Superconducting nanowire electronics for alternative computing

by

Emily Toomey

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2020

Author .....
Department of Electrical Engineering and Computer Science
May 14, 2020

Certified by .....
Karl K. Berggren
Professor of Electrical Engineering
Thesis Supervisor

Accepted by .....
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Superconducting nanowire electronics for alternative computing

by

Emily Toomey

Submitted to the Department of Electrical Engineering and Computer Science
on May 14, 2020, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering

Abstract

With traditional computing systems struggling to meet the demands of modern technology, new approaches to both hardware and architecture are becoming increasingly critical. In this work, I develop the foundation of a power-efficient alternative computing system using superconducting nanowires. Although traditionally operated as single photon detectors, superconducting nanowires host a suite of attractive characteristics that have recently inspired their use in digital circuit applications for amplification, addressing, and memory. Here, I take advantage of the electrothermal feedback that occurs in resistively shunted nanowires to develop two new technologies: (1) A multilevel memory cell made by incorporating a shunted nanowire into a superconducting loop, allowing flux to be controllably added and stored; and (2) An artificial neuron for use in spiking neural networks, consisting of two nanowire-based relaxation oscillators acting analogously to the two ion channels in a biological neuron. By harnessing the intrinsic dynamics of superconducting nanowires, these devices offer competitive energy performance and a step towards bringing memory and processing closer together on the same platform.

Thesis Supervisor: Karl K. Berggren
Title: Professor of Electrical Engineering

Acknowledgments

There are so many people who have helped me along this journey, but in particular I would like to thank:

My advisor, Prof. Karl Berggren, for his encouragement and guidance, and for giving me the freedom to explore unusual problems.

Adam McCaughan and Qing-Yuan Zhao, for their invaluable mentorship during my first two years of graduate school. They made me the researcher I am today, and continue to inspire me with their selflessness, persistence, and creativity.

Dorothy Fleischer and Lara Ranieri, for their generosity and for keeping all of us organized.

Qing Hu for his mentorship and support throughout my time in graduate school.

Ken Segall, Tom Ohki, and Nancy Lynch for sharing their knowledge with me and providing invaluable advice about the technical aspects of my work.

My collaborators at Raytheon BBN Technologies, especially Graham Rowlands, Andrew Wagner, Guilhem Ribeill, and Minh-Hai Nguyen, not only for being talented scientists to work with and learn from, but more importantly for being wonderful people.

James Daley and Mark Mondol, for their technical assistance and for cheering up even the most distressing days in the cleanroom.

Leslie Kolodziejski, for being the heart and soul of this department, and for caring so genuinely about the graduate student body.

Navid Abedzadeh, for being my friend, confidant, and fellow artist since our first day at MIT together, and for giving me invaluable support, without which I wouldn't have survived these past five years.

Simon Vecchioni, for starting me off on this whole career path since our first project together seven years ago. The world is undoubtedly a better place for having researchers like Simon in it.

Marco Colangelo, for embarking on numerous projects with me and for making the research process all the more enjoyable.

Di Zhu, Brenden Butters, Owen Medieros, and all members of the QNN group for their research advice, daily discussions, and camaraderie.

Jenna Norton, for putting up with me as a roommate since our sophomore year of college and for keeping me grounded.

Most of all to my parents, my brothers, and my sisters-in-law. If I accomplish anything, it's because I have the best support system anyone could ask for. They also aren't afraid to remind me that a broken nanowire is far from the end of the world.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The natural dynamics of devices at the material level have long inspired new systems and architectures for computing. Transistors, for example, use the manipulation of energy bands in silicon to control the flow of electrons, allowing them to act as amplifiers or digital switches. This deceptively simple behavior is responsible for the astonishing growth in technology over the past fifty years, setting the foundation for digital computing as we know it today.

However, transistors are now at a point where they can no longer support continued expansion of the technologies they initially created. Critical performance metrics like the maximum operation frequency of microprocessors once steadily improved every year, but flatlined around 2005 due to issues including scaling limitations, heating costs, and parasitics. Additionally, the physical separation between memory and processing elements in traditional von Neumann architectures has created a notorious "memory bottleneck" that constrains the minimum achievable access time [1].

These performance plateaus have serious implications in light of the growing need for more energy-efficient computation. With new applications like autonomous vehicles and the Internet of Things demanding faster computation with more data, estimates project that our system-level energy costs of computation will surpass the world's total energy production in less than 20 years [2]. To keep up with the data revolution, it is therefore critical to look beyond silicon for new materials and architectures with superior energy efficiency that can grow to support the needs of the

next generation of technology.

1.1 Alternative architectures: neuromorphic computing

One major branch of research that has grown out of this need is neuromorphic computing, which aims to reach the performance and energy efficiency of the human brain by mimicking its structure and dynamics. Unlike traditional architectures, the brain communicates using electrical spikes called action potentials and is organized into hundreds of thousands of parallel connections. Furthermore, biological neurons unite both a memory (synapse) and processing element (soma) together within a single cell. These characteristics enable the brain to perform complex tasks with impressive energy costs; for instance, the human brain's typical power budget for conducting simultaneous tasks like movement and recognition is around 20 W, while a standard computer uses more than ten times that amount to classify just 1,000 objects [3]. As a result, brain-inspired architectures are highly promising alternatives that could vastly expand computing power while decreasing the costs.

Spiking neural networks (SNNs) are among the most bio-realistic approaches to neuromorphic computing. Whereas other architectures like those used in deep learning adopt only select aspects of the brain like a multilevel hierarchy, SNNs reproduce actual spiking dynamics. Since spiking signals allow for event-based computation, SNNs may achieve superior energy efficiency and could be key in applications like autonomous vehicles with event-based sensors.

Directly encoding spiking dynamics in software allows for precise control, but is often too computationally expensive to be practical for large-scale parallel networks. To realize an extensive network, researchers are instead turning to hardware devices and circuits that can naturally generate spiking behavior on their own. While spiking has been explored in a variety of materials, the existing devices face numerous limitations. For instance, SNNs made from CMOS [4] allow for scalability and ease

of integration with control circuitry, but suffer from high power dissipation and require many components to achieve biological realism. Magnetic materials using the spin-torque effect can also generate spiking [5], but often have small on/off ratios [3]. Memristors based on filament formation have shown promise as the synaptic memory component of neural networks [6][7], but are vulnerable to device and cycle variations, and require additional CMOS circuitry to produce spikes. These shortcomings motivate the need for a scalable, power-efficient device that naturally generates spiking and integrates easily with existing CMOS controls.

1.2 Alternative hardware: superconducting electronics

Devices made from superconducting materials are frequently championed for their potential for fast, energy-efficient computation. Most superconducting circuits use Josephson junctions (JJs), devices consisting of two superconducting layers sandwiching a thin insulating or normal barrier [8]. JJs are quantum mechanical devices, relying on the tunneling of superconducting electrons or Cooper pairs through the barrier to create rapid ($>$100 GHz) voltage pulses with very low energy dissipation (~ 0.1 aJ). These finite pulses have inspired a suite of JJ-based logic, including rapid single flux quantum (RSFQ) logic [9], reciprocal quantum logic (RQL) [10], and adiabatic quantum flux parametron (AQFP) logic [11]. In addition to their low switching energies, JJs and superconducting materials in general have no static power dissipation in their interconnects, allowing them to maintain competitive energy performance even after accounting for cooling costs [12].

Despite these advantages, there are significant limitations to implementing a large-scale cryogenic computer based on JJs. For one, JJ pulses are often below the energy sensitivity of transistors and outside the bandwidth of CMOS circuits, making them incompatible for integrating with CMOS controls [13]. Furthermore, the output impedance of a JJ in its resistive state is typically on the order of a few ohms, which

is insufficient for driving transistors or supporting fan-out. Solutions to work around these issues, such as stacking multiple JJs to create an amplified signal, are often bulky, energetically expensive, and highly sensitive to fabrication variations, making them unsuitable for the development of a cryogenic computing architecture on their own. Additional issues are introduced by JJ logic circuits requiring large-scale distribution of bias current to each gate, which is typically achieved using bias resistors attached to a DC voltage rail. The power dissipated by bias resistors may be ten times more than the switching dissipation of a JJ, substantially increasing the overall static power dissipation costs and producing parasitic heat loads [14] [15]. Furthermore, large DC bias currents can produce magnetic fields that penetrate JJ circuits and reduce their operating margins [16]. Some architecture modifications that have been adopted to avoid these effects include replacing the bias resistors with additional JJs [15], or using an AC bias with an inductive coupling network [14].

1.2.1 Superconducting nanowire devices

Although traditionally operated as single photon detectors, superconducting nanowires monolithically patterned as strips on thin films have started to be used in circuit applications where JJs fall short. Unlike the tunneling that occurs in JJs, the resistive state of nanowires is dominated by a Joule heated "hotspot" with an impedance on the order of $\sim$1–10 kΩ, making them ideal for amplification and fan-out. This high impedance also allows nanowire devices to connect JJ circuits to external CMOS electronics, highlighting their unique role as an intermediary between silicon and superconducting platforms.

Figure 1-1(a) shows a simplified circuit schematic of a superconducting nanowire and its current-voltage characteristics. When biased below their critical current I_c, nanowires transport current without developing a resistance or voltage. Once the current exceeds I_c, superconductivity breaks down and the aggregation of hot quasiparticles spurs the growth of a resistive hotspot. The nanowire remains in the resistive state until the bias current has been reduced, and the hotspot relaxes. In addition to switching by surpassing a critical current density J_c, nanowires can also be made

resistive by applying a critically high magnetic field H_c or temperature T_c, or through the arrival of a photon.

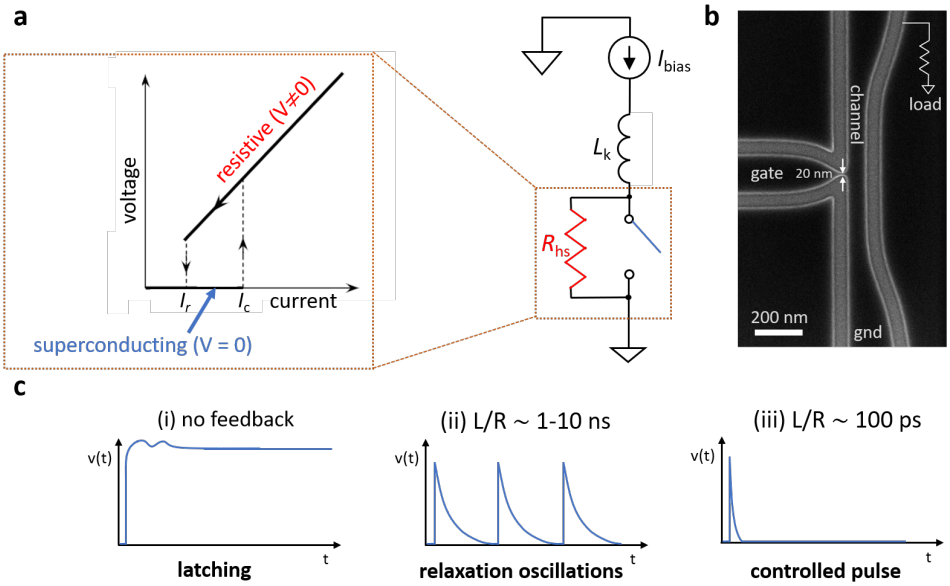

Figure 1-1: Superconducting nanowire devices. (a) Simplified circuit schematic of a superconducting nanowire, modeled as a hotspot resistance R_{hs} placed in parallel with a switch shorting to ground, in series with a kinetic inductance L_k. (b) Scanning electron micrograph of an nTron. The dark regions are the superconducting film, while the lighter outlines are the underlying substrate. (c) Different modes of operation due to electrothermal feedback. (i) No feedback circuit leads to latching. (ii) A series inductance producing an L/R time constant on the order of ∼1 ns creates relaxation oscillations. (iii) A smaller inductance with a shorter L/R time constant on the order of ∼100 ps allows for controlled finite voltage pulses.

Devices made from superconducting thin films manipulate hotspot formation to serve their unique functionality. Starting in 1956, Dudley Buck proposed a device made from intertwined superconducting wires called the cryotron [17], in which current passing through one wire induced a critically high magnetic field that penetrated the other, causing it to switch. A design for this device in a planar geometry was proposed a few years later [18] [19], reinforcing the technique of using the propagation of hotspot boundaries in thin films to create new device operations [20].

While superconducting thin-film devices took a backseat after the invention of JJs in 1962, they have gained new momentum in recent years due to improvements in patterning techniques like electron beam lithography, which have enhanced our

ability to control hotspot formation at the nanoscale. For instance, superconducting nanowire single photon detectors (SNSPDs) are made from meandered wires that are 100-nm wide or narrower so that the energy dissipated by a single photon is enough to create a localized hotspot [21][22]. The kiloohm impedance of the hotspot then creates a large voltage signal that is used to communicate the arrival of a photon.

More recently, a nanocryotron (nTron) building off of Buck's previous work was developed for use in digital circuits [23]. As shown in Fig. 1-1(b), the nTron is a three-terminal device that uses localized hotspot formation in a narrow ($\sim$ 20 nm) gate region to switch a wider channel wire into the resistive state, creating a high impedance region that generates a large output current. This switching mechanism allows the nTron to support fan-out, drive high impedances, and amplify signals. A derivative of the nTron in which the gate and channel are electrically disconnected but thermally coupled has also recently been used as a bit-select element for magnetic memory arrays, illustrating the compatibilty of superconducting nanowire devices with other materials in large-scale circuits [24]. Other recent nanowire devices have found use as read-out sensors [25], memory cells [26] [27], and frequency mixers [28].

1.2.2 Kinetic inductance

In addition to the formation of a high-impedance normal domain, an equally unique characteristic of nanowire devices is the contribution of kinetic inductance [29]. Unlike magnetic inductance, kinetic inductance relates to the inertia of electrons. While kinetic inductance is negligible in normal metals at sub-terahertz frequencies due to electron scattering, superconductors have infinitely large scattering times, causing kinetic inductance to be significant.

An intuitive sense of kinetic inductance L_k can be derived by relating the kinetic energy of Cooper pairs to the inductive energy:

$$(2 \cdot \frac{1}{2} m_e v^2)(n_s l A) = \frac{1}{2} L_k I^2 \tag{1.1}$$

where $2m_e$ is the mass of the two superconducting electrons in a Cooper pair, v is their

velocity, n_s is the density of superconducting electrons, l is the length of the structure, and A is the cross-sectional area. The current I can be expressed as $I = (2e)vn_sA$ to give a final expression for the kinetic inductance:

$$L_k = \frac{m_e l}{2e^2 n_s A} \tag{1.2}$$

As seen by Eq. 1.2, the inductance of a nanowire scales with its length for a fixed width. From a simplified circuit perspective, a nanowire can therefore be modeled as a variable resistor in series with an inductor whose magnitude depends on the device geometry, as shown in Fig. 1-1(a).

Previous work demonstrated that the presence of kinetic inductance allows nanowires to produce similar behavior to JJ-based circuits. For instance, when placed in parallel with a shunt resistor R_s, nanowires emit rapid pulses (∼100 MHz–1 GHz) as a result of electrothermal feedback between the nanowire and the shunt with a time constant of L_k/R_s [28]. Although these speeds are slower than those of JJs, the longer pulses are large enough to drive CMOS gates and support fan-out, and can easily be viewed on an oscilloscope for diagnostics. Additionally, nanowire loops based on kinetic inductance have been used as cryogenic memory cells [30][27], analagous to superconducting quantum interference devices (SQUIDs) in JJ circuits. As a result, nanowires show promise as a suitable platform for developing low-power alternative computing architectures that can be integrated with CMOS controls.

1.3 Thesis goal

The electrothermal feedback between a nanowire and resistive shunt presents a new opportunity to use nanowires as dynamic switching elements whose behavior can be changed by tuning the kinetic inductance [31]. As shown in Figure 1-1(c), a nanowire can operate in three main regimes depending on the inductor's magnitude. In a case when there is no feedback circuit, a nanowire will latch when it switches into the resistive state, and will remain resistive until the bias current is removed. However, when a feedback circuit is introduced and the inductor gives an L/R time constant on

the order of nanoseconds, the nanowire can produce sustained relaxation oscillations with a bias-dependent frequency. Here, R is the sum of the shunt resistance and the nanowire resistance, while L is the kinetic inductance of the nanowire. If the inductance is further reduced to give an L/R time constant on the order of $\sim$100 ps, the nanowire can emit discrete, narrow voltage pulses without oscillating.

This thesis takes advantage of the different regimes of electrothermal feedback to develop two new devices for computing: (1) A multilevel nanowire memory cell that uses low-inductance resistive shunting to controllably trap flux in a superconducting loop; and (2) An artificial neuron based on nanowire relaxation oscillators to serve as a hardware spiking element in neural networks. Together, these devices aim to set the foundation for a superconducting neuromorphic computing architecture that brings memory and processing closer together on the same platform with improved energy efficiency.

This thesis will be organized as follows:

Chapter 2— A multilevel memory cell. In this chapter, we present the design of a multilevel memory cell, which uses resistive shunting to controllably add flux into a superconducting nanowire loop. Through both simulations and experimental testing, we show that the number of states in the memory can be tuned simply by changing the circuit parameters, creating a flexible system that could enable denser cryogenic memory.

Chapter 3—The nanowire neuron: design. This chapter describes the design of a nanowire-based artificial neuron for spiking neural networks. Using simulations conducted in LTspice, we show how two nanowire relaxation oscillators can be used analagously to the two ion channels in a biological neuron. We demonstrate that the nanowire neuron is capable of reproducing multiple biorealistic behaviors and supporting fan-out, while maintaining competitive energy performance.

Chapter 4—The nanowire neuron: demonstration. This chapter summarizes the fabrication and testing of a nanowire soma based on the designs presented in Chapter 3. Measured spiking characteristics are used to obtain a realistic circuit model of the device, which is used in two applications of spiking neural networks: a simple image recognition task, and winner-takes-all competition.

Chapter 5—Fabrication studies. While developing the electronics discussed in this thesis, fabrication studies on the compatibility of electron-beam resist processes with superconducting films were conducted. This chapter describes two specific investigations. First, we show that a common developer called tetramethylammonium hydroxide reacts with niobium nitride thin films to form a salt that acts as an etch barrier. Afterwards, a study on ma-N 2400 series photoresist as an electron-beam resist for patterning superconducting nanoscale devices is presented.

Chapter 6—Conclusion and outlook. In this chapter we review the central findings presented in this thesis and describe future investigations that may be built off of these results.

Chapter 2

A multilevel memory cell

Computers using superconducting elements promise to achieve competitive energy performance, even while accounting for cooling costs, due to advantages like low switching energies and negligible static power dissipation in their interconnects. While the active development of superconducting logic has produced a suite of logic families, such as rapid single flux quantum (RSFQ) logic [9] and reciprocal quantum logic (RQL) [10], less focus has been placed on the creation of a scalable superconducting memory [12]. As a result, systems are forced to rely on a combination of low-temperature logic and room-temperature memory, placing higher demands on interconnects and dissipating more power. This physical separation between memory and processing elements also worsens the "memory bottleneck" dilemma, limiting the minimum achievable access time.

As a step towards achieving a unified cryogenic computing system, we have developed a multilevel memory cell using superconducting nanowires. This work is built off of our previous design of a binary memory cell [30], with the goal of obtaining a denser cryogenic memory. In addition to allowing for higher density as the limits of physically shrinking a unit cell are approached, multilevel operation may allow for more energy savings, since it provides more information given the same peripheral circuitry costs as a binary cell.

The text that follows includes reprints from a work that was originally published in *Physical Review Applied* [32]. I would like to acknowledge Murat Onen for the

simulations he contributed to this work, Marco Colangelo for his help with fabrication, and Brenden Butters and Dr. Adam McCaughan for their assistance with experimental design.

2.1 Binary nanowire memories

One of the most unique characteristics of superconductors is their ability to store quantized flux in the form of persistent current. This functionality stems from the Meissner effect, which describes how the magnetic field within a superconductor is always zero. Physically speaking, when a magnetic field penetrates a superconducting ring above its critical temperature, it triggers a circulating current within the material that exactly cancels the field once the ring is cooled down. When the external field is turned off, the induced current (also called persistent current) will continue to circulate indefinitely, since superconductors are perfect conductors. A schematic of this effect is shown in Figure 2-1.

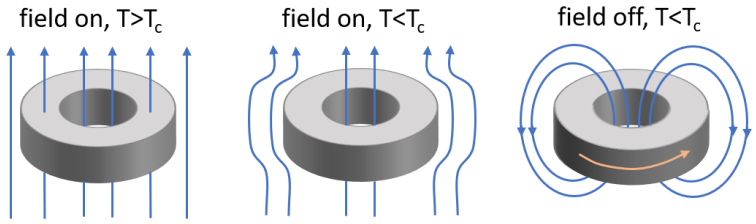

Figure 2-1: Schematic of the Meissner effect in a superconducting ring. A field is applied above the critical temperature, leaving a persistent current once the material is cooled down and the field is removed. Blue lines indicate magnetic fields, while the orange line represents the persistent current.

Past superconducting memories have taken advantage of this effect by intentionally storing persistent current in superconducting loops. In these cases, the "1" state is usually denoted by current circulating in one direction, while the "0" state is either the absence of current or the flow of current in the opposite direction. While this operation can be achieved through single flux quantum circuits, JJ-based memories struggle from scalability issues due to the large areas of individual JJs (usually on the order of 1 μm^2), the amount of space taken up by the transformers and SQUID

amplifiers used to operate the memories, and the need to keep individual cells far enough apart to avoid crosstalk from magnetic coupling [30].

Given these limitations, superconducting nanowires have recently been proposed as alternative platforms for making cryogenic memories [30][27][33][34]. In general, the kinetic inductance of superconducting nanowires is two orders of magnitude larger than the magnetic inductance, allowing nanowire memory loops to be scaled down. Rather than through the application of a magnetic field, current is typically stored in nanowire memory cells by applying a bias current that breaks superconductivity on one side of the loop [30][33]. When that side of the loop switches from superconducting to normal, the portion of the bias current that initially flowed through it is diverted into the other side. This diverted current then charges the kinetic inductance of the remaining portion of the loop, creating a change in flux that induces a current. When the switched side of the cell heals, a persistent current remains.[1] Unlike JJ-based memories, nanowire memories do not have the same issues of crosstalk from magnetic coupling, since magnetic fields typically penetrate through the thin films the memories are patterned from.

Although recent nanowire memories have achieved impressive performance metrics, such as a bit error rate less than 10^{-7} [30], they are limited to binary operations due to the nature of their writing; the resistive portion of the cell expels nearly all current from its high-impedance hotspot into the loop once it switches, storing the maximum amount of current. As a result, it is not possible to successively store smaller amounts of current over multiple writing events to obtain additional states.

Here, we present a superconducting nanowire memory cell based on thermal principles that demonstrates a controlled, moderated output. Unlike other nanowire memories, the device responds proportionally to the magnitude of an input signal and can be operated to achieve multiple discrete states. The key to this design is the introduction of a low-inductance resistive shunt placed close to the write port that absorbs the majority of the diverted bias current, allowing a smaller, controlled

[1] In this thesis, the persistent current will also be referred to as trapped flux, in units of fluxons. However, it should be noted that the more accurate term for quantized persistent current is "fluxoids".

amount of flux to be stored in the superconducting loop in quantities of $n\Phi_0$, where n is an integer and Φ_0 is the magnetic flux quantum (2.07×10^{-15} Wb). We experimentally show that the amount of flux per event n is dictated by circuit parameters and validate these results with electrothermal simulations. We anticipate that this device will serve as the foundation for a nanowire multilevel memory or multilevel-logic circuit elements.

