, R. P. Mirin, and S. W. Nam, Superconducting optoelectronic loop neurons, J. Appl. Phys. **126,** 044902 (2019).

[21] Johnny Moughames, Xavier Porte, Michael Thiel, Gwenn Ulliac, Laurent Larger, Maxime Jacquot, Muamer Kadic, and Daniel Brunner, Three-dimensional waveguide interconnects for scalable integration of photonic neural networks, Optica **7,** 640 (2020).

[22] D. Psaltis and N. Farhat, Optical information processing based on an associative-memory model of neural nets with thresholding and feedback, Opt. Lett. **10,** 98 (1985).

[23] H. John Caulfield, Jason Kinser, and Steven K. Rogers, Optical neural networks, Proc. IEEE **77,** 1573 (1989).

[24] Alexander N. Tait, Student Member, Mitchell A. Nahmias, Bhavin J. Shastri, and Paul R. Prucnal, Broadcast and Weight: An Integrated Network for Scalable Photonic Spike Processing, J. Light. Technol. **32,** 3427 (2014).

[25] Yichen Shen, Nicholas C. Harris, Scott Skirlo, Mihika Prabhu, Tom Baehr-Jones, Michael Hochberg, Xin Sun, Shijie Zhao, Hugo Larochelle, Dirk Englund, and Marin Soljacic, Deep learning with coherent nanophotonic circuits, Nat. Photonics **11,** 441 (2017).

[26] J. Feldmann, N. Youngblood, C. D. Wright, H. Bhaskaran, and W. H. P. Pernice, All-optical spiking neurosynaptic networks with self-learning capabilities, Nature **569,** 208 (2019).

[27] Jeffrey M. Shainline, Sonia M. Buckley, Richard P. Mirin, and Sae Woo Nam, Superconducting Optoelectronic Circuits for Neuromorphic Computing, Phys. Rev. Appl. **7,** 034013 (2017).

[28] Jeffrey M. Shainline, Sonia M. Buckley, Adam N. McCaughan, Jeff Chiles, Amir Jafari-Salim, Richard P. Mirin, and Sae Woo Nam, Circuit designs for superconducting optoelectronic loop neurons, J. Appl. Phys. **124,** 152130 (2018).

[29] J. M. Shainline, Fluxonic processing of photonic synapse events, IEEE J. Sel. Top. Quant. Electron. **26,** 7700315 (2020).

[30] Sonia M. Buckley, Adam McCaughan, Alex N. Tait, Jeff Chiles, Sae Woo Nam, Richard P. Mirin, and Jeffrey M. Shainline, Progress in superconducting opto-electronic neuromorphic networks (to be published).

[31] Sonia Buckley, Jeffrey Chiles, Adam N. McCaughan, Galan Moody, Kevin L. Silverman, Martin J. Stevens, Richard P. Mirin, Sae Woo Nam, and Jeffrey M. Shainline, All-silicon light-emitting diodes waveguide-integrated with superconducting single-photon detectors, Appl. Phys. Lett. **111,** 141101 (2017).

[32] A. N. McCaughan, V. B. Verma, S. M. Buckley, J. P. Allmaras, A. G. Kozorezov, A. N. Tait, S. W. Nam, and J. M. Shainline, A superconducting thermal switch with ultrahigh impedance for interfacing superconductors to semiconductors, Nat. Electron. **2,** 451 (2019).

[33] Christof P. Dietrich, Andrea Fiore, Mark G. Thompson, Martin Kamp, and Sven Höfling, GaAs integrated quantum

photonics: Towards compact and multi-functional quantum photonic integrated circuits, Laser Photon. Rev. **10,** 857 (2016).

[34] Simone Ferrari, Carsten Schuck, and Wolfram Pernice, Waveguide-integrated superconducting nanowire single-photon detectors, Nanophotonics **7,** 1725 (2018).

[35] Ali W. Elshaari, Wolfram Pernice, Kartik Srinivasan, Oliver Benson, and Val Zwiller, Hybrid integrated quantum photonic circuits, Nat. Photonics **14,** 285 (2020).

[36] Faraz Najafi, Jacob Mower, Nicholas C. Harris, Francesco Bellei, Andrew Dane, Catherine Lee, Xiaolong Hu, Prashanta Kharel, Francesco Marsili, Solomon Assefa, Karl K. Berggren, and Dirk Englund, On-chip detection of non-classical light by scalable integration of single-photon detectors, Nat. Commun. **6,** 5873 (2015).

[37] F. Marsili, V. B. Verma, J. A. Stern, S. Harrington, A. E. Lita, T. Gerrits, I. Vayshenker, B. Baek, M. D. Shaw, R. P. Mirin, and S. W. Nam, Detecting single infrared photons with 93% system efficiency, Nat. Photonics **7,** 210 (2013).

[38] V. B. Verma, B. Korzh, F. Bussières, R. D. Horansky, S. D. Dyer, A. E. Lita, I. Vayshenker, F. Marsili, M. D. Shaw, H. Zbinden, R. P. Mirin, and S. W. Nam, High-efficiency superconducting nanowire single-photon detectors fabricated from MoSi thin-films, Opt. Express **23,** 33792 (2015).