2.2 Device description

Figure 2-2 summarizes the device architecture and its basic characteristics. As shown in the scanning electron micrograph (SEM) in Fig. 2-2(a) and the circuit model of Fig. 2-2(b), the device is comprised of three superconducting nanowire elements: a narrow 60-nm-wide constriction, a storage loop, and a nanoscale readout tool known as the yTron [25]. All three elements are fabricated together on an approximately 20-nm-thick niobium nitride (NbN) film on a silicon oxide substrate using electron-beam lithography. In addition to the nanowire components, a resistive metal shunt is patterned in parallel with the constriction to reduce Joule heating and provide damping, similar to the purpose served in resistively shunted JJs. Previous attempts to shunt nanowires found that the series inductance between the shunt and the constriction plays a critical role in the effectiveness of the damping; high series inductance produces relaxation oscillations, while increasing the inductance even further makes the shunt resistor almost completely ineffective [28]. To avoid this problem, the series inductance between the write constriction and the metal shunt must be minimized to allow for discrete, controlled pulses rather than oscillations.

Fig. 2-2(c) shows the current-voltage characteristics of an isolated shunted nanowire patterned alongside the device with dimensions identical to those of the constriction. The absence of hysteresis, as shown by the lack of separation between the switching and retrapping currents, indicates that the shunt resistor is able to reduce Joule heating through the constriction by effectively diverting the bias current, thereby reducing power dissipation in the nanowire and allowing it to regain the superconducting state

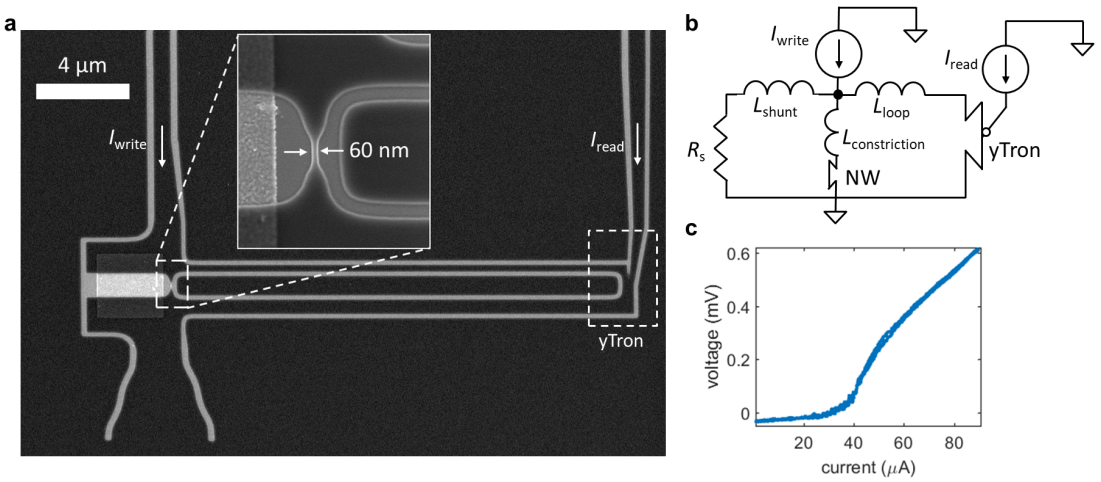

Figure 2-2: Device design and characterization. (a) Scanning electron micrograph of the device. The dark area is NbN film, while the light outlines are the underlying substrate. The inset shows an enlarged view of the 60-nm-wide constriction in parallel with the resistive shunt. The righthand side of the loop is connected to a yTron with arm widths equal to 300 nm. (b) Circuit schematic of the device. (c) Current-voltage characteristics of an isolated shunted nanowire of width = 60 nm, $R_s = 5\ \Omega$.

more quickly [35]. Observation of the amplified RF output of the device within a bandwidth of 2 GHz did not reveal any relaxation oscillations, suggesting that the shunt inductance was sufficiently low [28].

2.2.1 Writing operation

To trap flux into the loop, a bias write current I_{write} is inductively split to the nanowire constriction in the amount of αI_{write}, where $\alpha = L_{loop}/(L_{constriction} + L_{loop})$, L_{loop} represents the inductance of righthand side of the memory cell and $L_{constriction}$ is the inductance of the constriction on the lefthand side. For the device shown in Fig.2-2, $L_{constriction} = 284$ pH and $L_{loop} = 1.87$ nH, leading to $\alpha = 0.87$. Once the sum of αI_{write} and any existing current circulating in the loop surpasses the critical current of the constriction $I_{c,NW}$, the nanowire switches and the bias current is diverted away from the constriction to the shunt resistor and the righthand side of the loop. By shunting the majority of the bias current, the resistor allows the nanowire to recover the superconducting state more quickly and limits the amount of

current that charges L_{loop}, thus controlling the amount of flux that is trapped once the constriction heals. After the constriction heals, a persistent current circulates in the loop in quantized units of $n\Phi_0/L$, where n is an integer and L is the total inductance of the complete loop. Since the magnetic inductance of this device is $<$0.33 fH, the total loop inductance is dominated by the kinetic inductance of the superconducting nanowires. In this particular device, persistent current is quantized as approximately 0.95 μA/fluxon.

2.2.2 Reading operation

The amount of circulating current in the loop can be nondestructively read out using a nanowire device called the yTron. As described by McCaughan *et al.* [25], the yTron is a three-terminal nanoscale device with two adjoining arms whose switching currents depend on one another as a result of current crowding [36] around the narrow intersection point. In our device, the left arm of the yTron forms part of the superconducting loop so that the switching current of the right arm I_{switch} is a function of the amount of circulating current—a higher circulating current in the clockwise direction flowing through the yTron's left arm will result in a higher switching current in the yTron's right arm. Since the two arms of the yTron are electrically disconnected from one another, switching the right arm does not break superconductivity in the left. As a result, the state of the loop is undisturbed by the reading process in which I_{switch} is measured by applying a bias current I_{read} to the yTron's right arm until it switches and generates a voltage, allowing us to nondestructively sense the amount of circulating current in the loop. An enlarged view of the yTron and an example of its measured sensitivity curve is shown in Figure 2-3.

2.3 Demonstration of multilevel operation

Here we describe the fabrication process used to build the multilevel memory cell, followed by the measurements that were conducted to characterize its behavior. For comparison purposes, multiple cells of different circuit parameters were fabricated

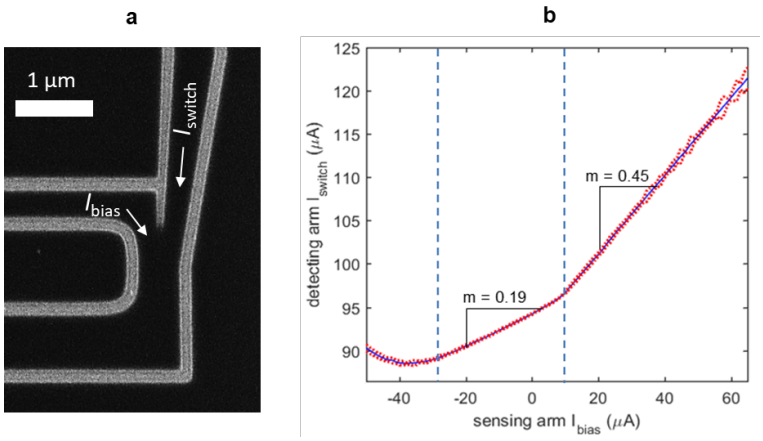

Figure 2-3: Example of a yTron. (a) Scanning electron micrograph showing the readout scheme for the memory cell. (b) yTron sensitivity curve showing the change in switching current of the readout arm as a function of bias current running through the lefthand side. Both arms were 300 nm wide. The dark blue trace shows the mean value of 100 sweeps, while the red dashed lines are ± one standard deviation from the mean.

together on the same chip, including an unshunted memory cell similar to those in prior literature.

2.3.1 Fabrication

The devices presented in this chapter were fabricated using a multistep lithography process illustrated in Figure 2-4. Initial efforts to fabricate the resistors on top of the NbN film failed due to contact resistance from the oxidized surface. As a result, the final process involved placing the resistor beneath the NbN film, as was done in similar work on Nb nanowires [37].

The shunt resistors and alignment marks were first patterned with electron-beam lithography (Elionix F125) using a bilayer resist process. We first spun polymethyl methacrylate (PMMA) copolymer EL6 (6% ethyl lactate) at 5 krpm for 60 s, then spun the positive-tone resist gL2000 (Gluon Lab LLC) at 6 krpm for 60 s. Following exposure, the resist was developed in o-xylene and MIBK:IPA in a 1:3 ratio at room temperature. A 10-nm Ti + 25-nm Au metal bilayer was then evaporated, after which lift-off was achieved in N-Methyl-2-pyrrolidone (NMP) heated to 60 °C for one hour.

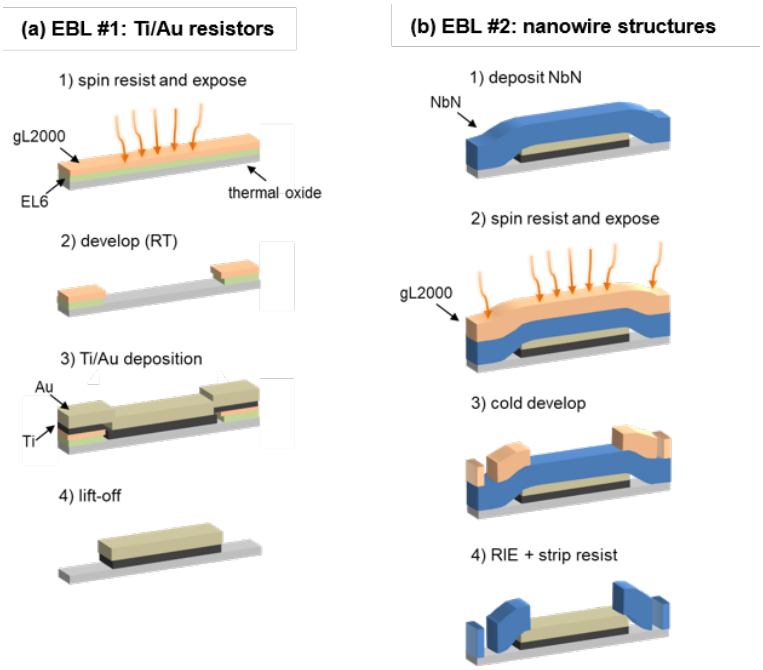

Figure 2-4: Fabrication process for the shunted memory cell (a) Fabrication of the alignment marks and resistors. (b) Fabrication of the nanowire structures.

An apporximately 20-nm-thick NbN film was then deposited in an AJA sputtering system, following the procedure described in Ref. [38]. The resulting sheet resistance was 150 Ω/sq and the critical temperature was 8.5 K.

Afterwards, the nanowire structures were patterned in a second electron-beam lithography step using gL2000, followed by cold development in o-xylene at 5°C. The patterns were transferred to the underlying superconducting film via reactive ion etching in CF_4 (Plasmatherm, RF power of 50 W, chamber pressure of 10 mTorr). To reduce the series inductance between the shunt and the write constriction, the two components were patterned as close to one another as possible, and the nanowire leads were made wide to reduce the number of squares of material.

2.3.2 Experimental setup

All measurements were performed with the devices submerged in liquid helium at 4.2 K. The devices were adhered to a printed circuit board, and electrical connections were made using aluminum wire bonds. The circuit board ports were connected to

room-temperature electronics outside of the liquid helium dewar through CMP cables.

DC characterization

Current-voltage characteristics like those shown in Fig. 2-2(c) were measured by applying a sinusoidal bias current with an arbitrary waveform generator (Agilent AWG33622a) at a sweep frequency of 10 Hz with a 10 kΩ series resistor. The DC output voltage was read by a 2-GHz, real-time oscilloscope (LeCroy 620Zi) after amplification through a low-noise preamplifier (Stanford Research Systems SRS560).

Memory loop operation

Figure 2-5 shows the experimental setup for operating the memory loop with a pulsed input. To apply a write pulse, a pulsed-voltage waveform of widths ranging from 5 ns to 100 μs and heights ranging from 50 to 550 mV were sent to the constriction input through a 30 dB attenuator after being passed through a pulse splitter that connects to the oscilloscope. To apply a DC write bias instead of a pulse, a battery source (Stanford Research Systems SIM928) could be connected to the write port through a 100 kΩ series resistor and a DC-1.9 MHz coaxial low-pass filter (MiniCircuits).

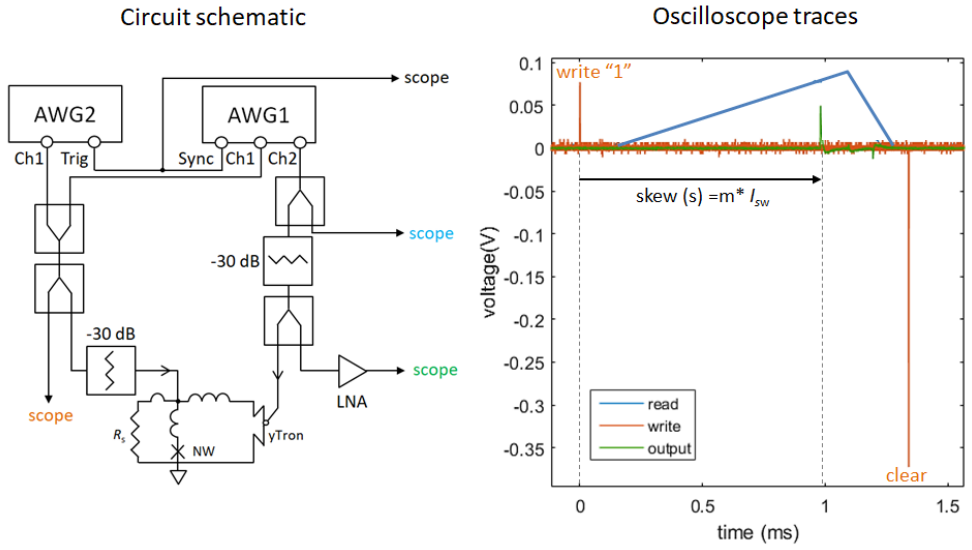

Figure 2-5: Experimental setup for measuring the response of the device to a pulsed input. (a) Circuit schematic of the setup.(b) Corresponding oscilloscope traces. In this case, a negative clearing pulse was also applied to the write port after reading.

For readout, the switching current of the yTron was measured by applying a

voltage pulse through a 30 dB attenuator to the device, and measuring the skew between the oscilloscope trigger riding edge and the time at which a voltage output from the yTron is recorded, signifying a switching event. The voltage pulse had a frequency of 500 Hz, width of 650 μs, rising edge of 400 μs, and height of 560 mV. The yTron output was sent through a pulse splitter and a low-noise amplifier (RF Bay LNA-2000, bandwidth: 10 kHz-2000 MHz, gain: 26 dB) before being read by the oscilloscope. The skew was then converted to units of switching current based on the slope of the bias waveform.

2.3.3 Response to pulsed input: controlled dynamics

Figure 2-6 shows the response of a shunted memory cell to an input voltage pulse of varying amplitude and width; the response is compared to that of an otherwise identical device lacking a resistive shunt. For these measurements, a large negative pulse (width = 10 μs, height = -1.3 V) was sent to the constriction to reset the loop after each read operation.

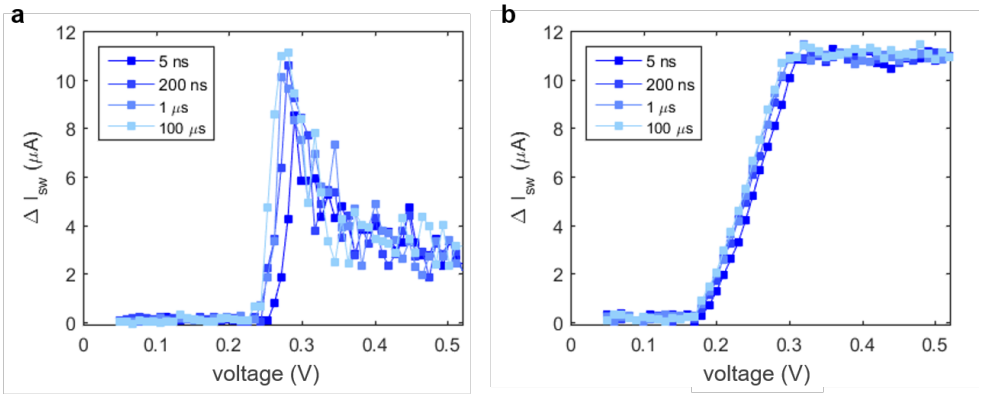

Figure 2-6: Response to the voltage amplitude and pulse width of the write input. The switching current of the yTron readout arm is plotted in terms of $\Delta I_{sw} = I_{switch} - I_{switch}(v_{write} = 0)$. (a) Unshunted memory loop. (b) Memory loop with a shunt resistance of 5 Ω.

As shown in Fig. 2-6(a), the amount of stored current in the unshunted device sharply increases with increasing input voltage, but then abruptly drops off, suggesting instability. This response was also observed in the memory cell reported in Ref.

[33], and was speculated to be due to overheating of the constriction, causing flux to be lost. In contrast, the response of the shunted device in Fig. 2-6(b) shows that the amount of flux stored in the loop increases proportionally with input voltage. Unlike the unshunted constriction, in the shunted device there was no sudden loss of stored flux or signatures of unstable oscillations, implying that heating in the constriction is moderated by the presence of the resistive shunt.

2.3.4 Response to DC input: multilevel behavior

To demonstrate the flux-shuttling capabilities of the shunted device, we measured its dependence on the previously written state by ramping a DC bias current on the write port without resetting the loop, and recording the switching current of the yTron readout at every bias point. As shown in Figure 2-7 (a), ramping the write current on the shunted memory cell produces either increasing or decreasing steps in the switching current of the yTron readout, signifying a sudden addition or subtraction in the amount of trapped flux. The horizontal lines show that the steps can be categorized into seven distinguishable states, revealing that the successive switching of the constriction produces controlled, incremental changes in the amount of circulating current in the loop, rather than storing the maximum amount of current every time. In contrast, Fig. 2-7(b) displays the results from repeating the measurement on an unshunted device of the same geometry. In this case, no intermediate states are observed, and the loop traps nearly its maximum amount of circulating current whenever the constriction switches. Thus, it is not possible to achieve distinguishable intermediate states without the presence of a resistive shunt.

The slight variation in the position of the seven states of the multilevel memory occurs due to instability in the plateaus, representing when the loop current is nearly maximized (close to $I_{c,NW}$) and may have lost a small amount of flux to achieve stability. Despite the small shifts at the plateaus, the seven states have well-separated mean values including consideration of their standard deviations, as shown in Figure 2-8(a).

Although the separate states are distinguishable, they appear to be unevenly

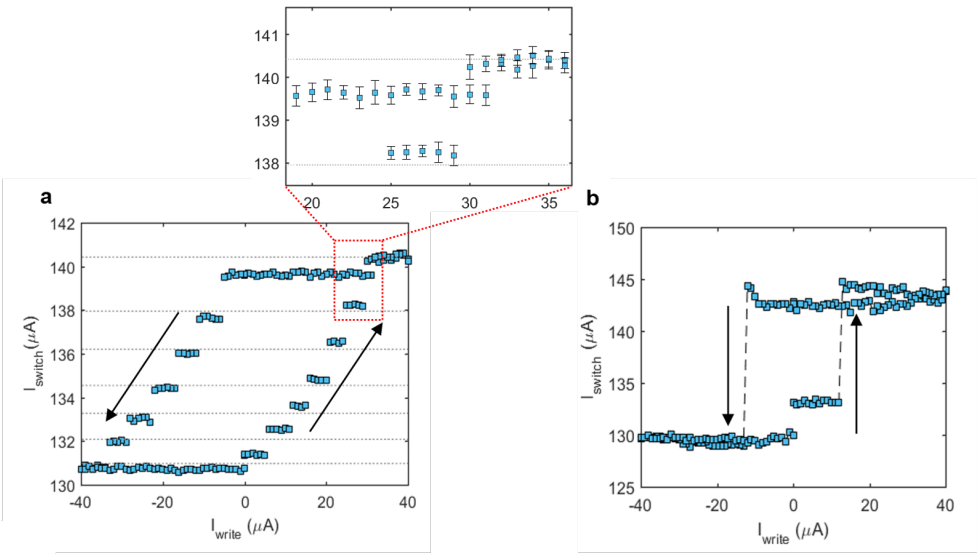

Figure 2-7: Demonstration of controlled flux shuttling. Switching current of the yTron readout arm in response to a DC bias ramp of ± 40 μA applied to I_{write}. Each point is the mean of 10 measurements of the yTron switching current, which had standard deviations ranging from 0.04–0.33 μA. Each bias current step in I_{write} is applied for 100 ms before being turned off during the reading operation. (a) Memory cell with a shunt resistance of 5 Ω. The inset shows an enlarged view of the region where the standard deviations were the largest, as illustrated by the error bars. (b) Unshunted memory cell.

spaced in terms of yTron switching current. This nonlinear spacing can be attributed to nonlinearities in the yTron's sensitivity curve. Depending on the geometry of the yTron, there may be a nonlinear relationship between the amount of circulating current and the induced change in switching current, like the trend shown in Fig. 2-3 (b). Additionally, the sensitivity of the yTron is highly dependent on the intersection point between its two arms, which has a radius of curvature < 5 nm, leaving room for fabrication variability and thus differences in sensitivity between yTrons of identical design. As a result, while the yTron is an effective tool for sensing changes in circulating current, we found that it can be imprecise for extracting the exact amount of current stored in the memory loop.

To bypass this shortcoming, we used the yTron only to sense when a change in trapped flux occured, and examined the corresponding bias current at each of the points of change in order to infer the magnitude of the loop current. Fig. 2-8(b) shows

the bias currents at each of the first 14 steps of the plot in Fig. 2-7(a). The bias current at each step can be used to estimate the amount of circulating current remaining in the loop, given that the transition occurs when the nanowire constriction switches, or when $|\alpha I_{write} + I_{loop}| \geq I_{c,NW}$. The average zero flux state ($I_{loop} \approx 0$) occurs at $I_{c,NW} \approx 20.48 \pm 1.45 \mu$A over a set of eight write bias ramps. While the seven levels in terms of yTron switching current are spaced unevenly, the steps in terms of write bias current occur at nearly equal intervals of approximately 5 μA, corresponding to about 5 Φ_0 of trapped flux for this design. Through this perspective, it is possible to infer that the loop gains or loses about 5 Φ_0 of trapped flux every time the constriction switches. Repeating this measurement over eight ramping cycles with a finer sweep produced an average of 4.77 Φ_0 per step in circulating current, with variation from an integer number ($n = 5$) expected to be caused by noise in the measurement setup.

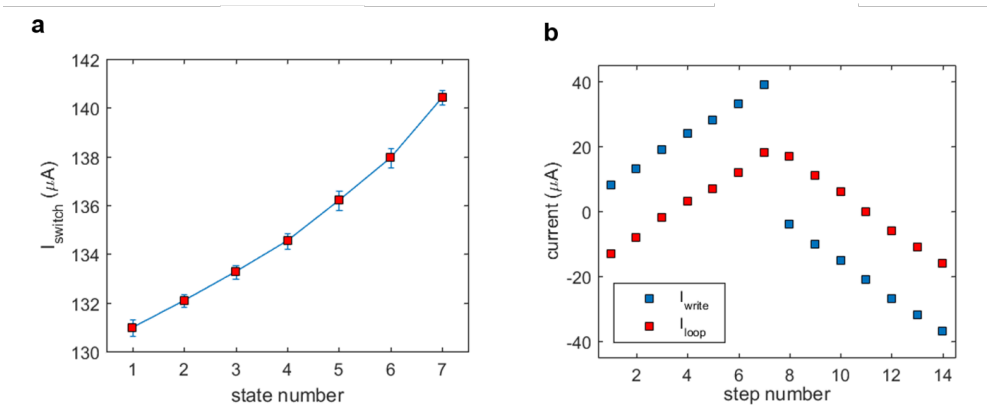

Figure 2-8: Multilevel behavior of the shunted memory cell. (a) The seven states observed in Fig. 2-7(a). Red squares are the mean of eight repeated cycles, and the error bars are the standard deviation. (b) The calculated stored current (red squares) for each of the first 14 steps in yTron switching current in Fig. 2-7(a), using $I_c = 21$ μA. Blue squares represent I_{write} at the initial point of each of the steps.

Table 2.1 summarizes the results of repeating this measurement on two other devices with varying values of L_{loop} and R_s. All other geometries and parameters were kept the same. For each device, the bias write current was ramped to $\pm$ 40 μA as before, in increments of about 10% of the amount of current per fluxoid, or $0.1\Phi_0$. The results from Table 2.1 show that increasing R_s decreases the number of memory states and increases the average number of fluxons trapped per switching event, while

decreasing the loop inductance slightly reduces the amount.

Table 2.1: Change in flux per switching event

	L_{loop}	R_s	No. of states	μ	σ
Device 1	1.87 nH	5 Ω	7	4.77 Φ_0	1.23 Φ_0
Device 1	1.87 nH	7.8 Ω	5	7.63 Φ_0	1.5 Φ_0
Device 1	0.66 nH	7.8 Ω	3	5.87 Φ_0	1.02 Φ_0

Mean change of circulating current is measured over eight complete ramping cycles, and represented in terms of an average change in flux μ with a standard deviation of σ.

2.4 Simulations

To better understand how the device parameters influence the amount of flux n trapped per switching event, we simulated the dynamics of the system. The circuit simulator is described in Ref. [39], and uses a superconducting nanowire model implemented in MATLAB that includes thermoelectric dynamics of hotspot formation and decay [31]. Physical parameters for the materials were derived from prior literature [40], and were adjusted to match experimental results. Flux quantization in the superconducting loop was enforced following each transient, which was found to be sufficient in explaining the observed behavior.

Figure 2-9 shows the basic circuit, highlighting the main branches through which current is divided after the constriction switches: the shunt resistor (Fig. 2-9(b)), the constriction itself (Fig. 2-9(c)), and the loop inductor (Fig. 2-9(d)). In a memory cell with an unshunted constriction, represented as having $R_s = 1$ MΩ, nearly all of the bias current is diverted to the loop inductor after a switching event, causing the maximum amount of current to be stored in the loop once the hotspot collapses and the constriction regains the superconducting state. In constrast, shunting the constriction with $R_s = 5$ Ω allows the majority of the bias current to be diverted instead to the resistor, minimizing the amount of flux trapped through the loop inductor. Figures 2-9b(i) – 2-9d(i) show the results over a longer timescale, where continuous switching of the shunted constriction brought on by a steadily increasing

write current ramp adds flux to the loop in increments of 5 Φ_0, thus confirming the experimental observations displayed in Fig. 2-8 for Device 1.

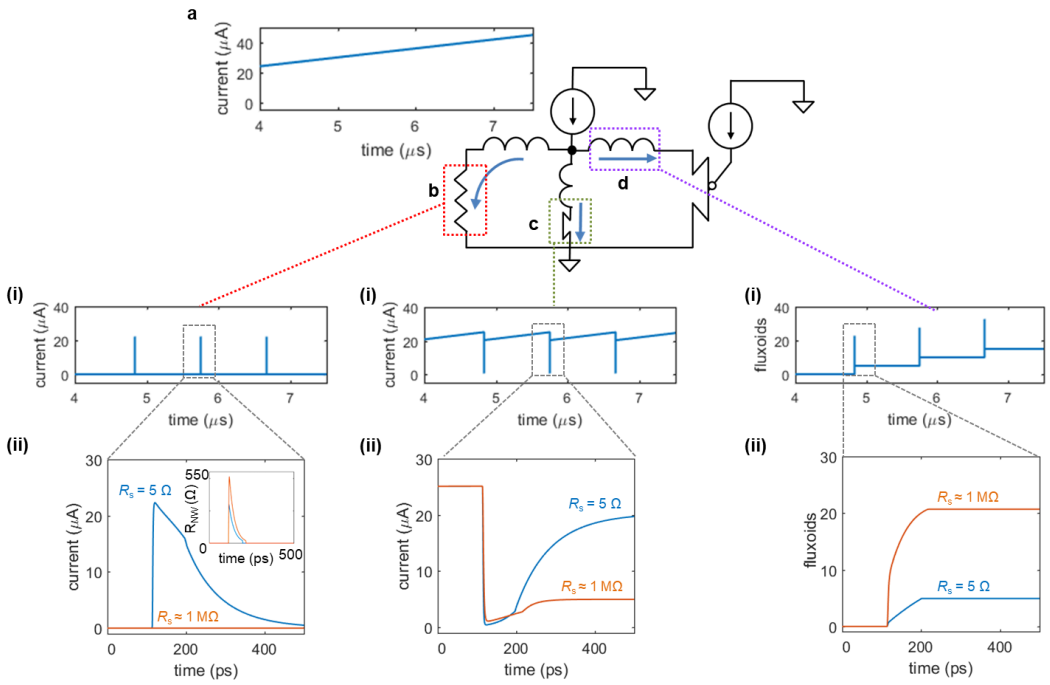

Figure 2-9: Time domain simulations of the circuit highlighting the three branches through which the bias writing current is diverted. (a) Bias current ramp sent to the write port of the device. (b) Current through the shunt resistor. (c) Current through the constriction. (d) Current through the inductor, represented in terms of trapped fluxoids. Panels labeled (i) show the simulation over a long time domain, while panels labeled (ii) show the diversion of current over a single switching event, with time on the x-axis shifted to start from $t = 0$.

Figure 2-10 displays the amount of trapped flux per switching event resulting from simulating devices of varying circuit parameters. Fig. 2-10(a) shows that the amount of flux increases proportionally with increasing shunt resistance and shunt inductance, which agrees with the experimentally observed shift caused by increasing R_s. Fig. 2-10(b) suggests a slightly more complex relationship between R_s and L_{loop}, with plateaus occurring due to the limitations on the maximum loop current with respect to the critical current of the constriction—for example, if $I_{c,NW} = 20$ μA, a loop inductance leading to a ratio of 2 μA of circulating current per fluxon cannot have more than 10 fluxons per switching event. Fig. 2-10(c) compares the simulated trends

for varying R_s with $L_{loop} = 1.85$ nH and $L_{loop} = 0.65$ nH to the experimental results for the three memory cells listed in Table 2.1. Data points representing the three measured devices show that the electrothermal simulations are in good agreement with the experimental measurements.

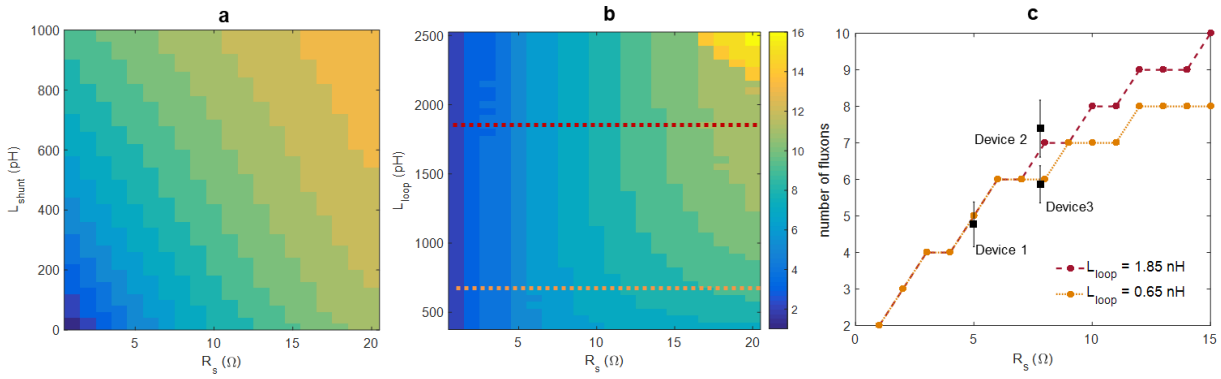

Figure 2-10: Simulated effect of circuit parameters on the amount of trapped flux. (a) Simulated number of trapped fluxons per switch as function of varying R_s and L_{shunt}. (b) Simulated number of trapped fluxons per switch as a function of varying R_s and L_{loop}. (c) Comparison of experimental and simulated number of trapped fluxons per switching event for $L_{loop} = 0.65$ nH and $L_{loop} = 1.85$ nH.

While both of the colormaps in Fig. 2-10 rely on the electrothermal parameters included in the simulation, an intuitive sense of their trends can be gained by simply calculating the trapped flux predicted by current division between the shunt resistor and the loop inductor after the nanowire switches. The total impedance of the shunt resistor pathway is the sum of the shunt resistance and the impedance of the series inductance, $Z_{shunt} = R_{shunt} + 2\pi f L_{shunt}$, while the impedance for the loop inductor is $Z_{loop} = 2\pi f L_{loop}$. The current charging the loop inductor, which dictates the amount of flux trapped after the loop heals, is roughly $I_{loop} \approx I_{c,NW} \times Z_{shunt}/(Z_{shunt} + Z_{loop})$. The total number of trapped fluxons is then $I_{loop} L_{loop}/\Phi_0$.

Figure 2-11 shows the result of varying R_{shunt} and L_{shunt}, and varying R_{shunt} and L_{loop} in the general expression for trapped flux above. For simplicity, $I_{c,NW}$ was set equal to 25 μA, and f was arbitrarily set to 1.8 GHz. The shapes of the trends correspond to those in the simulations of Fig. 2-10; however, it is clear that the electrothermal factors included in the more complex simulations of Fig. 2-10 are

needed to reflect the true timing dynamics of the nanowire and the thermal effects on the shunt resistor on the hotspot relaxation time. Nonetheless, the general shape of the trends is clearly a result of the current division between the two impedances.

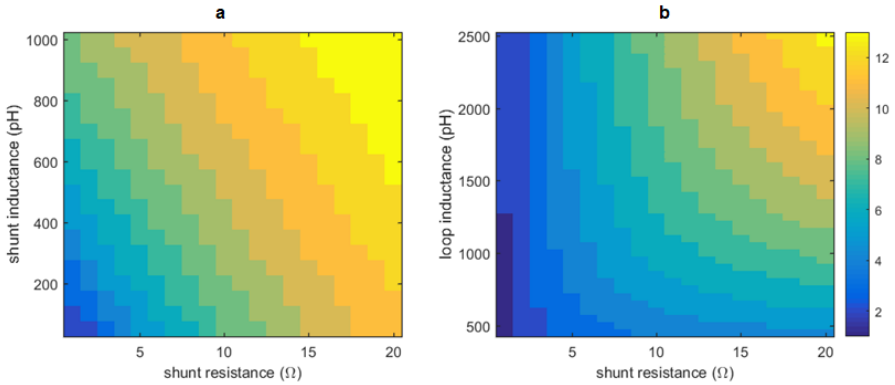

Figure 2-11: Trapped flux dependencies on circuit parameters calculated from simple bias current division. (a) Varying the shunt resistance and shunt inductance. The loop inductance was set constant at 1.87 nH. (b) Varying the shunt resistance and loop inductance. The shunt inductance was set constant at 50 pH.

2.5 Outlook and applications

The measurements and simulations shown in this chapter demonstrate that the output of this device may be tuned through simple circuit parameters and tailored to meet specific design requirements. As a result, the device is a promising platform for the development of a multilevel memory, with the number of states dictated by the critical current of the constriction and the amount of flux per switching event. While the devices shown here have a maximum of seven states, recent work building on this device showed that more levels could be achieved by adjustments like decreasing the shunt inductance [41], as suggested by the simulations in Fig. 2-10. This multilevel operation is significantly different from previously reported nanowire-based memories, which have thus far been binary devices [30][33][27]. While the proof-of-concept devices reported here had a maximum size of 3×25 μm^2, the device could be scaled down by introducing a high kinetic inductance wiring layer for the loop, given that the magnetic inductance is inconsequential. Further scaling improvements could be

made by fabricating the loop as a stacked structure, as was suggested with previous nanowire-based memories [30].

These devices have also recently been extended to applications beyond cryogenic memory. For instance, using the write port as a single photon detector allows the loop to act similarly to a CCD, where the stored loop current is proportional to the number of incident photons [42]. Additionally, the device has been incorporated into a crossbar architecture for deep neural networks [41], where it was shown that a nanowire memory with 33 discrete states can support image recognition with high testing accuracy. These examples highlight how the number of required states in the memory can vary depending on the application, illustrating the nanowire memory's utility as a diverse, tunable element that could be used in many multilevel operations.

Chapter 3

The nanowire neuron: design

In addition to memories like the device introduced in Chapter 2, computing architectures also require processing components to perform operations. As discussed previously, neural networks made from naturally spiking hardware are actively being explored for their potential for fast, energy-efficient computation. In these schemes, spikes serve as the tokens of information, while neurons act analogously to logic gates, producing a single output in response to a combination of multiple inputs. Synapses, which connect neurons, also play a crucial role in the overall architecture by serving as local memories and allowing networks to learn and adapt.

Superconducting nanowires are logical candidates for building spiking neural networks given their intrinsic ability to generate low-power pulses from the nonlinear switching between the superconducting and resistive states. Additionally, their compatibility with external silicon circuits could allow for integration of neural networks with more traditional systems in the initial phases of architecture development.

Here, we describe the design of an artificial neuron that uses the coupling between two shunted nanowires to generate spikes. Unlike the shunted nanowire in the memory cell of Chapter 2, the artificial neuron relies on nanowires with large enough series inductances to support relaxation oscillations. Using electrothermal circuit simulations, we demonstrate that the nanowire neuron can replicate multiple characteristics of biological neurons, and that kinetic inductance can be used to create a tunable synapse that is capable of supporting fan-out with an energy figure of merit four

orders of magnitude better than that of competing technologies.

The text that follows includes reprints from a work that was originally published in *Frontiers in Neuroscience* [43]. I would like to acknowledge Prof. Ken Segall for his guidance and advice throughout the development of this work.

3.1 The nanowire neuron model

In this section, we describe the nonlinear dynamics of relaxation oscillations in superconducting nanowires, and then present the architecture of the nanowire-based neuron.

3.1.1 Relaxation oscillations

The intrinsic nonlinearity of superconducting nanowires makes them ideal candidates for the hardware generation of spiking behavior. As discussed in Chapter 1, when a bias current flowing through a superconducting nanowire exceeds its critical current I_c, superconductivity breaks down and the nanowire becomes resistive, producing a voltage. The nanowire only switches back to the superconducting state once the bias current is reduced below the retrapping current I_r, and the resistive hotspot cools down. When the nanowire is placed in parallel with a shunt resistor, this switching process participates in electrothermal feedback with the shunt; if the series inductance between the nanowire and shunt is sufficiently large, the feedback produces relaxation oscillations [28].

Action potentials in biological neurons can also be viewed as relaxation oscillations. Figure 3-1 shows a basic schematic of an action potential in a biological neuron, highlighting the two main voltage-gated ion channels that control its dynamics. The electrochemical potential across a neuron's membrane initially sits at its resting potential ($\sim$ 70 mV), but rises in response to a sufficiently high stimulus. This increase in membrane potential triggers the opening of the sodium (Na+) ion channel, which causes Na+ ions to flood into the membrane, further increasing its potential (see Fig. 3-1(a)). Once the membrane potential rises to a new higher level, the potassium

(K+) ion channel opens up, and K+ ions rush out of the cell, restoring the potential back to its resting state (see Fig. 3-1(b)). Other dynamics also come into play, such as the sodium-potassium pump that expends energy to restore the charge imbalance by transporting three Na+ ions out of the cell for every two K+ ions into the cell (a process called "active transport"). However, for our purposes it is sufficient to just consider the simplified neuron model of an action potential generated by the two voltage-gated ion channels.

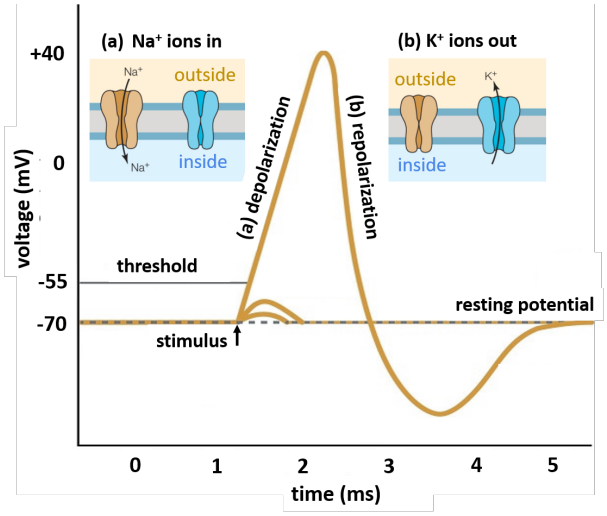

Figure 3-1: Schematic of an action potential in a biological neuron. The rising edge (a) is controlled by Na+ ion channel influx. The falling edge (b) is controlled by the K+ ion channel outflux, restoring the membrane potential to its resting state and allowing the neuron to fire again.

Like the Na+ influx and K+ outflux currents of a neuron, the influx and outflux currents from the nanowire to the shunt resistor are governed by different timescales, τ_1 and τ_2. As shown in Figure 3-2(a), the rising edge of the output voltage is defined by $\tau_1 = L/(R_s + R_{hs})$, where L is the inductance of the nanowire, R_s is the shunt resistance, and R_{hs} is the resistance of the nanowire hotspot, usually on the order of $\sim$1–10 kΩ. Conversely, the outflux current occurs when the nanowire is no longer resistive and the bias current is redirected from the shunt; this reduced resistance results in a slower time constant $\tau_2 = L/R_s$, which defines the falling edge of the output voltage. For typical nanowire devices, $\tau_1 \sim 100$ ps and $\tau_2 \sim 1$ ns. The two currents are "gated," as shown by the insets in Fig. 3-2(a), by the state of the

nanowire—when the state is resistive (producing a voltage), the influx current flows into the shunt, and when it is superconducting, the outflux current flows back into the nanowire.

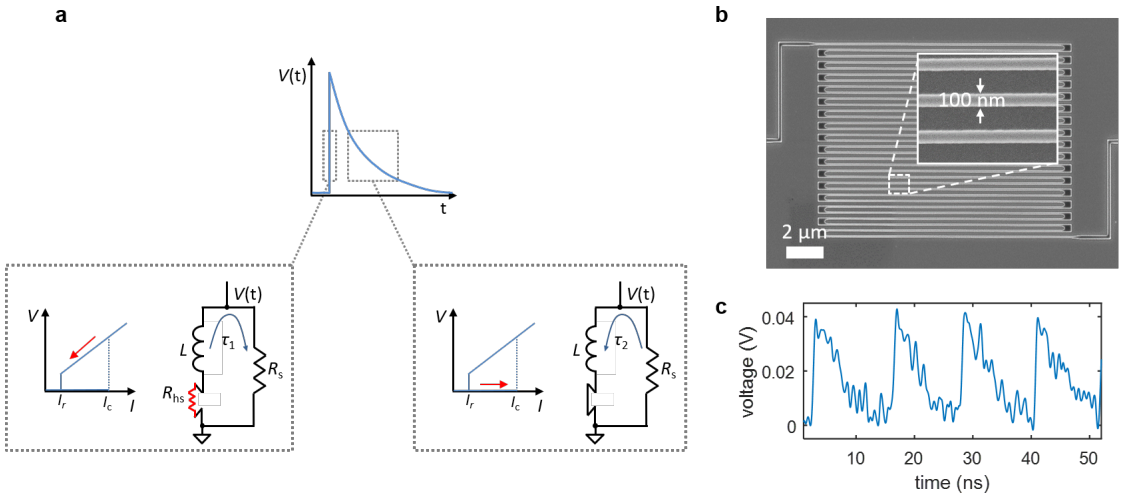

Figure 3-2: Relaxation oscillations in superconducting nanowires. (a) Simplified model of a relaxation oscillation, specifying the two time constants that define the output voltage. (b) Example of a superconducting nanowire in a long, meandered design for obtaining a high kinetic inductance. (c) Experimentally measured relaxation oscillations in a long superconducting nanowire shunted by 50 Ω.

Since kinetic inductance is defined per unit length of a nanowire with a fixed width, a long nanowire can be used to create a sufficiently high series inductance between a constriction and shunt resistor to support relaxation oscillations. Figure 3-2(b) shows a scanning electron micrograph of a typical superconducting nanowire with a meandering geometry designed for maximizing the total device inductance. An example of experimentally observed relaxation oscillations for such a device is displayed in Figure 3-2(c).

3.1.2 The basic model

Although a shunted nanowire on its own produces oscillations analogous to action potentials, the output signal eventually accumulates a voltage offset as the bias current increases. This effect deviates from true neuron behavior, since the cell must maintain a constant resting potential as described above. To overcome this difference,

we have implemented a neuron architecture based on one that was first proposed for Josephson junctions [44], as shown in Figure 3-3. The circuit consists of two shunted nanowires, referred to as the main oscillator and control oscillator, linked together in a superconducting loop. A bias current I_{bias} is applied to both oscillators such that they are each biased right below their critical currents, but in opposite directions. To trigger an action potential, a small input current pulse I_{in} (Fig. 3-3(a)) is applied that sums with the bias current to exceed I_c of the main oscillator, causing it to switch (Fig. 3-3(d)). The control does not fire since the input opposes the direction of its bias.

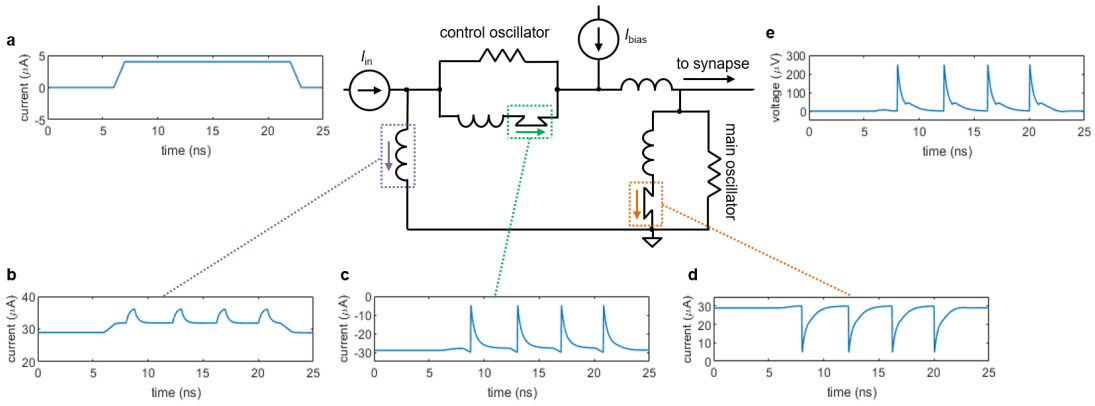

Figure 3-3: Circuit simulations of the two-nanowire soma. (a) Input pulse. (b) Current through the loop inductor. (c) Current through the control nanowire. The control nanowire reduces the amount of counterclockwise current circulating in the loop, allowing the main nanowire to fire again. (d) Current through the main nanowire. (e) Output voltage pulse that is sent to the synapse. For these simulations, the critical current of the control nanowire and of the main nanowire is $I_c = 30$ μA, the bias current $I_{bias} = 58.6$ μA, and the input current $I_{in} = 4$ μA.

Once the main oscillator switches, current is added to the superconducting loop in the counterclockwise direction (Fig. 3-3(b)), which sums with the bias current to fire the control oscillator (Fig. 3-3(c)). The control oscillator removes counterclockwise current from the loop, allowing the main oscillator to fire again. Without the presence of the control oscillator, the main oscillator would only be able to fire once, since the counterclockwise current added to the loop would reduce the total current through the nanowire of the main oscillator below its I_c. The voltage from the main oscillator

node (Fig. 3-3(e)) is a superposition of the two firing oscillators and serves as the spiking output that is carried down to the next neuron via the synapse. Unlike the output from the single shunted nanowire, the output of the two-nanowire circuit does not accumulate a bias offset, making it a suitable spiking signal.