[39] Emma E. Wollman, Varun B. Verma, Adriana E. Lita, William H. Farr, Matthew D. Shaw, Richard P. Mirin, and Sae Woo Nam, A kilopixel array of superconducting nanowire single-photon detectors, Opt. Express **27,** 35279 (2019).

[40] Y. P. Korneeva, D. Y. Vodolazov, A. V. Semenov, I. N. Florya, N. Simonov, E. Baeva, A. A. Korneev, G. N. Goltsman, and T. M. Klapwijk, Optical Single-Photon Detection in Micrometer-Scale NbN Bridges, Phys. Rev. Appl. **9,** 064037 (2018).

[41] J. Chiles, S. M. Buckley, A. Lita, V. B. Verma, J. Allmaras, B. Korzh, M. D. Shaw, J. M. Shainline, R. P. Mirin, and S. W. Nam, Superconducting microwire detectors based on WSi with single-photon sensitivity in the near-infrared, Appl. Phys. Lett. **116,** 242602 (2020).

[42] Ronan Gourgues, Iman Esmaeil Zadeh, Ali W. Elshaari, Gabriele Bulgarini, Johannes W. N. Los, Julien Zichi, Dan Dalacu, Philip J. Poole, Sander N. Dorenbos, and Val Zwiller, Controlled integration of selected detectors and emitters in photonic integrated circuits, Opt. Express **27,** 3710 (2019).

[43] C. Schuck, X. Guo, L. Fan, X. Ma, M. Poot, and H. X. Tang, Quantum interference in heterogeneous superconducting-photonic circuits on a silicon chip, Nat. Commun. **7,** 10352 (2016).

[44] Svetlana Khasminskaya, Felix Pyatkov, Karolina Słowik, Simone Ferrari, Oliver Kahl, Vadim Kovalyuk, Patrik Rath, Andreas Vetter, Frank Hennrich, Manfred M. Kappes, G. Gol'tsman, A. Korneev, Carsten Rockstuhl, Ralph Krupke, and Wolfram H. P. Pernice, Fully integrated quantum photonic circuit with an electrically driven light source, Nat. Photonics **10,** 727 (2016).

[45] Oliver Kahl, Simone Ferrari, Vadim Kovalyuk, Andreas Vetter, Georgia Lewes-Malandrakis, Christoph Nebel, Alexander Korneev, Gregory Goltsman, and Wolfram Pernice, Spectrally multiplexed single-photon detection with

hybrid superconducting nanophotonic circuits, Optica **4,** 557 (2017).

[46] V. Kovalyuk, O. Kahl, S. Ferrari, A. Vetter, G. Lewes-Malandrakis, C. Nebel, A. Korneev, G. Goltsman, and W. Pernice, On-chip single-photon spectrometer for visible and infrared wavelength range, J. Phys.: Conf. Ser. **1124,** 051045 (2018).

[47] J. P. Sprengers, A. Gaggero, D. Sahin, S. Jahanmirine-jad, G. Frucci, F. Mattioli, R. Leoni, J. Beetz, M. Lermer, M. Kamp, S. Höfling, R. Sanjines, and A. Fiore, Waveguide superconducting single-photon detectors for integrated quantum photonic circuits, Appl. Phys. Lett. **99,** 181110 (2011).

[48] W. H. P. Pernice, C. Schuck, O. Minaeva, M. Li, G. M. Goltsman, A. V. Sergienko, and H. X. Tang, High-speed and high-efficiency travelling wave single-photon detectors embedded in nanophotonic circuits, Nat. Commun. **3,** 1325 (2012).

[49] V. Kovalyuk, W. Hartmann, O. Kahl, N. Kaurova, A. Korneev, G. Goltsman, and W. H. P. Pernice, Absorption engineering of NbN nanowires deposited on silicon nitride nanophotonic circuits, Opt. Express **21,** 22683 (2013).

[50] Carsten Schuck, Wolfram H. P. Pernice, and Hong X. Tang, Waveguide integrated low noise NbTiN nanowire single-photon detectors with milli-Hz dark count rate, Sci. Rep. **3,** 1 (2013).

[51] Patrik Rath, Oliver Kahl, Simone Ferrari, Fabian Sproll, Georgia Lewes-Malandrakis, Dietmar Brink, Konstantin Ilin, Michael Siegel, Christoph Nebel, and Wolfram Pernice, Superconducting single-photon detectors integrated with diamond nanophotonic circuits, Light: Sci. Appl. **4,** e338 (2015).

[52] Oliver Kahl, Simone Ferrari, Vadim Kovalyuk, Gregory N. Goltsman, Alexander Korneev, and Wolfram H. P. Pernice, Waveguide integrated superconducting single-photon detectors with high internal quantum efficiency at telecom wavelengths, Sci. Rep. **5,** 10941 (2015).