In the context of the two-channel simplified neuron model, the main oscillator acts analagously to the Na+ influx current by adding flux to the superconducting loop in the form of a circulating current. The control oscillator acts analagously to the K+ outflux current by reducing the circulating current, resetting the neuron and allowing the main oscillator to fire again. As described in Ref. [28], the rate at which each oscillator fires depends on the magnitude of the bias current, paralleling the voltage-dependent rate constants of ion gates in the Hodgin-Huxley model for a biological neuron [45].

3.2 Single neuron characteristics

Neurons display a wide variety of traits unique to certain populations, allowing them to achieve varied and complex tasks when they interact collectively. While no single neuron possesses all possible traits, the basic functionality of an artificial neuron can be evaluated by demonstrating some common bio-realistic characteristics. Here we present multiple neuron behaviors that can be achieved with the nanowire neuron, using electrothermal circuit simulations conducted in LTspice. The simulations implement material-specific characteristics and nanowire hotspot dynamics as described in previous literature [31][40], and have been shown to reliably reproduce experimental data pertaining to nanowire relaxation oscillations [28].

3.2.1 Threshold response

A general characteristic of biological neurons is their inability to fire unless the input signal exceeds a certain threshold. Figure 3-4(a) shows the threshold voltage response of the nanowire neuron when the bias current is held constant and the input current is varied. As evident in the plot, the neuron does not begin firing until the input

current passes a threshold, defined by when the sum of the bias and input current through the main nanowire exceeds its I_c. Above the threshold, the peak voltage of the spike output is essentially constant. However, biological neurons have a threshold that may be varied by previous activity, such as an inhibiting input that reduces it [46]. Figures 3-4(b) and (c) illustrate this process ("threshold variability") in the nanowire neuron; an initial subthreshold input pulse fails to elicit a spike, while a later input pulse of the same magnitude triggers a spike after a smaller negative pulse reduces the firing threshold. It should be noted that when the preceding pulse was of the opposite polarity, no spike was triggered. This behavior is consistent with the expectations of Ref. [46].

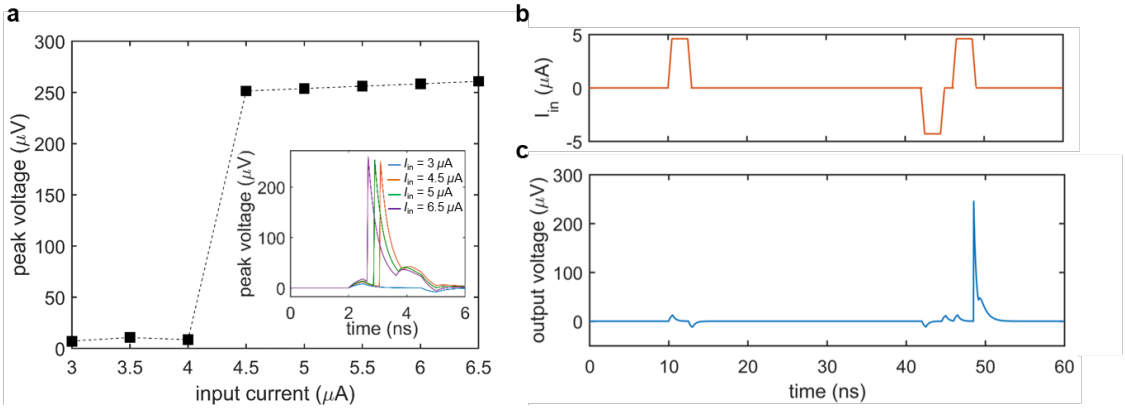

Figure 3-4: Firing threshold of the nanowire neuron. (a) Peak output voltage as a function of input current under a constant bias (I_{bias} = 58.6 μA). Inset shows the time domain voltage output of the neuron for different input currents. (b) Input to the neuron, leading to a reduction in firing threshold by a preceding negative pulse. (c) Spiking output of the neuron in response to the inputs of (b), demonstrating that the nanowire neuron's firing threshold is variable. For this simulation, I_{bias} = 57.62 μA, the positive input I_{in} = 4.6 μA, and the negative input I_{in} = -4.3 μA.

3.2.2 Refractory period

In addition to exhibiting a firing threshold, the nanowire neuron displays a refractory period, which we define to be the minimum time between two input pulses such that both pulses elicit a spike. Figure 3-5 illustrates this response. When two pulses are separated enough in time so that the main oscillator is biased close to its critical

current when the second input pulse arrives, then the second pulse will cause a spike (Fig. 3-5(a)). However, if the second pulse arrives before the bias current has fully returned to the main oscillator, then the sum of the second input pulse and the bias will not be sufficient to switch the nanowire and trigger the neuron (Fig. 3-5(b)). As a result, the refractory period is limited by the time it takes to fully bias the main oscillator again, which is a function of the L/R time constants of the circuit.

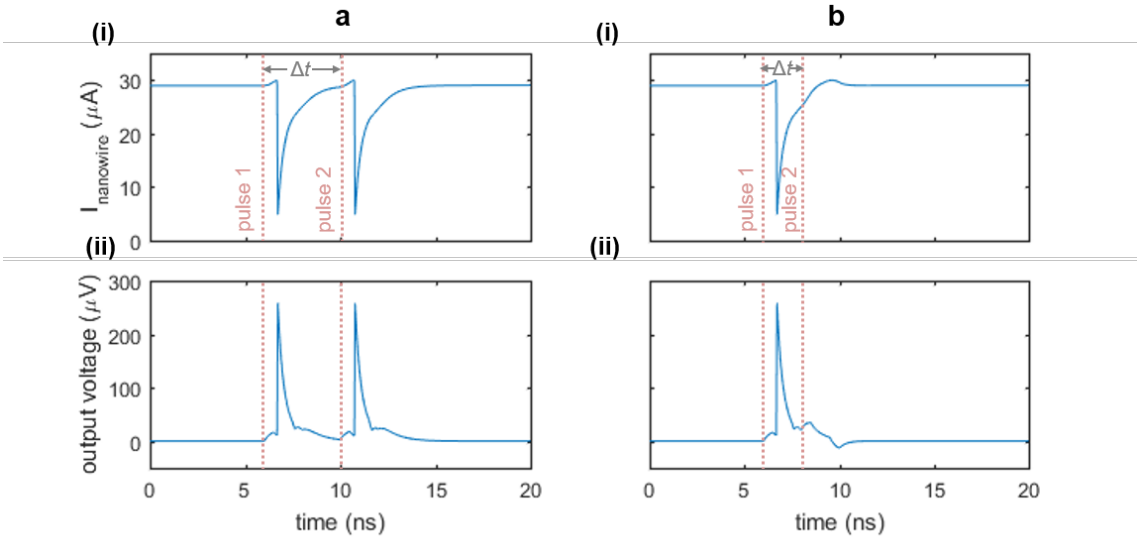

Figure 3-5: Refractory period of the nanowire neuron. (a) Response when there is sufficient time between two inputs to each elicit a separate spike. Parameters: I_{bias} = 58 μA, I_{in} = 6 μA, Δt = 4 ns. The pink dashed lines indicate the beginning of the rising edge of each pulse. (b) Response when there is insufficient time between two input pulses, causing the neuron to fire only one. Parameters are the same as in (a), except Δt = 2 ns. For both cases, panel (i) displays the current through the nanowire of the main oscillator, while panel (ii) displays the output voltage of the neuron.

3.2.3 Class I behavior

Biological neurons differ in their response to varying signal strengths. Whereas Class I neurons have a spiking frequency that increases with increasing input strength, Class II neurons maintain a constant firing rate [46][44]. Figure 3-6 illustrates the spiking behavior of the nanowire neuron at different levels of bias current. Figure 3-6(a) shows the time-domain voltage output of the neuron as the bias current is increased, and suggests an increase in spiking frequency. This response is confirmed by observing

the voltage output's frequency spectrum (Fig.3-6(b)), which shows a shift in the spiking frequency to higher values with increasing bias. Consequently, the nanowire neuron has Class I behavior. Although this effect can also be achieved by changing the input current, as shown by the frequency shift in Figure 3-7, the modulation is less pronounced. Nonetheless, the change in spiking frequency by the bias and input currents demonstrates that the frequency of the nanowire neuron output may be used to glean information about its input conditions.

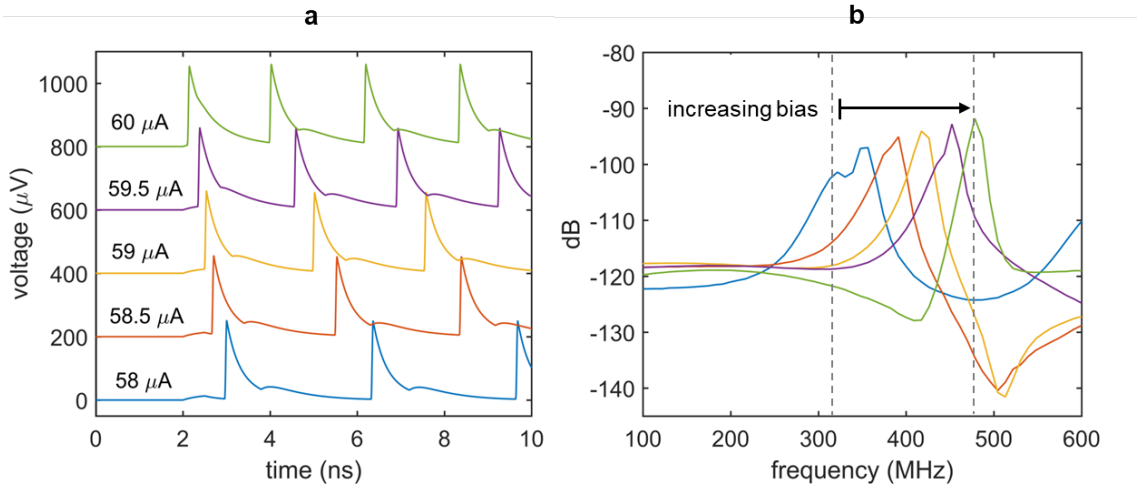

Figure 3-6: Effects of bias current on spiking frequency. (a) Time domain simulations of the nanowire neuron with different bias currents ($I_{in} = 6$ μA for all simulations). Traces have been shifted from one another in the y-axis for clarity. (b) Fourier transform of the voltage output for each biasing condition. The shift in peak frequency with bias current indicates that the circuit acts like a Class I neuron.

3.2.4 Parabolic bursting

Some neurons, such as thalamic and dopaminergic neurons [47], display a unique mode of behavior called bursting, in which the cell alternates between the resting and firing states. The transition between states may be dictated by slow changes in low levels of intracellular calcium ions, which influence the conductance of K+ [48]. As a result, a small, slowly varying signal controls the rapid dynamics of the action potential, leading to alternating periods of resting and firing. This process is considered to be an important aspect of electrical activity in the brain.

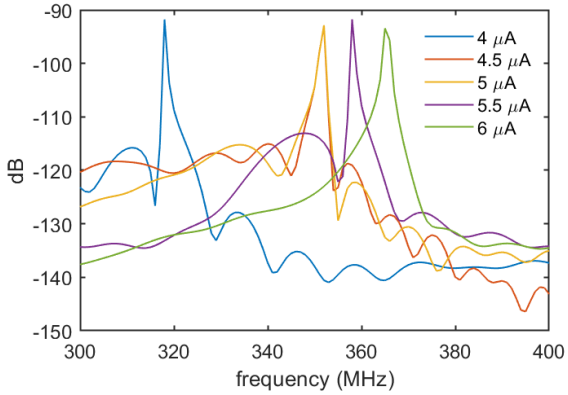

Figure 3-7: Effects of input current on spiking frequency. The Fourier transform of the voltage output for each input current shows that the spiking frequency increases with increasing signal strength. For all simulations, $I_{bias} = 58$ μA.

Due to the significance of bursting in biological neurons, past work has sought to replicate similar behavior in platforms such as digital silicon models [49] and Josephson junction models [50] by injecting a low-frequency ac signal into the system. Figure 3-8 shows the result when a similar technique is applied to the nanowire neuron. In this case, a weak ac signal ($f = 50$ Hz, $I_{ac} = 4$ μA) shown by the dashed red line in Fig. 3-8(a) is coupled into the bias port of the neuron, causing the neuron to alternate between the resting and firing states, as reflected in the resulting output voltage signal. A close examination of the timing between adjacent spikes (see Fig. 3-8(b)) shows that the spiking frequency increases and then decreases over the firing period, a phenomenon known as parabolic bursting [51]. This behavior was first observed in neuron R15 of the abdominal ganglion of Aplysia [52] [53], and has since been demonstrated in many other cells. The ability of the nanowire neuron to replicate similar dynamics may be therefore be useful for performing a wider range of functions.

3.2.5 Axon: transmission line characteristics

After an action potential occurs in a biological neuron, the output signal propagates down the axon as if sent through a delay line [54]. This delay is valuable in that it preserves time domain information, potentially facilitating behaviors that rely on the recognition of specific spatio-temporal patterns [46]. Such pattern recognition is

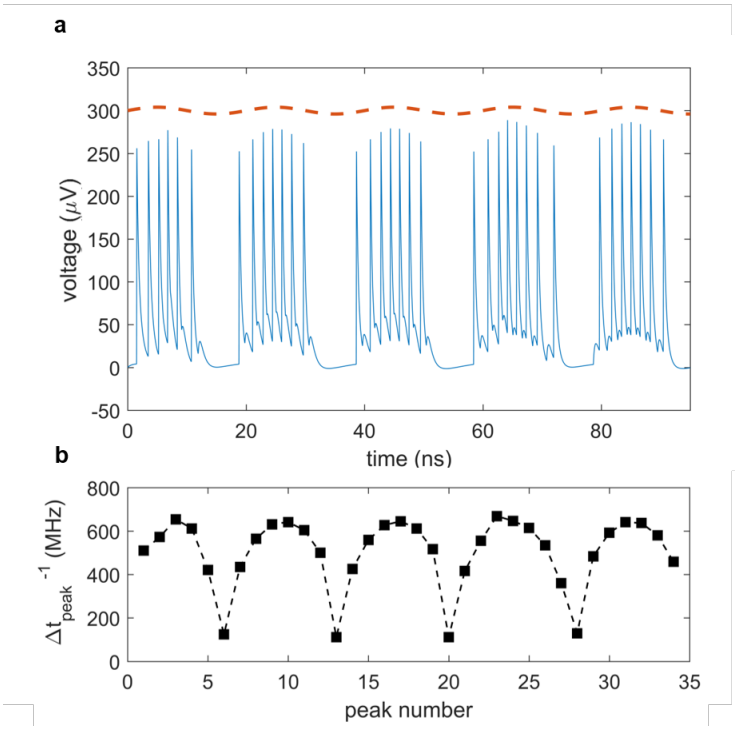

Figure 3-8: Parabolic bursting in the two-nanowire neuron. (a) Output voltage of the nanowire neuron when the bias is coupled to a weak sinusoidal drive, indicated by the red dashed curve and shifted in the y-axis for clarity. (b) The inverse of the time between adjacent peaks shows that the time difference follows a parabolic form. Parameters: $I_{bias} = 59$ μA, $I_{in} = 6$ μA.

often not possible in spiking neural networks with traditional wiring, since signals travel too rapidly for time information to be maintained [54]. This is not the case for superconducting nanowires that are designed to act as transmission lines. Recent work has shown that the high kinetic inductance of superconducting transmission lines, like those made out of niobium nitride, results in propagation speeds of $\sim 2\% c$, where c is the speed of light in vacuum [55]. For example, a typical 300-nm-wide NbN microstrip with a thickness of 5 nm has a kinetic inductance of roughly $L_k = 2.7 \times 10^{-4}$ H/m or $212\mu_0$, where μ_0 is the permeability constant ($\sim 12.57 \times 10^{-7}$ H/m), and a capacitance of roughly $C = 1.9 \times 10^{-10}$ F/m or $21\epsilon_0$, where ϵ_0 is the permittivity constant ($\sim 8.85 \times 10^{-12}$ F/m). The resulting velocity is then $v = 1\sqrt{L_k C} = 4.4 \times 10^6$ m/s or approximately 1.5% c. For comparison, the typical propagation speed of a signal through a coaxial cable is about 66% c.

As illustrated in Figure 3-9, a simulated nanowire neuron output sent through a

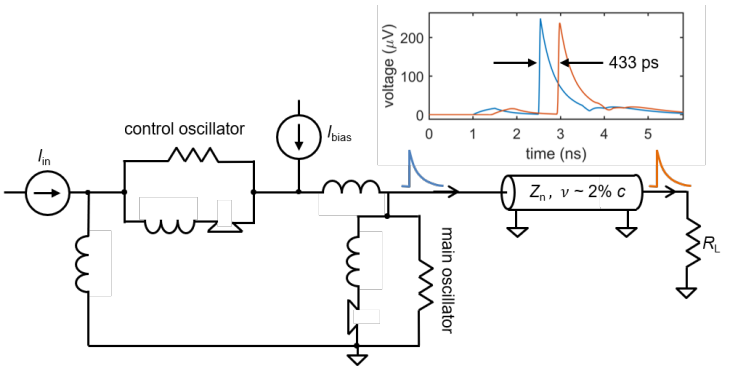

Figure 3-9: A superconducting transmission line as an axon. Simulations of a superconducting transmission line show that the spikes can be delayed on the order of ~ 5 ns, depending on the length of the structure. This delay could enable the preservation of timing information. Transmission line parameters: $L_k = 0.3$ nH/μm, $C = 0.1$ fF/μm, $v = 1.9\%$ c, transmission line length $l = 2.5$ mm. Shorter transmission lines on the order of 800 μm still had delays of $\sim$ 140 ps.

superconducting transmission line model is delayed by ∼100–500 ps, close to the full width of an action potential. In mammalian brains, axonal delays like the cortico-cortical delay [56] are also on the same timescale as the full width of an action potential [57](typically a few milliseconds), suggesting that the relative delay in our system with respect to the spike duration is appropriate. If longer delays are needed, the transmission line can simply be made longer.

It is worth noting that, just as not all biological neurons have meter-long axons with delays lasting tens of milliseconds, not all nanowire neurons need to have millimeter long transmission lines. As with biological systems, whether the temporal delay is needed depends on the application. While some spiking neural networks strive to be as bio-realistic as possible in order to shed light on how our minds work, others are applied to problems like image recognition, where (so far) time information has not been used. Our goal is to illustrate that nanowires are a flexible platform for achieving either type of behavior, and that they offer a way of slowing down pulses to preserve time information if the application requires it.

3.3 The synapse

The collective dynamics of a neural network depend on the ability of a neuron to influence the behavior of another downstream neuron via a synapse. In biological systems, fast electrical action potentials trigger the slow release of chemical neurotransmitters that bind to receptors on a downstream neuron, causing ion channels on its membrane to open or close. In the excitatory case, the binding of neurotransmitters causes the downstream neuron to fire by raising its membrane potential. Conversely, neurotransmitters can have an inhibitory effect where they make the target neuron's membrane potential even more negative, making it less likely to fire.

In previous work using Josephson junctions, a slow synaptic response was achieved through the gradual charging and discharging of a large capacitor [44]. The same approach can be taken with the nanowire neuron. Figure 3-10 shows an example of a capacitive synapse. In the case of excitatory control, the upstream (main) neuron is positively biased, accumulating a positive voltage on the synapse capacitor (see panel (ii)). When the capacitor discharges, a positive current then flows into an underbiased downstream target neuron, causing it to fire. To inhibit an overbiased downstream neuron, the main neuron can be biased in the negative direction, accumulating a negative voltage on the capacitor that suppresses the firing of the target.

3.3.1 Inductive synapse

The same effect of a slowly charging and discharging capacitive synapse can be achieved with an inductor by creating an approximately equivalent circuit, as shown in Figure 3-11. To get a rough idea of the magnitude of the synapse inductance, the input impedances and time constants of both circuits can be compared in the extreme limits of $\omega \to 0$ and $\omega \to \infty$, where ω is the frequency (see Appendix A for complete derivation). We found that it was also necessary to add a series resistance to the left of the synapse inductor in order to prevent backflow into the main neuron and to keep the input impedances of the two circuits equivalent at low frequencies. The estimated synapse inductance ($\sim$0.1 μH) was two orders of magnitude higher than

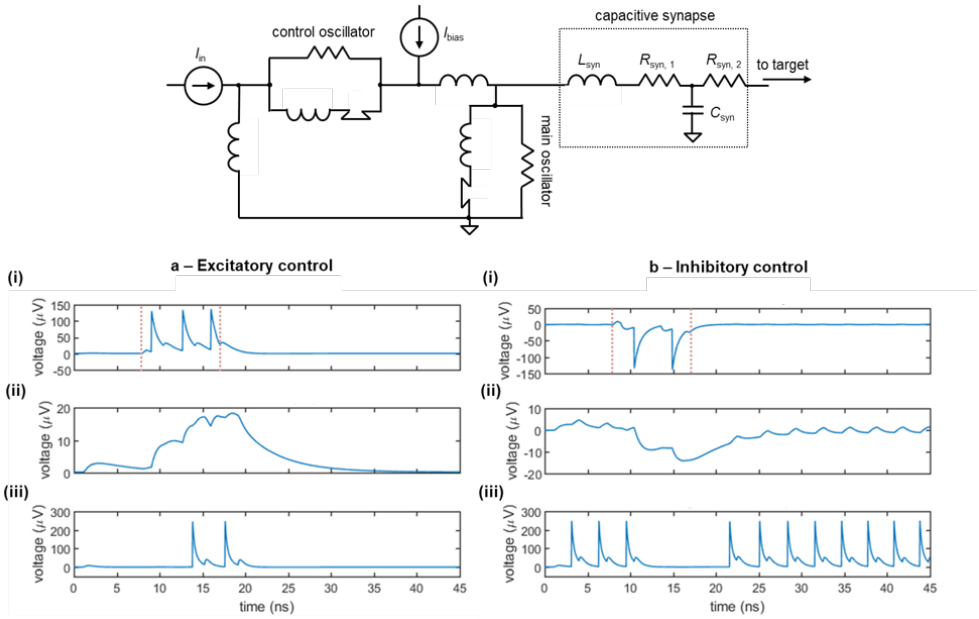

Figure 3-10: The nanowire neuron using a capacitive synapse. (a) Excitatory control. (b) Inhibitory control. Panel (i) shows the output of the main neuron. Panel (ii) shows the voltage on the synapse capacitor. Panel (iii) shows the output of the downstream target neuron. Parameters: $L_{syn} = 0.1$ pH, $C_{syn} = 1$ nF; $R_{syn,1} = 10\ \Omega$, $R_{syn,2} = 10\ \Omega$.

the equivalent synapse capacitance (∼1 nF), which was to be expected by setting the two time constants equal to one another ($R_{C,syn}C = L/R_{L,syn}$).

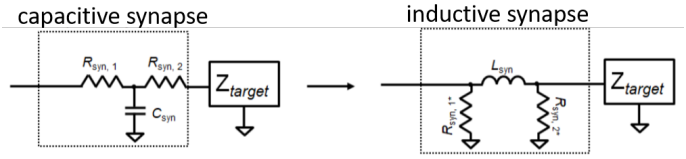

Figure 3-11: Transforming a capacitive synapse into an inductive synapse.

Figure 3-12 shows simulations of the inductive synapse. Similar to the slow release of neurotransmitters in response to an action potential, the inductive synapse relies on the slow charging of a large inductor in response to the nanowire neuron's more rapid voltage spikes. The energy stored in the large synapse inductor is then discharged as current into the input port of the target neuron, modulating its behavior. Like the capacitive synapse, the polarity of the bias to the main neuron dictates whether the control is excitatory (Fig.3-12(b)) or inhibitory (Fig.3-12(c)).

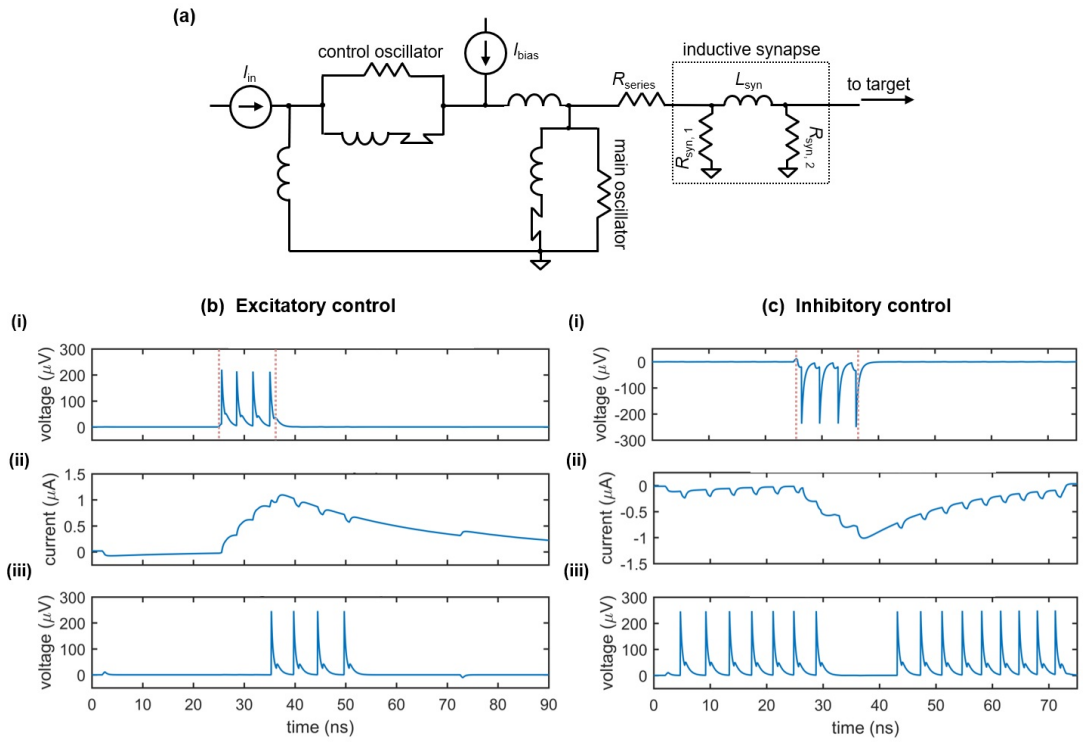

Figure 3-12: Nanowire neuron with an inductive synapse. (a) Circuit schematic. (b) Excitatory downstream control. Parameters: $L_{syn} = 0.265$ μH,$R_{syn,1} = 40$ Ω, $R_{syn,2} = 40$ Ω, $R_{series} = 14$ Ω. (c) Inhibitory downstream control.Parameters: $L_{syn} = 0.23$ μH,$R_{syn,1} = 45$ Ω, $R_{syn,2} = 40$ Ω, $R_{series} = 24$ Ω. Panel (i) shows the output of the main neuron. Panel (ii) shows the voltage on the synapse capacitor. Panel (iii) shows the output of the downstream target neuron.

3.3.2 Variable synaptic strength

Although the inductive synapse of Figure 3-12 can be engineered for both excitatory and inhibitory control of a downstream neuron, a fundamental property of artificial neural networks is the ability to modulate that control by adjusting synaptic strength. On an even more fundamental level, the connections between biological neurons are constantly changing whenever a behavior is learned or forgotten. It is therefore critical to consider ways in which synaptic strength can be adjusted for changing the control between two neurons.

One possible scheme for implementing such variability in the inductive synapse is to incorporate superconducting nanowires as a different circuit element: a tunable inductor. A nanowire's kinetic inductance increases with increasing bias current,

reaching an enhancement of 10–20% near I_c [58]. This modulation has been included in the circuit model of the superconducting nanowire used in these simulations [40]. By placing a high inductance nanowire with an ideal current source in parallel with the synapse inductor, the overall parallel inductance of the synapse can be modulated, which in turn changes the amount of current sent to the target neuron.

Figure 3-13 shows the simulated results for the case of an inhibitory synapse. When a higher modulating current I_{mod} is applied to the nanowire inductor, the overall parallel inductance increases, reducing the amount of current sent to the target. It is important to note that the polarity of the modulating current is not important, since the change in kinetic inductance depends only on the magnitude of the modulating current in relation to the nanowire's I_c. Figure 3-13(b) illustrates that the modulation in synaptic current for I_{mod} = 5 μA and I_{mod} = -5 μA is roughly the same. As a result, it is clear that the modulation is due to the change in kinetic inductance, and not simply the injection of current by I_{mod} in the opposing direction.

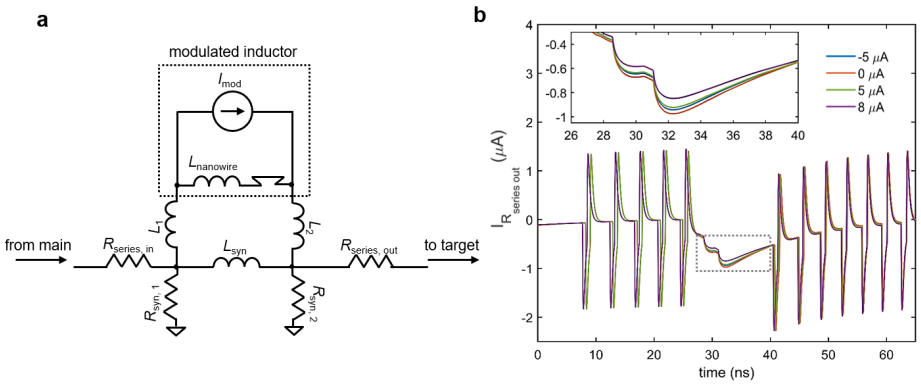

Figure 3-13: Modulation with the inductive synapse using a nanowire as a tunable inductor. (a) Circuit schematic. (b) Simulation of the current through $R_{series,out}$ in an inhibitory synapse as a function of different modulation currents. Spikes represent backaction from the firing of the target neuron. Parameters: L_{syn} = 0.45 μH,$L_{nanowire}$ = 0.275 μH,, $I_{c,nanowire}$ = 6 μA, $R_{syn,1}$ = 39 Ω, $R_{syn,2}$ = 40 Ω, $R_{series,in}$ = 25 Ω, $R_{series,out}$ = 0.1 Ω, $L_1 = L_2$ = 50 pH.

Figure 3-14 shows the result of tuning the inductive synapse when it is being used for excitatory control. By increasing the modulating current, the amount of current sent to the downstream target is reduced. When I_{mod} = 8 μA, the synaptic current

is no longer sufficient to trigger the downstream neuron, preventing it from firing.

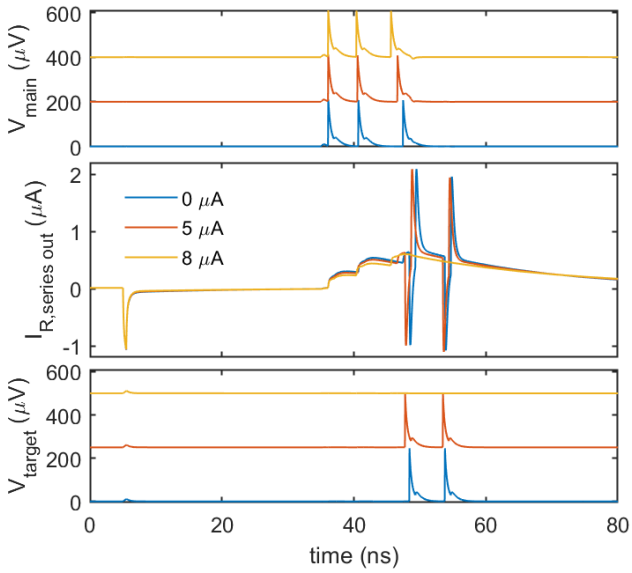

Figure 3-14: Example of the tunable inductive synapse with excitatory control. (Top panel) Voltage output of the main neuron. Traces have been shifted in the y-axis for clarity. The legend indicates the different values of I_{mod}. (Middle panel) Current through $R_{series,out}$ to the target neuron. (Bottom panel) Voltage output from the target neuron. Traces have been shifted in the y-axis for clarity. When $I_{mod} = 8$ μA, the synaptic current is too low to excite the target neuron. Parameters: $L_{syn} = 0.45$ μH, $L_{nanowire} = 0.275$ μH, $I_{c,nanowire} = 6$ μA, $R_{syn,1} = 39$ Ω, $R_{syn,2} = 40$ Ω, $R_{series,in} = 25$ Ω, $R_{series,out} = 5$ Ω, $L_1 = L_2 = 100$ pH.

In conjunction with the plasticity of synapses, the high degree of parallelism in the brain creates a densely connected, adaptable network able to optimize and adjust for different conditions. An example of fan-out in the nanowire neuron is shown in Figure 3-15. As shown in the circuit schematic, a single neuron is connected to four target neurons through four separate tunable synapses. When each of the four modulating currents is set to $I_{mod} = 0$ μA, the firing of all four targets is inhibited by the main neuron (Fig. 3-15(b)). To weaken the connection with one of the targets, in this case target #4, $I_{mod,4}$ is set to 8 μA, turning off the inhibiting action on target #4 but allowing inhibition to remain on the other three target neurons (Fig.

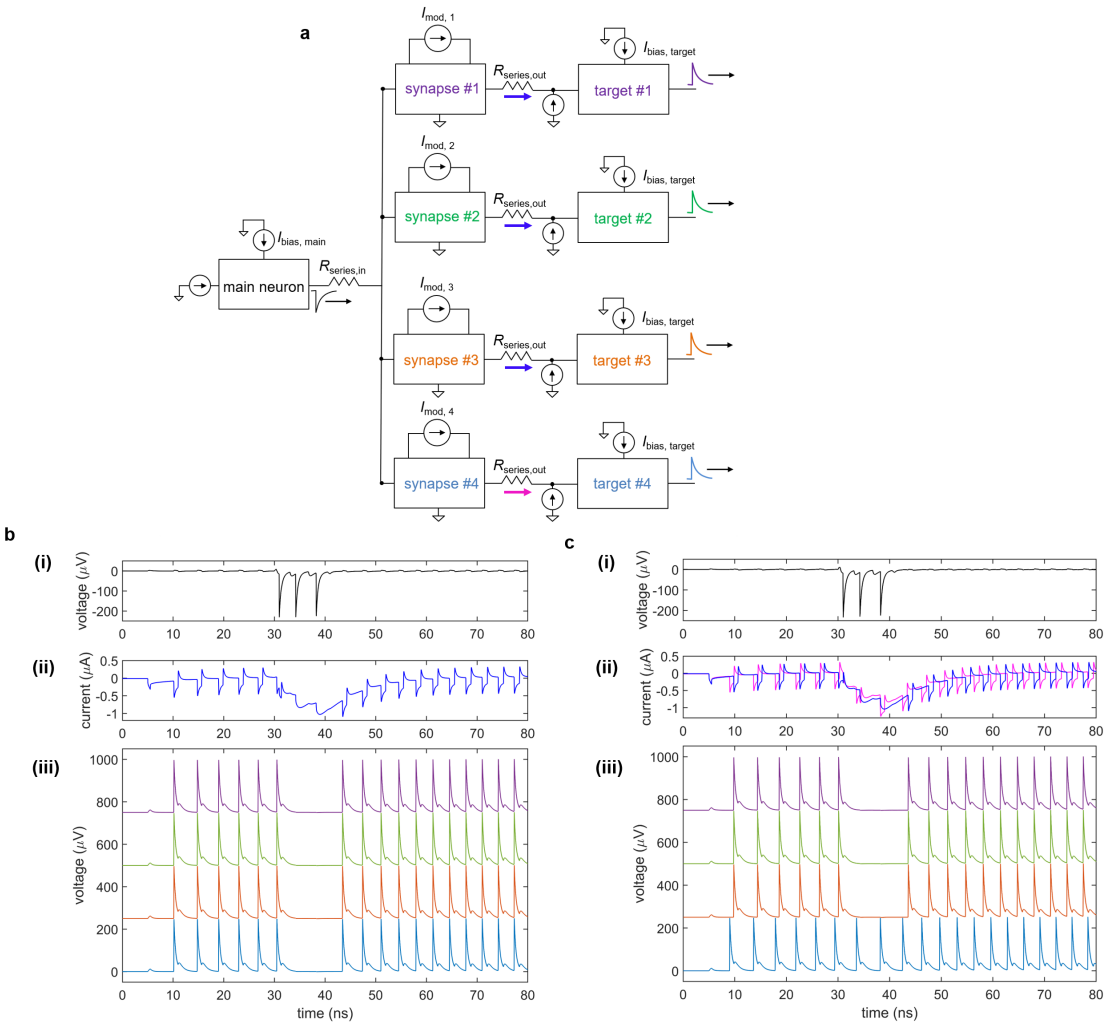

Figure 3-15: Fan-out of the nanowire neuron with a tunable inductive synapse. (a) Simplified circuit model. (b) Simulation when all modulating currents are turned off for the four target neurons. (c) Simulation when $I_{mod,1-3} = 0$ μA and $I_{mod,4} = 8$ μA. Panel (i) shows the output of the main upstream neuron. Panel (ii) shows the current through each of the four series resistors $R_{series,out}$ to each of the target neurons. Panel (iii) shows the output voltage of the four target neurons, shifted on the y-axis for clarity. Parameters: $L_{syn} = 0.4$ μH, $L_{nanowire} = 0.275$ μH, $I_{c,nanowire} = 6$ μA, $R_{syn,1} = 300$ Ω, $R_{syn,2} = 300$ Ω, $R_{series,in} = 7$ Ω, $R_{series,out} = 1$ Ω, $L_1 = L_2 = 100$ pH.

3-15(c)). Comparison of the synaptic currents for each of the four targets shows that the synaptic current for target #4 is reduced as a result of the modulated inductance. This modulation demonstrates that the nanowire neuron is able to be used in a parallel network where the strength of individual synaptic connections can be adapted without

disrupting the rest of the circuit.

Fan-out may be one area where nanowires will be an improvement over Josephson junctions [44]. Both nanowires and JJs have quantized flux outputs (flux here being defined as the time integral of the voltage). However, Josephson devices have outputs of only a single flux quantum, while nanowires typically have outputs with many more flux quanta (e.g. 70 flux quanta for the device shown in Figure 3-3). For instance, typical niobium nitride nanowires with critical currents in the range of 100 μA and film thicknesses around 10 nm have flux outputs ranging from 50 to 100 flux quanta. In pushing the fan-out and fan-in to larger systems, one expects to be eventually limited by parasitic inductances and thermal noise. In such a case, the more substantial signal of the nanowire neuron would permit a larger fan-out and fan-in, leading to a higher degree of parallelism. In addition to fan-out, the nanowire voltage signals are long enough and large enough to be digitized directly on an oscilloscope, in contrast to their JJ counterparts, allowing for more direct analysis and readout. Finally, power-control devices like the nTron [23] could be fabricated alongside the neurons and used to boost signals and match impedances.

3.3.3 Benchmarking synaptic energy and speed

The energy dissipation of both the neuron and the synapse can be calculated using LTspice by taking the time integral of the current-voltage product of each circuit element. Performing this analysis, we find that the nanowire neuron as an energy dissipation of about 0.05 fJ for each action potential, of which the synapse contributes less than 0.005 fJ (or less than 10%). In large systems, the synapses will dominate in comparison to the neurons; typically if there are O(N) neurons, there will be O(N^2) synapses. Hence, even though the neuron dissipates more energy, the synapses will dominate the power consumption of a large network.

In a spiking neural network, the energy dissipated increases with the speed of the system; spiking twice as often dissipates twice the energy, assuming the energy per spike is constant. As a result, the common figure of merit to compare different technologies takes the ratio of speed and power. IBM [59] introduced the figure of merit

synaptic operations per second per watt (SOPS/W). For our system, we also include a constant factor of about 400 W/W to account for the 4.2 K cooling costs in the nanowire neuron. Table 3.1 compares both the energy per spike and the SOPS/W for the nanowire neuron, the human brain, and two CMOS technologies. We acknowledge that the estimate for the nanowire neuron is a projection from a simulation of a single component, whereas the other entries in the table have actually been measured on large systems. However, the comparison shows that the nanowire neuron has the potential to be a competitive technology from an energy perspective, especially when one accounts for additional benefits like the lack of static power dissipation in superconducting interconnects. As discussed in the Introduction, large-scale superconducting circuits using current-biased devices struggle with dissipation costs from the resistors that deliver the bias current. As nanowire neurons are expanded into larger networks, it will therefore be critical to investigate alternative biasing techniques like the inductive distribution schemes used to minimize power dissipation in JJ logic circuits. In this case, the biasing network must also be made compatible with the higher impedances of our nanowire devices.

Table 3.1: Neural network energy comparisons

	Human brain	NeuroGrid	TrueNorth	Nanowire neuron
Energy/switch	10 fJ	100 pJ	25 pJ	0.05 fJ
SOPS/Watt	1e14	1e10	4e10	5e14

Energy values of the human brain, NeuroGrid, and TrueNorth are taken from Ref.[60].

3.3.4 Discussion on possible synapse alternatives

Although the fan-out achieved here is far from the level of parallelism in the human brain, where each neuron connects to thousands of neighbors, it suggests that nanowires may serve a unique purpose in the development of future superconducting neural networks, which have thus far struggled to support fan-out to more than one or two Josephson junctions. The maximum fan-out of the nanowire neuron may be further increased by changes like increasing the critical currents of the oscillators in order to increase the overall synaptic current, or using a tree structure. Given that

nanowires can also interface with both CMOS and Josephson junction circuits [26], it may be possible for nanowire neurons to serve as intermediary devices in a network with both platforms. A recently proposed neural network with hybrid technologies employed superconducting nanowires as photon detectors, relying instead on optical signals for facilitating high fan-out [61] [62]. Although our work uses nanowires solely as electrical components, they can easily be biased to act as photon detectors [21] [22] as well, illustrating that the two different architectures would be compatible for integration. These two schemes thus demonstrate the diverse ways in which superconducting nanowires can be used in neural networks, suggesting that a combination of the two technologies may be possible.

Furthermore, while the inductive synapse proposed in this chapter relies on an external modulating bias current, it may be possible to use a superconducting memory cell like that proposed in Chapter 2 to store the value of this modulating current, or replace the bias source directly in the circuit. Incorporating these programmable memory cells [32] into the modulating elements of the inductive synapse could enable storage of a varying synaptic strength in each synapse. It may also be possible for nanowire neurons to connect through alternative synapse designs, such as one based on inductive coupling. Finally, a synapse that implements firing-based tunability would allow the nanowire neuron to be applied to many more complex applications involving unsupervised learning.

3.4 Outlook

By taking advantage of intrinsic nonlinearities in superconducting nanowires, we designed a platform for a low-power artificial neuron where the coupling of two nanowire-based oscillators acts analagously to the two ion channels in a simplified neuron model. Using electrothermal circuit simulations, we showed that the nanowire neuron is able to reproduce universal characteristics of biological neurons, such as a firing threshold, as well as unique characteristics distinct to certain classes, such as parabolic bursting. Furthermore, we suggested that a nanowire transmission line with a propagation

speed of 2% c may be used as an axon delay line, potentially allowing spatio-temporal information to be accessed. These collective behaviors may enable the nanowire neuron to be used in a rich variety of operations.

In addition to harnessing the nonlinearity of the nanowire's switching dynamics, we relied on the nonlinearity of the nanowire's kinetic inductance in order to develop a variable inductive synapse. This scheme proved to be capable of fan-out, providing an advantage over other superconducting platforms. An energy analysis of the nanowire neuron in comparison to other spiking networks illustrated that it has competitive performance in the dynamic firing state and a figure of merit four orders of magnitude better than certain alternative technologies, while the static state benefits from the lack of power dissipation by superconducting elements.

Looking forward, networks of superconducting nanowires could be the basis for powerful new computer hardware. The ultimate goal would be a large-scale neuromorphic processor which could be trained as a spiking neural network to perform tasks like pattern recognition or used to simulate the spiking dynamics of a large, biologically-realistic network. The combination of speed, lower power dissipation, and biological realism with only a few components suggests that nanowires could outperform or complement other existing and developing neuromorphic hardware technologies. Finally, the possibility of integrating the nanowire neuron with the multilevel memory cell introduced in Chapter 2 opens the door to uniting memory and processing elements together on the same platform and even within the same lithography step. While future work is needed to realize this union, it presents an exciting opportunity to drastically reduce the memory access time bottleneck, and illustrates the creative ways in which nanowire dynamics can be exploited to achieve new functionality.

Chapter 4

Nanowire neuron: demonstrations

The previous chapter introduced the design of a nanowire-based artificial neuron and simulated its basic functionality. This chapter moves towards making the nanowire neuron a reality by describing the fabrication and testing of a soma made from coupled nanowire relaxation oscillators, which serves as the neuron's spiking element. To demonstrate the neuron's potential as a useful technology, we also simulate two applications using a network of nanowire neurons: 1) pattern recognition of simple nine-pixel images; and 2) winner-takes-all (WTA) competition based on stochastic firing.

The text that follows constitutes a preliminary write-up of work that will be submitted to a journal for publication. I would like to acknowledge Prof. Ken Segall for his guidance throughout the development of this work, Dr. Mike Schneider for his advice on the nine-pixel image recognition, Prof. Nancy Lynch for her guidance with the winner-takes-all network, and Matteo Castellani for his assistance with inductive coupling simulations.

4.1 Soma experiments

The soma is comprised of two identical nanowire oscillators connected together in a superconducting loop. We first describe the fabrication and basic characterization of the nanowire soma, and then experimentally demonstrate its bio-realistic capabilities.

4.1.1 Basic characterization

Fabrication

Before fabricating the soma, we considered several critical aspects of its design. First, each nanowire oscillator must have a sufficiently large series inductance such that the L/R time constant is on the order of nanoseconds, allowing for relaxation oscillations. As a result, long meandered nanowire inductors (∼200 squares) were added between the 60-nm-wide switching elements and the shunt resistor in order to get a series inductance on the order of nanohenries for a typical NbN film with a sheet inductance of 20–50 pH/sq. Additionally, the operation of the soma relies on the bias current splitting evenly between the two branches of the nanowire loop so that both oscillators are biased identically. To ensure that the oscillator biases were equal, COMSOL simulations of both pathways were used to check that the number of squares on either side of the bias port were the same (∼634 squares).

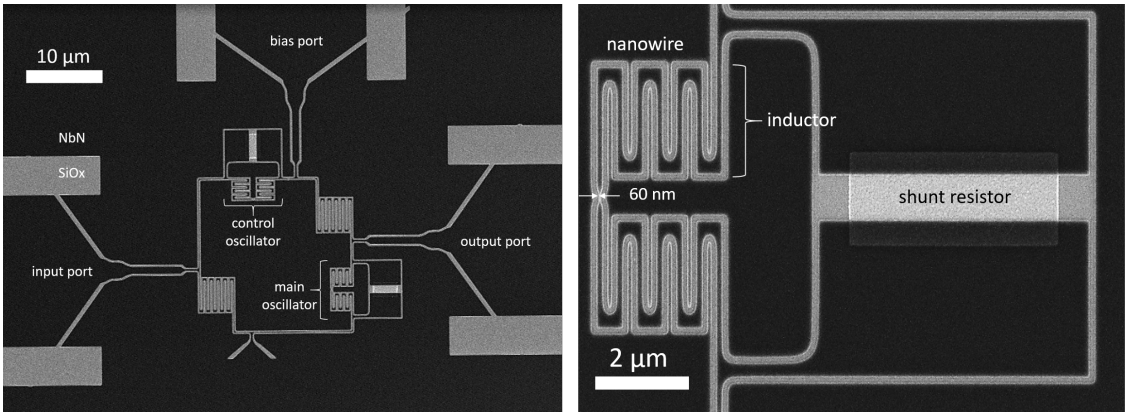

Figure 4-1: Scanning electron micrographs of a fabricated soma. (Left) The complete soma with labeled signal ports. (Right) Enlarged view of a single relaxation oscillator, indicated the 60-nm-wide switching element, the inductor, and the shunt resistor.