[53] Oliver Kahl, Patrik Rath, Simone Ferrari, Andreas Vetter, Christoph Nebel, and Wolfram H. P. Pernice, High efficiency on-chip single-photon detection for diamond nanophotonic circuits, J. Lightwave Technol. **34,** 249 (2016).

[54] Simone Ferrari, Vadim Kovalyuk, Andreas Vetter, Changhyoup Lee, Carsten Rockstuhl, Alexander Semenov, Gregory Gol'tsman, and Wolfram Pernice, Analysis of the detection response of waveguide-integrated superconducting nanowire single-photon detectors at high count rate, Appl. Phys. Lett. **115,** 101104 (2019).

[55] M. Haussler, M. Yu. Mikhailov, M. A. Wolff, and C. Schuck, Amorphous superconducting nanowire single-photon detectors integrated with nanophotonic waveguides, APL Photonics **5,** 076106 (2020).

[56] Sonia Buckley, Adam N. McCaughan, Jeff Chiles, Richard P. Mirin, Sae Woo Nam, Jeffrey M. Shainline, Grant Bruer, James S. Plank, and Catherine D. Schuman, in *2018 IEEE International Conference on Rebooting Computing, ICRC 2018* (Institute of Electrical and Electronics Engineers Inc., Tysons, VA, USA, 2019).

[57] Michael L. Schneider, Christine A. Donnelly, Stephen E. Russek, Burm Baek, Matthew R. Pufall, Peter F. Hopkins,

Paul D. Dresselhaus, Samuel P. Benz, and William H. Rippard, Ultralow power artificial synapses using nanotextured magnetic Josephson junctions, Sci. Adv. **4,** e1701329 (2018).

[58] Miguel Romera, Philippe Talatchian, Sumito Tsunegi, Flavio Abreu Araujo, Vincent Cros, Paolo Bortolotti, Juan Trastoy, Kay Yakushiji, Akio Fukushima, Hitoshi Kubota, Shinji Yuasa, Maxence Ernoult, Damir Vodenicarevic, Tifenn Hirtzlin, Nicolas Locatelli, Damien Querlioz, and Julie Grollier, Vowel recognition with four coupled spin-torque nano-oscillators, Nature **563,** 230 (2018).

[59] Miao Hu, Catherine E. Graves, Can Li, Yunning Li, Ning Ge, Eric Montgomery, Noraica Davila, Hao Jiang, R. Stanley Williams, J. Joshua Yang, Qiangfei Xia, and John Paul Strachan, Memristor-based analog computation and neural network classification with a dot product engine, Adv. Mater. **30,** 1705914 (2018).

[60] Stefano Ambrogio, Pritish Narayanan, Hsinyu Tsai, Robert M. Shelby, Irem Boybat, Carmelo Di Nolfo, Severin Sidler, Massimo Giordano, Martina Bodini, Nathan C. P. Farinha, Benjamin Killeen, Christina Cheng, Yassine Jaoudi, and Geoffrey W. Burr, Equivalent-accuracy accelerated neural-network training using analogue memory, Nature **558,** 60 (2018).

[61] Tifenn Hirtzlin, Marc Bocquet, Bogdan Penkovsky, Jacques Olivier Klein, Etienne Nowak, Elisa Vianello, Jean Michel Portal, and Damien Querlioz, Digital biologically plausible implementation of binarized neural networks with differential hafnium oxide resistive memory arrays, Front. Neurosci. **13,** 1383 (2020).

[62] Jeffrey M. Shainline, Sonia M. Buckley, Nima Nader, Cale M. Gentry, Kevin C. Cossel, Justin W. Cleary, Miloš Popović, Nathan R. Newbury, Sae Woo Nam, and Richard P. Mirin, Room-temperature-deposited dielectrics and superconductors for integrated photonics, Opt. Express **25,** 10322 (2017).

[63] Lukas Chrostowski and Michael E. Hochberg, *Silicon Photonics Design* (Cambridge University Press, Cambridge, UK, 2015).

[64] Yi Zhang, Shuyu Yang, Andy Eu-Jin Lim, Guo-Qiang Lo, Christophe Galland, Tom Baehr-Jones, and Michael Hochberg, A compact and low loss Y-junction for submicron silicon waveguide, Opt. Express **21,** 1310 (2013).

[65] Richard J. Bojko, Jing Li, Li He, Tom Baehr-Jones, Michael Hochberg, and Yukinori Aida, Electron beam lithography writing strategies for low loss, high confinement silicon optical waveguides, J. Vac. Sci. Technol. B, Nanotechnol. Microelectron.: Mater., Process., Meas. Phenom. **29,** 06F309 (2011).

[66] Jeff Chiles, Sonia M. Buckley, Sae Woo Nam, Richard P. Mirin, and Jeffrey M. Shainline, Design, fabrication, and metrology of 10×100 multi-planar integrated photonic routing manifolds for neural networks, APL Photonics **3,** 106101 (2018).

[67] John R. Clem and Karl K. Berggren, Geometry-dependent critical currents in superconducting nanocircuits, Phys. Rev. B—Condens. Matter Mater. Phys. **84,** 174510 (2011).