Once the design was finalized, the soma was fabricated using a similar process as was used to make the memory cell in Chapter 2. However, instead of a bilayer liftoff for the metal resistors, we used a single layer liftoff with ZEP520. Prior to exposure, the resist was spun at 5 krpm for 60 s and baked for 2 min at 180°C. Following exposure, the pattern was developed in o-xylene at 0°C for 90 s and IPA at room temperature for 30 s. A layer of 10 nm Ti and 25 nm Au was evaporated and lifted

off in NMP for 1 hr at 60°C. Afterwards, an NbN film was deposited ($R_s = 150\ \Omega$/sq), and the nanowire structures were patterned with gL2000 using the same process as described in Chapter 2. A scanning electron micrograph of a fabricated soma and an enlarged view of one of the relaxation oscillators is shown in Figure 4-1.

Single oscillator

To understand the spiking characteristics of the soma, it was first necessary to measure the dynamics of an isolated oscillator. Figure 4-2 shows the oscillation frequency as a function of bias current for an individual oscillator identical to the main and control oscillators in the nanowire soma. To measure the oscillation frequency, a bias current was applied using a battery source in series with a 100 kΩ resistor, and the output voltage was sent through a 50 dB, 1 GHz-bandwidth amplifier (MITEQ AM-1309) and read-out on an oscilloscope. A Fast Fourier Transform (FFT) of the output signal was used to identify the frequency peak.

The frequency-versus-bias curve of Fig. 4-2 was fit to a simplified model for nanowire relaxation oscillations [28], where the frequency is dominated by the slower time constant of the signal's falling edge:

$$\frac{1}{f} \approx -\left(\frac{L}{R}\right) ln\left(\frac{I_{bias} - I_{sw}}{I_{bias}}\right) \tag{4.1}$$

The red curve in Fig. 4-2 shows a fit to this expression when $I_{sw} = 45\ \mu$A and $L/R = 4.92$ ns. Using the calculated number of squares in the inductor and the approximate film inductance of 30 pH/sq, we can estimate that $L \sim 6$ nH and $R_s \sim 1.22\ \Omega$.

Overdriving the bias port

The two-oscillator nanowire soma can also be driven as a single oscillator if it is overbiased from the bias port without applying an input pulse. In this operation, the two oscillators fire simultaneously, since there is no input signal to induce a phase shift between them. Figure 4-3 shows the oscillation frequency of the overbiased soma without the application of an input pulse. Whereas the single oscillator switched at $\sim 45\ \mu$A, the soma starts oscillating at $\sim 82\ \mu$A since the bias current

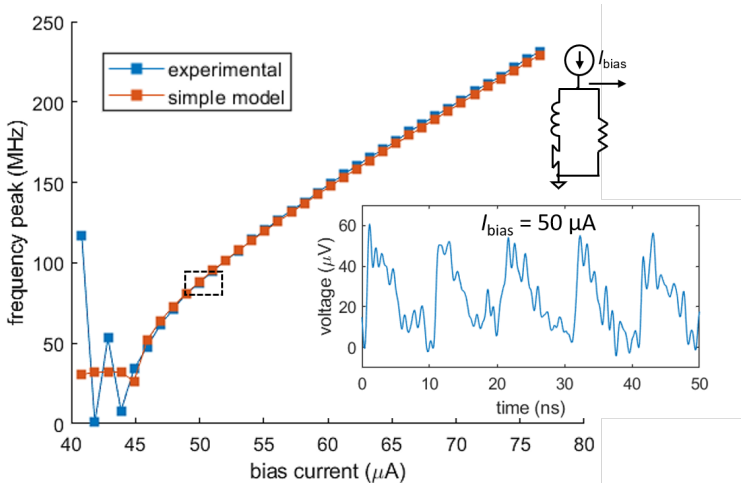

Figure 4-2: Frequency of a single oscillator as a function of bias current. Blue squares indicate experimentally measured points, while the red squares were derived from the simplified expression for oscillation frequency. The black box represents the 50 μA current bias, whose time domain characteristics are shown in the inset.

is split between the two branches of the loop. This frequency response can be well-explained by the soma model in LTspice, using the updated L/R time constants from the single oscillator in Fig. 4-2 and adjusting the switching current to match the experimental results. The red data points in Fig.4-3 show the simulated response of the nanowire soma when $I_c = 38.5$ μA, $L_{nanowire} = 6$ nH, and $R_s = 1.22$ Ω. By comparing both the frequency response as well as the time-domain characteristics of the experimental and simulated results, we can conclude that the LTspice soma model sufficiently reproduces real dynamics of the physical device.

Response to input pulse

With the updated LTspice model, we can experimentally test the soma's real operation and use simulations to better understand its behavior. Figure 4-4 shows the output of the soma in response to an input pulse when the bias current is 76.2 μA, below the point at which the soma acts as a single oscillator. For this measurement, the bias current was supplied through a DC battery source in series with a 10 kΩ resistor and a bias-tee with the RF port shorted to ground. The input pulse was generated using an Agilent waveform generator in series with a 100 kΩ series resistor, which was sent to both the device and the oscilloscope using a pulse splitter. The voltage output was sent through the RF port of a bias-tee and the MITEQ amplifier

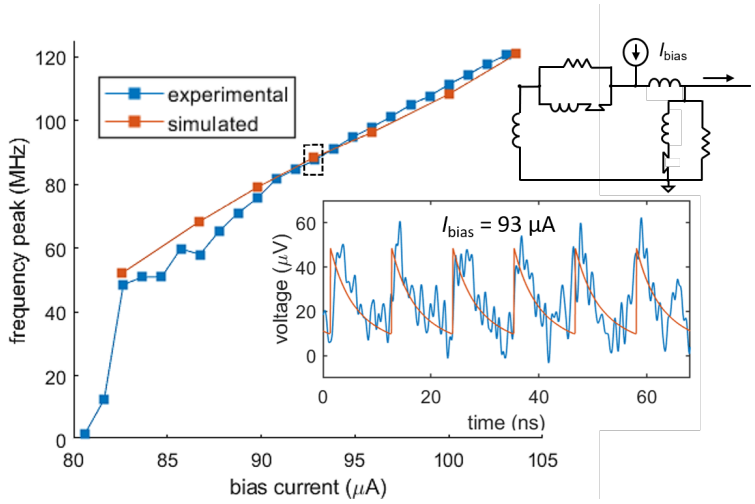

Figure 4-3: Frequency response of the soma when it is driven as a single relaxation oscillator (overdriven from the bias port with no input signal). Blue squares indicate experimentally measured points, while the red squares were simulated in LTspice using the L/R time constants derived from the single oscillator. The black box represents the 93 μA bias point, whose time domain characteristics for both the measurement and simulation are shown in the inset.

before being read-out by the oscilloscope. As demonstrated in the figure, the soma only spikes in response to the input pulse, in agreement with the expected operation. Comparison with the simulated results suggests that the timescale of the spikes also follows the expected response.

To ensure that the spikes are coming from the phase-shifted firing of both oscillators, we also examined the voltage signals of the input port and the bias port, and compared them to the simulated responses. Figure 4-5 (a) shows the output port (blue trace) and bias port (red trace) voltages in response to an input pulse, while the output and input port signals are shown in (b). The simulated responses for both cases are plotted in (c) and (d). By comparing Fig. 4-5(a) and (c), we observe that the signal from the bias port has one positive spike for each oscillator, while the traces in Fig. 4-5(b) and (d) show that the input port signal has one positive edge followed by one negative edge. The large spikes observed on the rising and falling edge of the input pulse in Fig. 4-5 (b) are likely absent from the simulation in (d), since the simulation does not account for the effects of the measurement setup, such as the amplifier and bias-tee used for readout. Despite the experimental results being

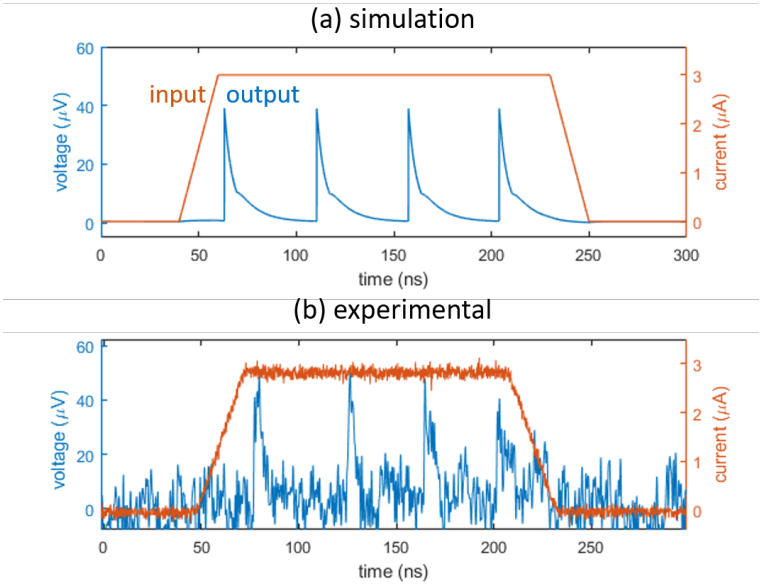

Figure 4-4: Spiking response of the soma to an input pulse. (a) Response from simulation. (b) Experimental measurement. For both cases, $I_{bias} = 76.2\ \mu$A.

noisier than the simulations, the overall agreement in the shapes of the input and bias signals indicates that the output port spikes are generated by the action of both oscillators in the loop.

Once we validated the origin of the spiking behavior, we characterized its reproducibility by gathering statistics from 100 sequential output voltage waveforms. Figure 4-6 displays a single waveform when the soma is subjected to a 2.5 μA, 150-ns-wide input pulse and a bias of 76.3 μA. The red dashed line indicates the threshold voltage used to determine the location of output peaks. Figure 4-6(b) shows histograms of the locations of these peaks for the 100 captured waveforms, while Figure 4-6(c) shows a histogram of the time between adjacent peaks, with a mean of 50.4 ns and a standard deviation of 6.46 ns. This spread is comparable to what has been observed in the interspike intervals of human motoneurons, where the standard deviation was roughly 5–10% of the mean interspike interval (the spread increased with increasing spike periods) [63].

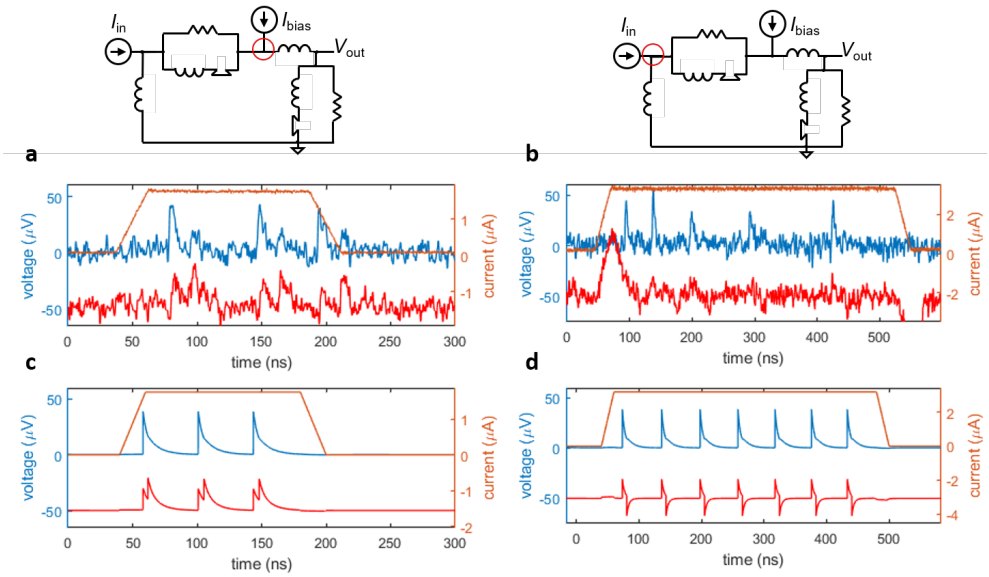

Figure 4-5: Output voltages of the input and bias ports. For all panels, the voltage from the output port is shown in blue, while the input/bias port output voltages are shown in red. Traces have been shifted on the y-axis for clarity. The orange trace shows the input current pulse. (a) Response from the bias port when I_{bias} = 75.7 μA, and I_{in} = 1.75 μA (pulse width = 150 ns). The bias port is indicated by the red circle in the circuit schematic above. (b) Response from the input port when I_{bias} = 75.9 μA, and I_{in} = 3.25 μA (pulse width = 480 ns). The input port is indicated by the red circle in the circuit schematic above. (c) Simulated response of the bias port. (d) Simulated response from the input port.

4.1.2 Neuron characteristics

As discussed in the previous chapter, biological neurons are capable of a wide variety of behaviors unique to different populations. Here, we focus on demonstrating certain characteristics that are universal to all neuronal classes.

Firing threshold

The firing threshold is defined as the minimum input signal required to initiate spiking for a given resting membrane potential. In the nanowire soma, the resting potential is dictated by the bias current—a larger bias current raises the resting potential and decreases the firing threshold.

Figure 4-7 shows the threshold response of the fabricated soma, measured as the mean voltage output of 500 sequential traces for a given input current. This measurement translates into firing probability, since the mean voltage of the 500

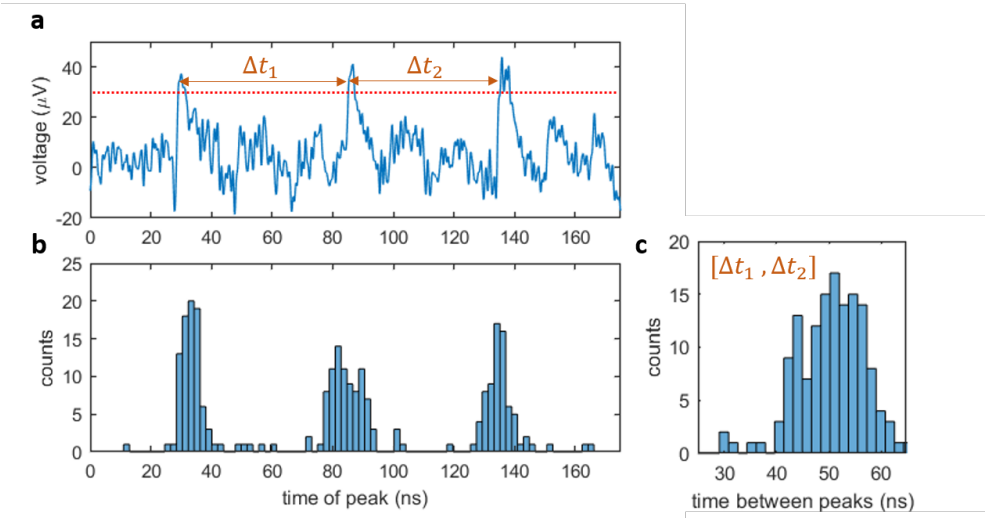

Figure 4-6: Reproducibility of spiking behavior in the soma. (a) Single waveform of the soma's output voltage in response to a 2.5 μA input current pulse (width = 150 ns, edge time = 20 ns) when the bias is 76.3 μA. The red dashed line indicates the threshold used to identify peak locations, while Δt_1 and Δt_2 denote the times between sequential peaks. (b) Histogram of the time stamps of voltage spikes for 100 captured waveforms. (c) Histogram of the time between sequential pulses for the 100 captured waveforms.

measurements will be higher if the soma spikes more often.

Unlike the simulated threshold behavior of Fig. 3-4, the experimentally measured data looks more like an "S" curve or sigmoid rather than a step response due to the probabilistic nature of nanowire switching in real measurements. Whereas the simulated nanowire always switches at the defined critical current I_c, real superconducting nanowires are susceptible to premature switching due to thermal and quantum fluctuations [64], leading to a switching probability that increases with bias current.

While stochasticity is usually avoided in electrical systems, biological neurons have firing probabilities that have inspired many neuromorphic applications that take advantage of probabilistic switching. Many learning algorithms, for instance, encode a sigmoidal probability into their spiking models in order to produce a gradual change in behavior for a small change in input. Since nanowires intrinsically possess probabilistic switching, they are a logical hardware platform for these types of applications. An example of harnessing stochastic firing for solving a real-life problem is discussed later in this chapter.

The different curves in Fig. 4-7 for various bias currents show how the threshold is altered by the "resting potential." These trends can be used to identify the optimal operating conditions for the soma, with a large difference in firing probability between the low and high input currents. Comparing the mean voltage outputs of the different curves at the same input current shows that the firing probability increases with increasing bias current, as expected. For low bias currents, we observed that the firing probability decreases at some point as the input current increases. One possible explanation for this phenomenon is if the inductances of the two branches are slightly unequal due to material inhomogeneities or other defects; in this scenario, the LI_c product of one oscillator would be greater than the other, causing flux to build up in the loop as the soma continues to fire. This accumulation would eventually prevent one of the oscillators from switching, reducing the firing probability.

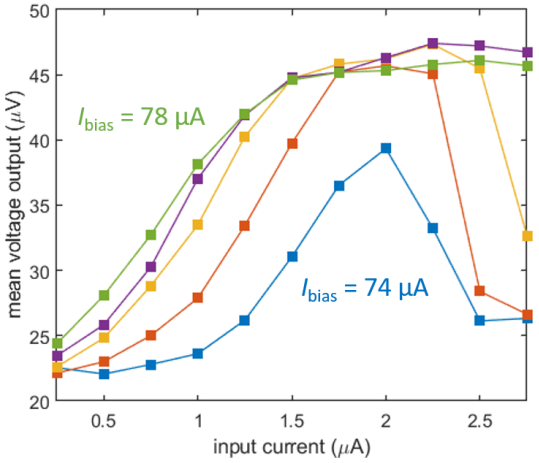

Figure 4-7: Firing probability as a function of input current, measured as the mean voltage output of 500 sequential traces. The different curves represent different bias currents, swept from 74 μA to 78 μA in 1 μA increments. Input pulses had a 150 ns pulse width and 20 ns edge times.

Refractory period

In addition to a threshold response, biological nanowires also display a refractory period, or a required "resting period" between two inputs so that both elicit their own output spikes. Figure 4-8 shows the effects of the refractory period on the fabricated soma in comparison to the simulated response. In both cases, the soma was biased at

75.9 μA while two identical input pulses were applied. The time between the input pulses was gradually decreased, and the output spikes from the soma were measured. In the first four panels, the time between input pulses is sufficient for both inputs to generate separate output spikes. In the fifth panel, however, the soma only spikes once. By looking at the simulation, it is apparent that the second input pulse ends right as the first spike relaxes, indicating that the main oscillator is not sufficiently biased during the second input pulse to fire again. As a result, it is clear that the soma has a finite refractory period.

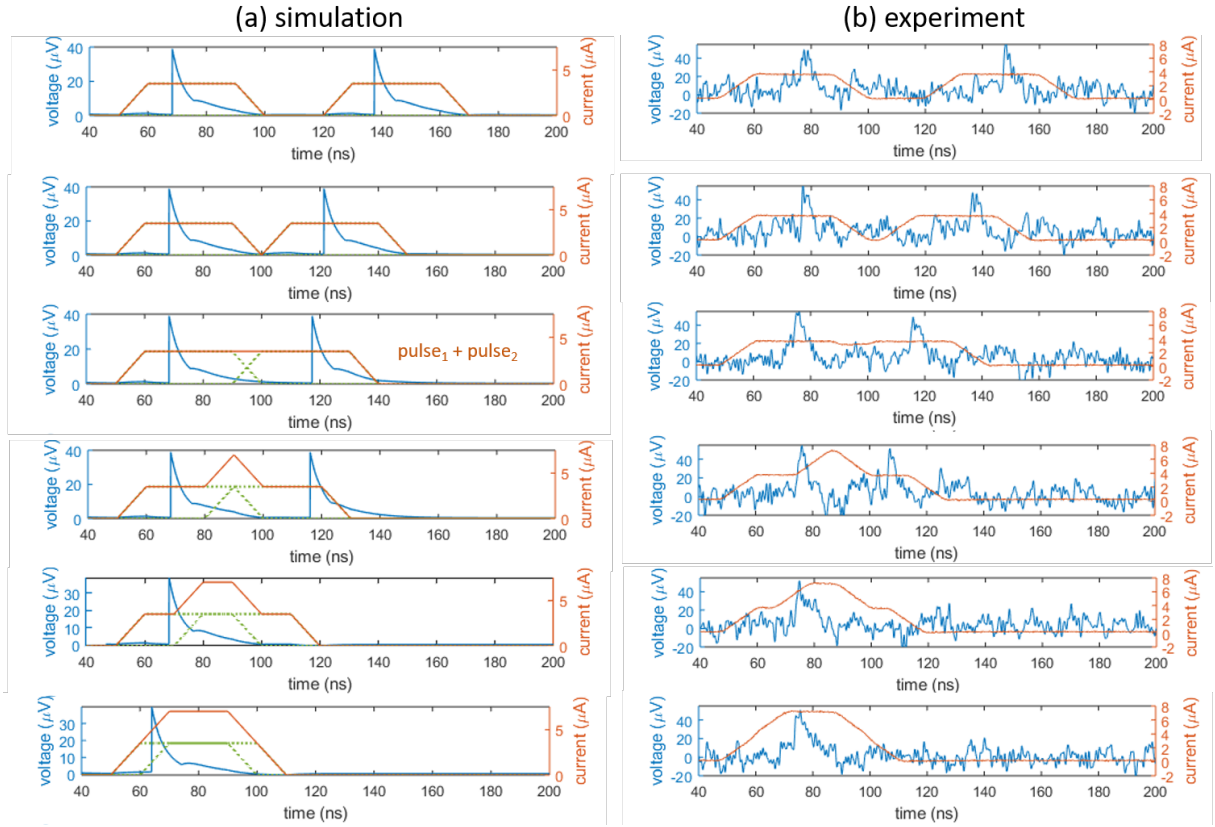

Figure 4-8: Evidence of refractory period in the nanowire soma. (a) Simulation. (b) Experimental measurements (single traces). The fifth panel shows the scenario when the time between input pulses is less than the refractory period, causing the soma to only spike once.

The refractory period measurement above can be repeated to get a sense of reproducibility. Figure 4-9 shows histograms of 200 measurements of the time of each output spike as the separation between input pulses is gradually reduced. The col-

lapse of two distinct histograms into one as the input pulses move closer together verifies the refractory period previously observed in the individual traces.

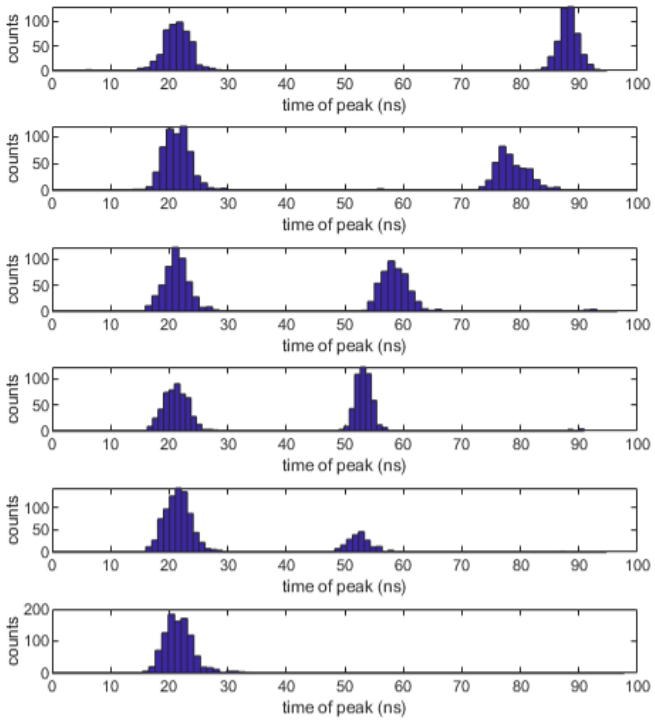

Figure 4-9: Reproducibility of the refractory period measurement. Histograms represent 200 measurements of the time at which a spike occurs. The collapse from two distinct histograms into one confirms the effect of the refractory period.

4.2 Non-stochastic application: image recognition

We can use the LTspice model of the nanowire neuron updated with the experimentally measured circuit parameters to build a simple neural network that can perform basic tasks. One of the most common applications of neural networks is image or pattern recognition.

4.2.1 Methods and results

Figure 4-10 shows a set of 3×3 pixel images originally developed for physical memristor circuits [65] and recently simulated in JJ neural networks [66]. The set consists of

three "ideal" images of the letters "z", "v", and "n", as well as images with single-pixel errors (nine error images per letter), leading to a total of 30 images.

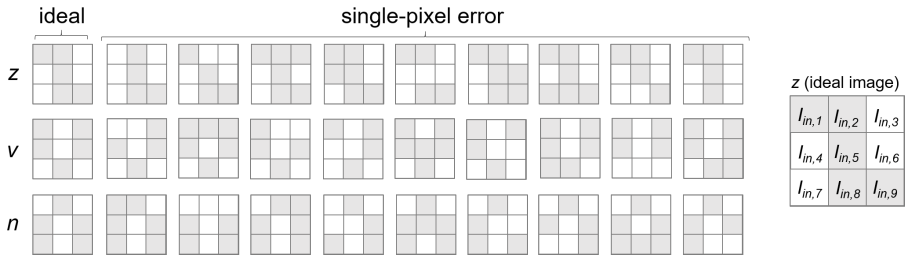

Figure 4-10: Set of the 30 test images. Each letter is represented by an ideal nine-pixel image and nine single-pixel-error images. The enlarged image on the right shows how each pixel is mapped to the input current of the nine pixel neurons. Figure is based on the abstraction in Ref. [65].

To identify these images, we can use a 9×3 neural network consisting of nine input pixel neurons and three output letter neurons, as shown in Figure 4-11. In our case, the pixel colors within each image determine the input current to each pixel, with grey pixels corresponding to an input current of 4.6 μA and white pixels corresponding to an input current of 0 μA. As discussed in Chapter 3, the weight of each connection maps to the magnitude of the synapse inductor, with higher weights interpreted as lower inductances, leading to more synaptic current.

Following the method described in Ref. [66], we used a basic neural network script [67] written in Python to build the image set and solve for the weights of a 9×3 neural network. Due to the limited size of the data set, we used all 30 images for both training and testing, leading to 100% identification. The code for generating the image set and solving for the weights may be found in Appendix B.1. Table 4.1 lists the calculated weight for each synapse, which was later linearly scaled to a corresponding inductance value in the circuit. For every synapse, the series resistor was 4 Ω and the parallel resistors were both 25 Ω.

To test all of the 30 images within one transient simulation, we used piecewise-linear current sources for the pixel inputs that changed to a new image every ∼50 ns. The Python code for creating the .txt piecewise-linear file for every pixel can be found in Appendix B.2.

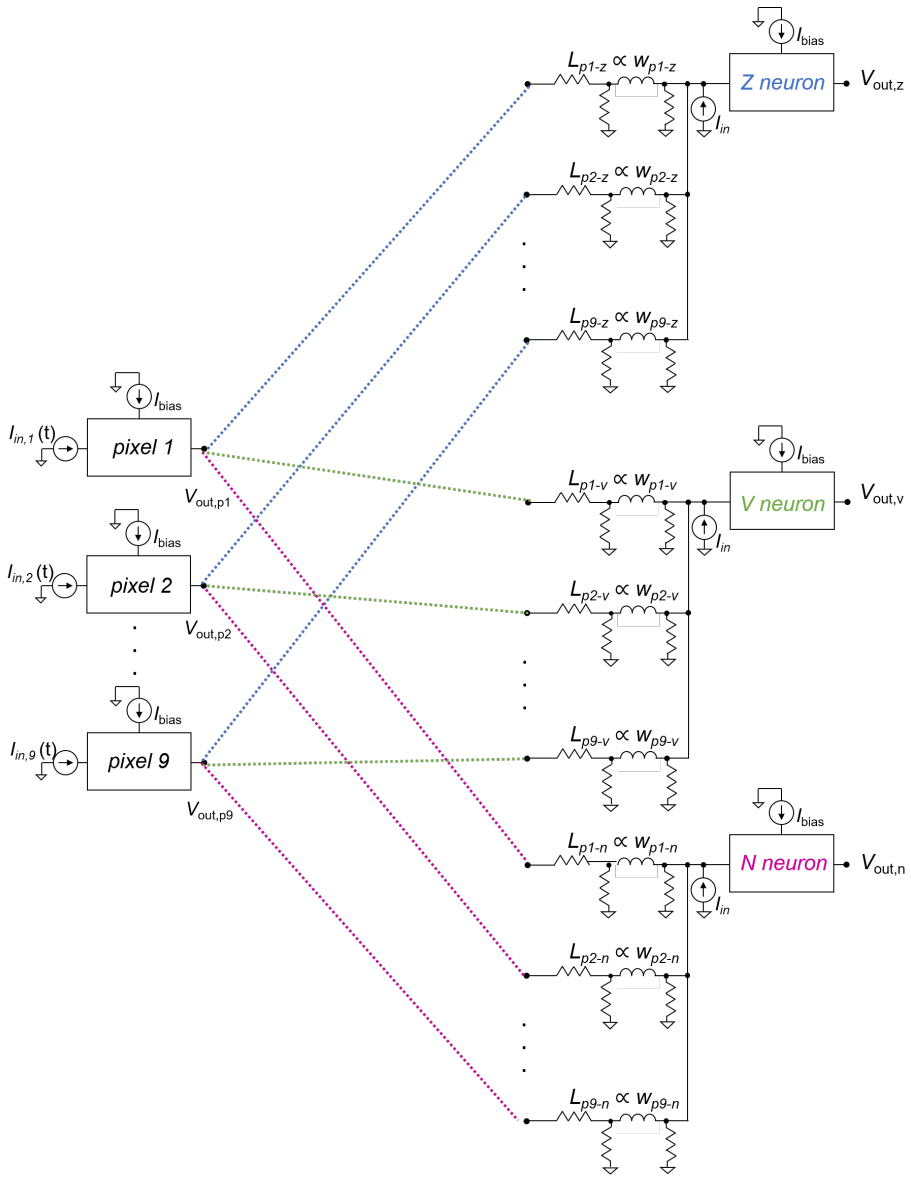

Figure 4-11: Simplified schematic of the image recognition circuit. The voltage of each pixel neuron is sent to a synapse connecting to each letter neuron using a behavioral voltage source. The weight of each synaptic connection determines the magnitude of the synapse inductor. Grey pixels correspond to an input current of 4.6 μA, while white pixels correspond to an input current of 0 μA.

Figure 4-12 shows the pixel input currents and the letter output voltages for all 30 images. The ideal cases for each letter are indicated with the red dashed lines, followed by the nine single-pixel error images. As shown by the network's output, each letter neuron fires 10 times, or once for each image, indicating correct classification of both the ideal and single-error images.

Table 4.1: Synaptic weights for the 9-pixel circuit

	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9
z	1.89	1.04	-0.97	-1.85	0.37	-2.33	-1.13	0.80	1.94
v	0.31	-1.98	1.98	0.50	-1.14	0.93	-0.80	1.46	-1.14
n	-1.53	0.99	-1.35	1.66	-0.81	0.14	2.23	-1.71	0.40

Weights are converted into inductances for each synapse. A higher weight is translated into a lower inductance, leading to more synaptic current.

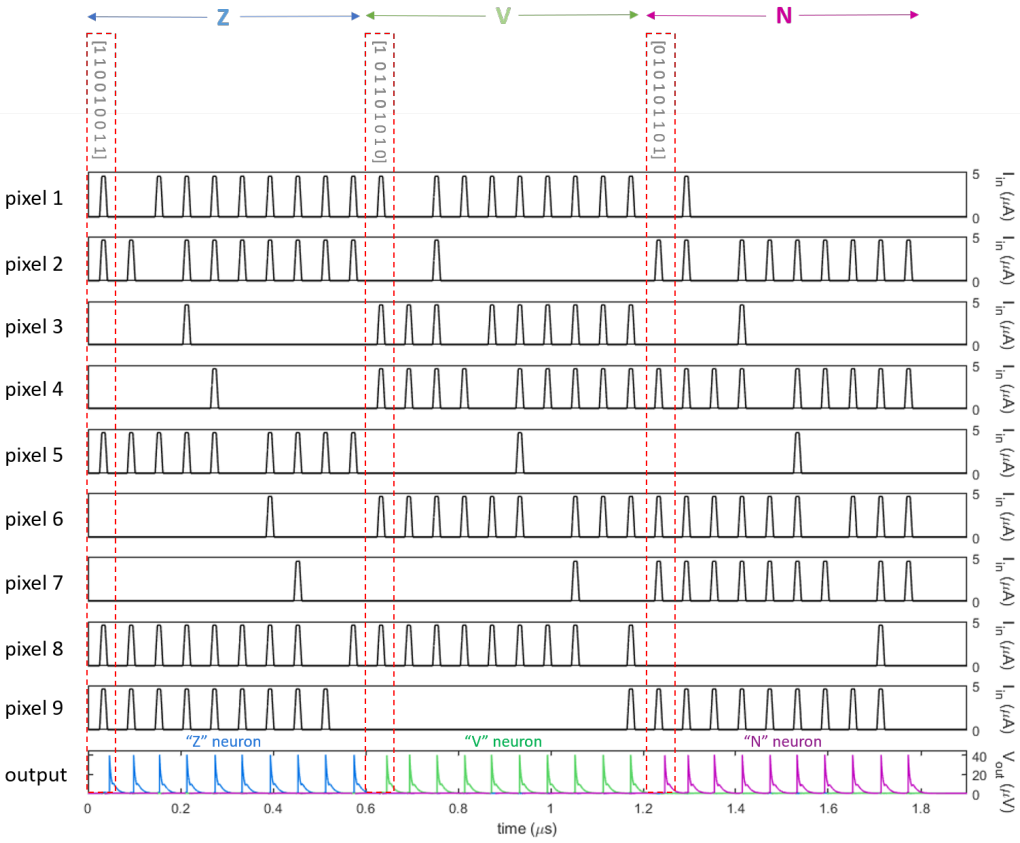

Figure 4-12: Simulation of each of the 30 test images. The top nine panels show the input current sent to each pixel neuron. The bottom panel shows the output voltage of each of the three letter neurons. Red dashed boxes indicate the "ideal" image for each letter, followed by their nine single-pixel error images. The "z" and "v" neurons were biased at 75.3 μA, and the "n" neuron was biased at 75.1 μA. As shown in the bottom panel, the output letter neurons correctly fire for each of their respective 10 images, and do not fire when the images of the other two letters are presented.

In this initial demonstration, we used behavioral voltage sources to pull the output voltage from each pixel to their synapses. However, it is also possible to transfer

the output of each synapse directly to the input of each pixel neuron using inductive coupling. Figure 4-13 shows a schematic of the 9×3 neural network that uses inductive coupling between layers, adding an inductor after the second parallel resistor of the synapse. Like the previous simulation, the weight of each synapse is encoded in the magnitude of the synapse inductor, as well as the magnitude of the coupling inductor.

Figure 4-14 shows the result of testing all 30 images on the 9×3 neural network connected through inductive coupling. Like the previous demonstration using behavioral voltage sources, the network correctly identifies all of the test images, showing that the implementation of weights as inductive strengths is suitable for these types of applications.

In the cases above, we used all 30 images for both training and testing because the data set is small, leading to 100% classification. However, for larger data sets, we would like to be able to solve for the synaptic weights using only a fraction of the total images as a training set, and then test the resulting network on all of the images. To show how our network would respond in this scenario, we repeated the training process using just 9 images that were randomly chosen from the set of 30, and then modified our circuit with these new synaptic weights. The results shown in Figure 4-15 indicate that the circuit correctly identifies 23 of the 30 images in the complete set, which is to be expected in comparison to the previous results since not all of the images were used during training. We found that the classification could be improved by non-randomly selecting 12 training images (one ideal image and three single-pixel error images, per letter), leading to correct identification of 27 of the 30 images. Given that our circuit design has not undergone any optimization, we believe that our classification results could be further improved by refining our parameters through processes like Bayesian optimization.

4.2.2 Discussion

Although the nanowire neuron does not yet have a scheme for unsupervised learning, the classification results above demonstrate that it can be used for inference in a hardware neural network designed to perform a specific task. Just like in the simple

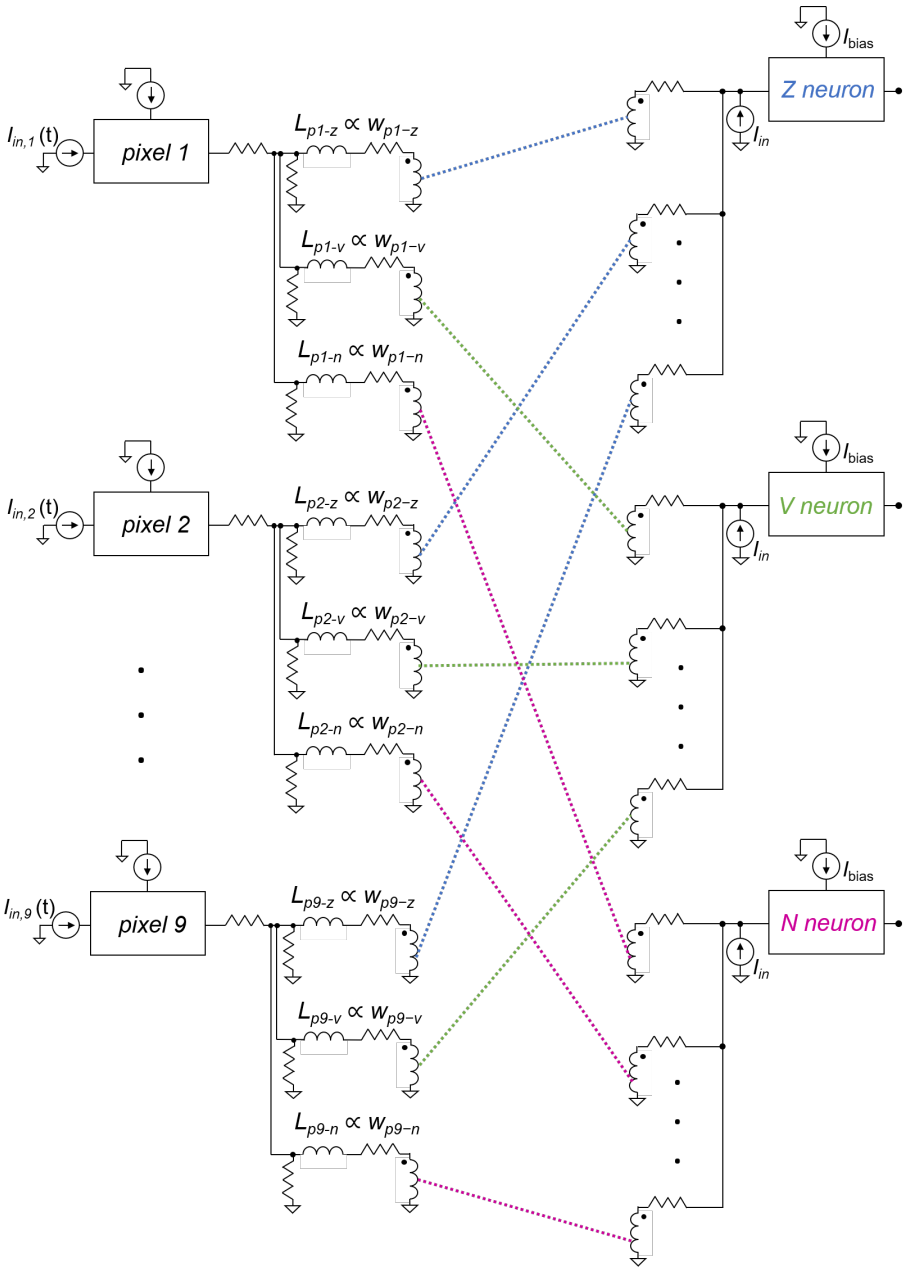

Figure 4-13: Simplified schematic of the image recognition circuit using inductive coupling. For each synapse connecting to a pixel, the parallel resistance is 25 Ω, while the left and right resistances are 5 Ω and 6 Ω. The synapse inductance and the coupling inductance are proportional to the weight, scaled to a baseline of 0.1 μH. The magnitude of the receiving coupling inductor attached to each letter neuron is twice the synaptic inductance, and is placed in series with a 0.5 Ω resistor.

3×3 image recognition problem, a network could be built for a targeted application and software algorithms could be used to calculate weights, which would determine the synapse inductors for the real system. A physical circuit could be then built with

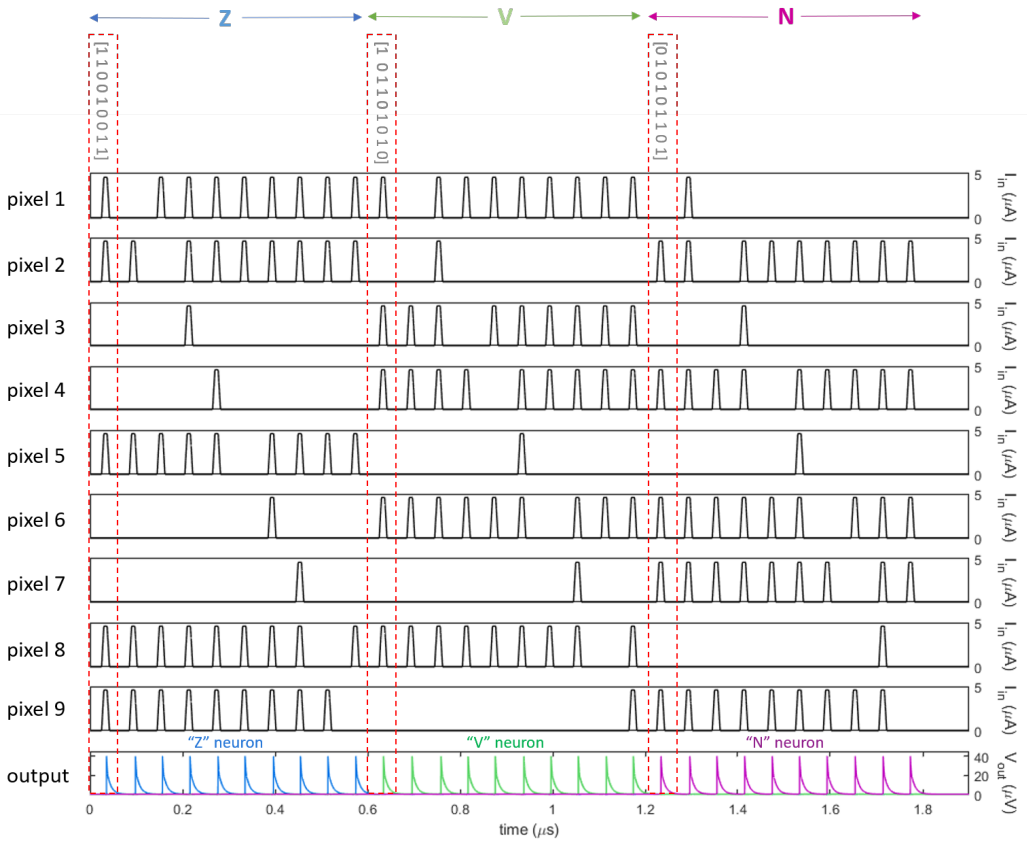

Figure 4-14: Simulation of each of the 30 test images in the inductive coupling circuit. The first nine panels show the input current sent to each pixel neuron. The last panel shows the output voltage of each of the three letter neurons. Red dashed boxes indicate the "ideal" image for each letter, followed by their nine single-pixel error images. The letter neurons were biased at 76.9 μA and had input currents of 5.46 μA.

these calculated parameters and used as an energy-efficient inference platform.

These simulations of a simple neural network using experimentally measured device characteristics also set the foundation for larger, more complex demonstrations with nanowire neurons. For example, greyscale images could be processed by mapping the pixel color to spiking frequency, which can be tuned with bias current, while finer image resolution may be obtained by increasing the number of pixel neurons.

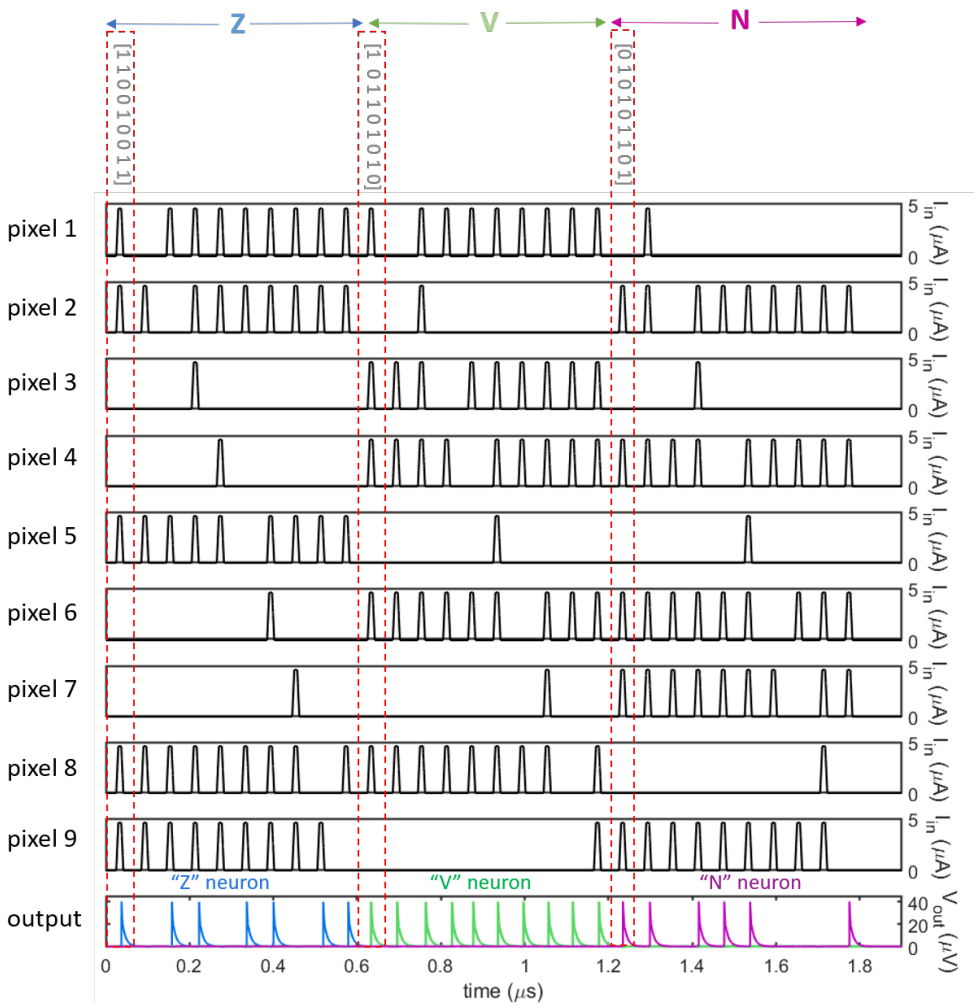

Figure 4-15: Simulation of each of the 30 test images in the inductive coupling circuit, using a training set of 9 randomly chosen images. 23 of the 30 images were correctly identified. The letter neurons were biased at 76.9 μA and had input currents of 5.45 μA.

4.3 Stochastic application: winner-takes-all (WTA)

The image recognition example above used a soma model that fires deterministically, or exactly when the current through the nanowire exceeds its I_c. However, the threshold measurements in Fig. 4-7 revealed the probabilistic nature of spiking in real superconducting nanowires due to the effects of noise and fluctuations. Here we discuss an application that relies on probabilistic firing.

4.3.1 Background

As mentioned previously, biological neurons also have firing probabilities, which has inspired theories of how stochastic behavior may play a critical role in the brain's operation. One key example of harnessing stochasticity is the winner-takes-all (WTA) theory of how the brain develops selectivity for a set of inputs [68]. With hundreds of thousands of neural connections, the brain must somehow make meaningful distinctions between many input signals in order to make decisions. In other words, if every neuron fired in response to a large number of input signals, the overall response would be meaningless.

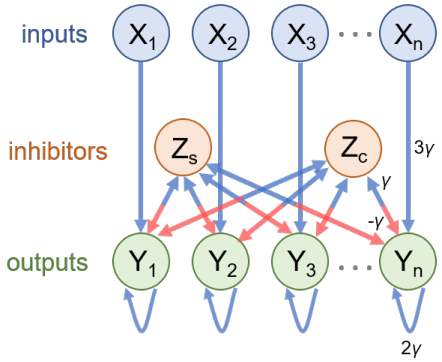

Figure 4-16: Diagram of the two-inhibitor WTA network. A set of inputs X connect to a set of outputs Y. Two inhibitors Z_c and Z_s foster competition between the outputs. Blue arrows denote excitatory connections (positive weights), red arrows denote inhibitory connections (negative weights). Weights for all connections are labeled on the n^{th} neuron and are scaled relative to the unitless factor γ. Figure is based on the diagram in Ref. [69].

The WTA theory suggests that the brain develops selectivity through competition between excitatory neurons that share a set of inhibitory connections. As the excitatory neurons start firing in response to a set of input signals, they trigger the inhibitors, which in turn start to suppress their firing. The excitatory neurons "compete" until just one is left in the firing state.

The idea of developing selectivity through competition has inspired the use of WTA in artificial networks for filtering, image recognition, and decision-making [70]. It has also been suggested that WTA subcircuits are repeated throughout the brain, creating selective responses that together dictate a cumulative behavior.

Figure 4-16 shows a schematic of a two-inhibitor WTA network based on the theory presented in Ref. [69]. When a set of input neurons $X_{1:n}$ fires, they trigger a set of stochastic output neurons $Y_{1:n}$ that each fire with a finite probability. Their competition is facilitated by two inhibiting neurons Z_s, the stability inhibitor, and Z_c, the convergence inhibitor. Z_s is biased so that it fires when at least one output neuron is firing, whereas Z_c is biased so that it only fires if more than one output is firing. Z_c eventually forces all but one neuron to stop firing, while Z_s continues to fire in order to stabilize the network and suppress all but the dominant neuron. Since $Y_{1:n}$ are identical and stochastic, they have equal probability of winning if all of their inputs are active.

4.3.2 Methods and results

To take advantage of the nanowire's intrinsic stochasticity for simulating a WTA competition, we can introduce noise sources into the soma model that cause it to spike with a bias-dependent probability. We amended our model to include Gaussian white noise sources at the bias port and the input port, where we expect the majority of noise to be injected. Keeping the noise bandwidth constant at 1 GHz, we varied the noise amplitude until the firing probability was similar to what we observed experimentally.

Figure 4-17 compares the simulated firing probability to the experimentally measured probability, converted from the mean voltage output curve of Fig. 4-7. The experimental bias was 74.5 μA, while the simulated bias was 76.6 μA with a noise amplitude of 800 nA. The simulated probability was calculated by recording the number of times the soma fired out of 50 trials for each input current level. The relative agreement between the simulated and experimental results suggests that the noise source is suitable for implementing a realistic firing rate in our model.

With these adjustments to our model, we designed a simple two-inhibitor WTA network with three inputs and three outputs. The inputs and inhibitors were deterministic, while the outputs included the noise sources described above in order to fire stochastically. Each output neuron was biased at 76.6 μA, while the bias to Z_s

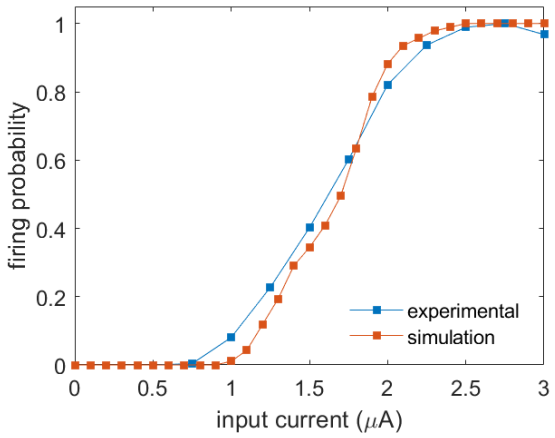

Figure 4-17: Simulated firing probability compared to experimental result. The experimental bias was 74.5 μA. The simulated bias was 76.6 μA.

was 76.8 μA and the bias to Z_c was 76.5 μA. The inductance of each synapse was calculated by scaling the unitless weights shown in Fig. 4-16 to a baseline inductance of 0.77 μH. As in the initial design of the image recognition circuit, output voltages were transmitted to synapses through behavioral voltage sources.

Figure 4-18 shows the time-domain voltage output for each neuron in an example when the first and third input neurons are active (overall input vector = [101]). Output neurons Y_1 and Y_3 both fire, triggering Z_s and Z_c. Eventually, Y_3 is turned off by the inhibition, and Z_c is no longer active. Y_1 continues to fire after it wins, in addition to Z_s which keeps the other two outputs suppressed.

Figure 4-19 shows a competition between all three output neurons (input vector = [111]). As expected, Z_c influences the competition by first eliminating Y_1 and then Y_3, turning off only once Y_2 wins. Z_s remains firing to stabilize the network.

Since the firing of all three output neurons is probabilistic, they should have equal chance of winning a competition if the three inputs are active and every other parameter is the same (e.g. the synapse between an inhibitor and a particular output neuron is not weighted differently from the others). To prove this is the case, we ran 49 competitions between all three inputs and recorded the winner for each trial. Figure 4-20 summarizes the outcome. The blue data represents wins that resulted from a competition in which all three outputs were firing at some point. The red data

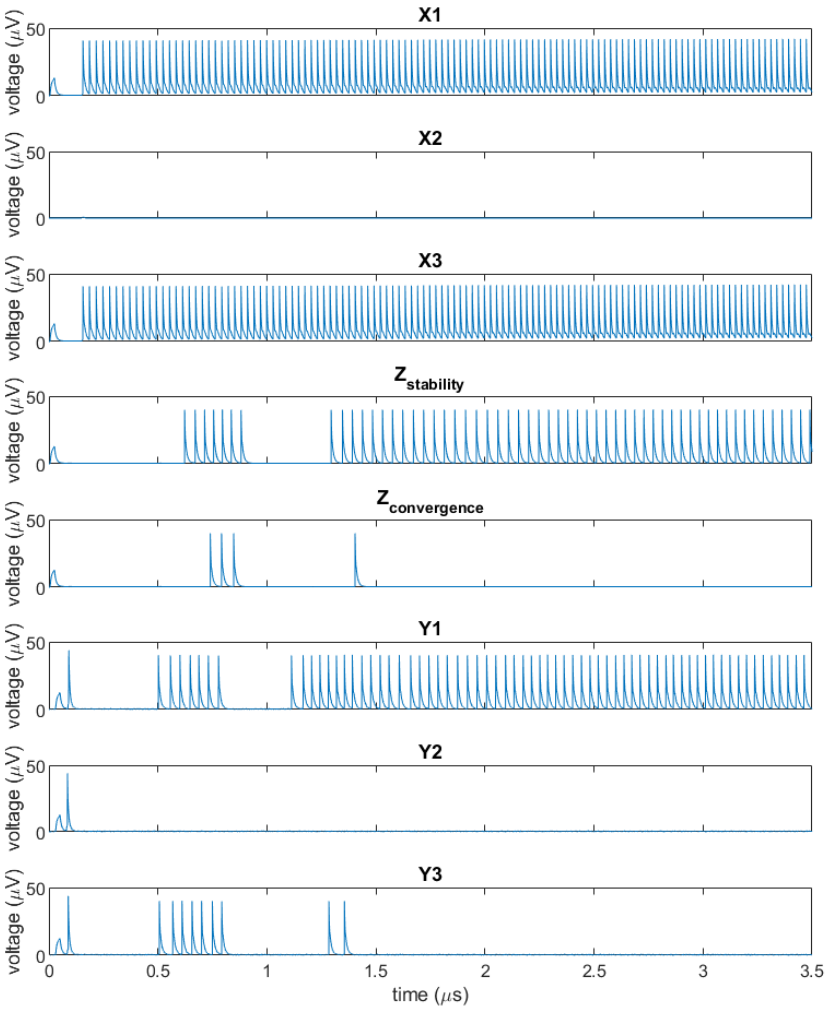

Figure 4-18: WTA competition between two output neurons, Y1 and Y3. The input to the second output neuron is turned off, so Y2 never fires. Y1 and Y3 both initially fire in response to their inputs, but Y3 is eventually turned off by the inhibitors and Y1 wins. The stability inhibitor $Z_{stability}$ continues to fire after the competition in order to keep Y3 suppressed.

represents wins in trials where only two outputs fired, indicating that the remaining output was suppressed early on. Six of the trials are not shown, since they either had no winner by the end of the simulation time, or a win occurred but not through the actions of the convergence inhibitor. Overall, Y_1 and Y_3 both had a total of 14 wins, while Y_2 had 15 wins. These results suggest there was an equal competition between

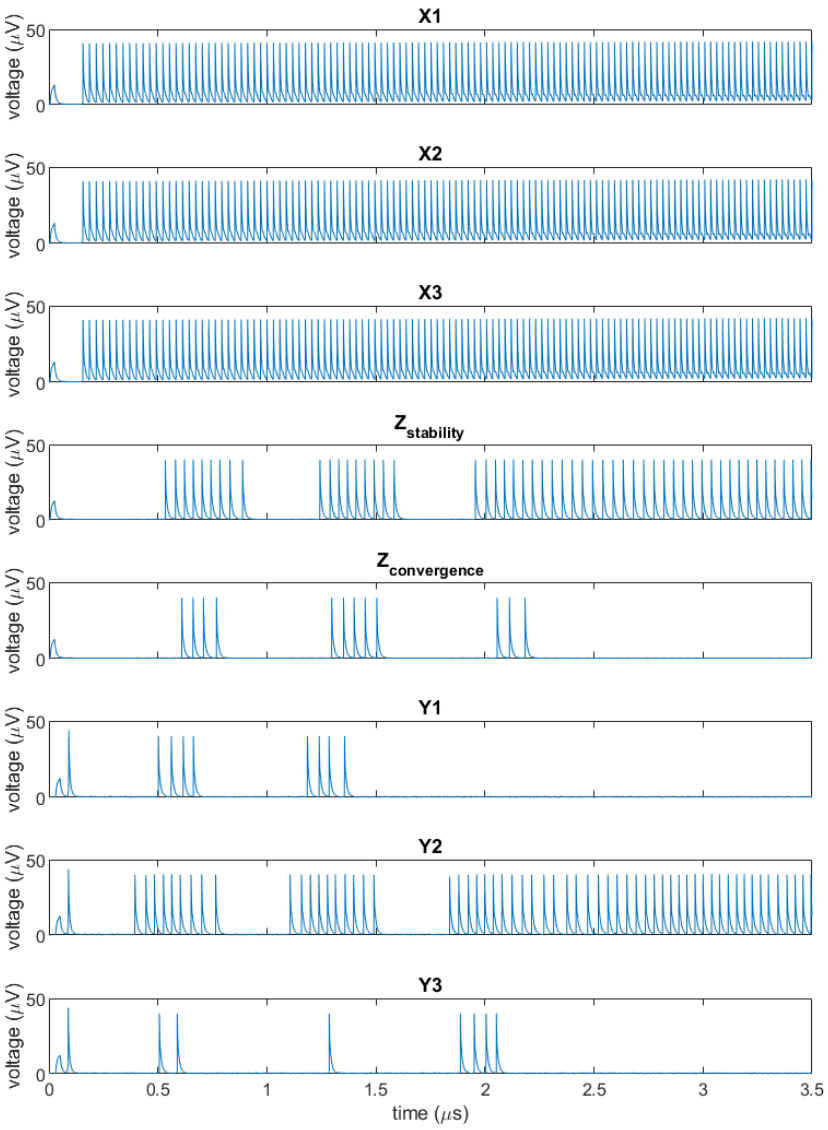

Figure 4-19: WTA competition between all three output neurons. Each of the three input neurons are active, causing the three outputs to fire. Y1 is the first output to be turned off, followed by Y3. Y2 wins, while the stability inhibitor continues to fire in order to keep the other two output neurons suppressed.

the output neurons.

Like in the image recognition example, it is also possible to use inductive coupling to connect synapse outputs to the inputs of neurons in the WTA circuit. However,

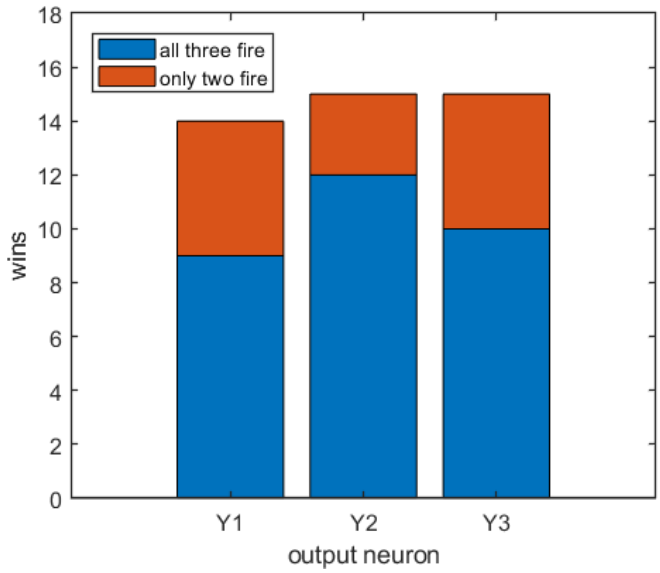

Figure 4-20: Result of 49 sequential simulations of WTA with three neurons. Blue bars indicate wins where all three neurons fired over the simulation. Red bars indicate wins where one neuron never fired, leading to a competition between the remaining two. 6 of the 49 simulations had errors, where either there was no winner, or the convergence inhibitor failed to fire.

we found that the circuit may not be as stable as the image recognition case due to the added noise sources. The amount of synaptic current being received by each output neuron is the same order of magnitude ($\sim 0.5\ \mu$A) as white noise sources. Consequently, when the neurons are biased high enough to be switched by the synaptic currents, they are also more likely to switch randomly due to noise [1].

Figure 4-21 shows an example of WTA competition between three output neurons whose synapses are connected by inductive coupling. In this particular competition, the first output neuron wins. As with the image recognition circuit, inhibitory connections were performed by reversing the polarity of the inductive couplers.

The inductive coupling simulations proved to be less stable than the simulations that used behavioral voltage sources to directly pull output voltages into synapses. Figure 4-22 displays the result of 100 repeated competitions between all three output

[1] It is worth noting that we refer to "bias" differently from typical neural network usage. In the case of superconducting nanowires, a higher bias current means a reduction in firing threshold, so the neuron will be more likely to fire. Neural networks often take the opposite meaning, where a higher bias means a higher threshold and a lower firing probability

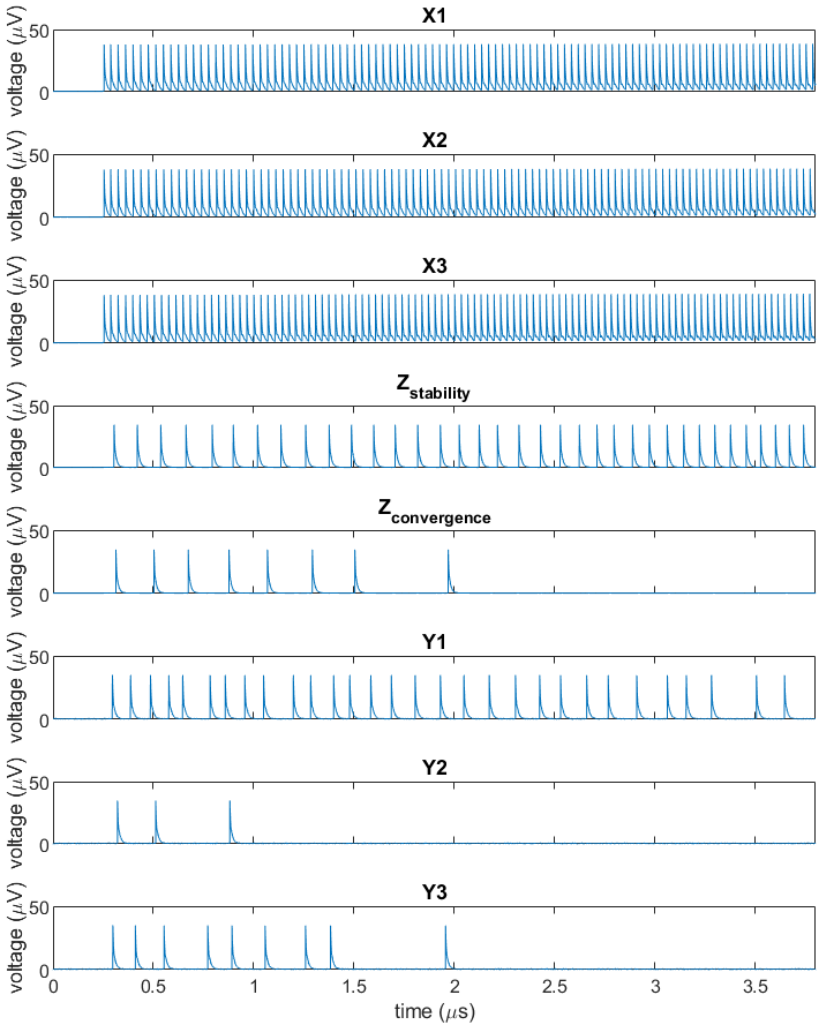

Figure 4-21: Example of a WTA competition between three output neurons when the synaptic connections are performed through inductive couplers. In this case, Y1 wins. For this simulation, the bias of the output neurons was 76.4 μA, the bias of Z_s was 76.98 μA, and the bias of Z_c was 76 μA. Synapse inductors were scaled relative to 0.77 μH, and the couplers had a 1:2 ratio.

neurons, where the noise to each output neuron was randomized at the beginning of every trial. Each neuron won between 20–27 total competitions (Fig. 4-22(a)), but 29 of the 100 competitions resulted in no winner, usually with all three output neurons being shut off (Fig. 4-22(b)). In some cases, this seemed to be due to continued firing of the stability inhibitor, despite the outputs being turned off, which could be caused

by instability resulting from biasing it so close to its critical current. This issue may be avoided by optimizing our circuit parameters to increase the synaptic currents so that the biases to the neurons can be lowered.

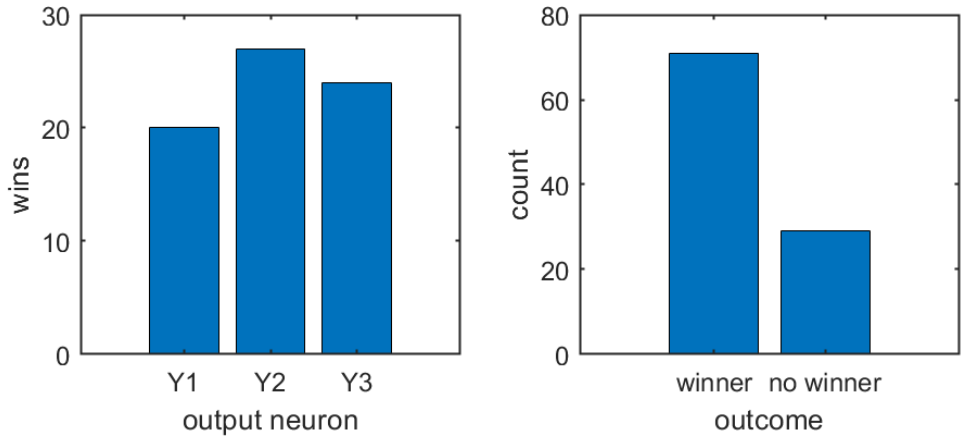

Figure 4-22: Result from 100 WTA competitions between all three output neurons using the inductive coupling scheme. (a) Number of wins per output neuron. Each neuron won between 20–27 of the 100 competitions. (b) 71 of the competitions resulted in a win, while the remaining 29 competitions either had no winner (i.e. either two or three neurons were still firing at the end of the simulation) or all three neurons were turned off.

4.3.3 Discussion

The WTA simulations showcase how the unique dynamics of superconducting nanowires can be harnessed for real applications and for mimicking behaviors observed in nature. While probabilistic switching is avoided in most digital circuits, neural-inspired architectures based on the stochasticity of biological neurons are prime opportunities to use nanowires to their full potential. Unlike other demonstrations of WTA competitions where winning results from a neuron having a higher intrinsic firing frequency or a weaker connection to the inhibitors, the scheme demonstrated in this thesis relies solely on probabilistic firing, demonstrating how probability can be a powerful tool in system dynamics. Our results show that experimentally measured device parameters can be incorporated into our model to build a network that uses stochastic firing to produce a meaningful output from a set of inputs, and that superconducting

nanowires have enough inherent stochasticity to support this type of competition. Looking forward, nanowire-based WTA circuits could be combined with memory elements that save the competition result as a means of establishing selectivity, while WTA subcircuits could be repeated throughout a large-scale network, like they are thought to in the brain.

Chapter 5

Fabrication studies

Superconducting nanowire devices like those presented in this thesis rely on the ability to pattern thin films with ∼10 nm precision. This high-resolution control is usually achieved with electron-beam lithography, involving the application of chemical resists and developers in order to create a desired pattern. It is therefore critical that the chemicals used in patterning a device are compatible with its underlying material in order to avoid device damage or degradation.

While the devices discussed thus far have been patterned using positive tone resists, negative tone resists are used in many processes. This chapter discusses two investigations of the compatibility of different negative tone processes with niobium nitride films. The first study shows that tetramethylammonium hydroxide (TMAH), a common developer for negative tone resists, reacts with niobium nitride thin films to form an etch barrier that limits device fidelity [71]. The second study looks at ma-N 2400 series photoresist as an alternative negative tone resist for patterning superconducting devices [72].

The text that follows includes reprints from works that were originally published in the *Journal of Vacuum Science & Technology B* in 2018 and 2019. I would like to acknowledge Marco Colangelo for conducting these studies with me, and Navid Abedzadeh for help with the scanning electron micrographs in Ref. [71].

5.1 Influence of TMAH on NbN

Precise control of superconducting materials at nanoscale dimensions demands both a high-resolution electron-beam lithography process and a superior electron-beam resist. Hydrogen silsesquioxane (HSQ) is a negative tone electron-beam resist often selected for these purposes due to its 5 nm resolution and minimal line edge roughness [73] [74]. After patterning HSQ, development is commonly done using solutions of the strong base TMAH; for instance, past processes have developed HSQ by submerging it in 25% TMAH for durations ranging from 1 to 4 min [75] [76] [77] [78]. TMAH has also been used to pretreat films before spinning resist in order to activate the film surface and promote HSQ adhesion [79] [80].

Despite the prevalence of TMAH in fabrication processes for superconducting devices, its potential negative effects on the superconducting film itself have not been fully investigated. As a result, it is unknown whether TMAH damages films in a way that somehow limits superconducting device performance. Prior literature showed that pretreating NbN films with 25 % TMAH for 4 min increased the sheet resistance R_s and decreased the critical temperature T_c, suggesting degradation and potential thinning of the superconducting material [79]. However, no further studies have been made to verify or explain this observation.

We were motivated to understand this phenomenon when we noticed that a thick NbN film (thickness d ~20 nm) patterned with HSQ took nearly twice as long to etch in CF_4 as an unpatterned control film, causing the HSQ mask to be almost entirely removed before the etch was complete. The effects of this over-etching are displayed in Figure 5-1. As this mask removal left the underlying NbN vulnerable to damage by reactive ion etching (RIE), it became clear that a deeper understanding of the lithography process was needed to guarantee successful pattern transfer and device fidelity, especially for thick films.

To address this problem, we investigated the HSQ development process with NbN films and demonstrated that the observed discrepancy in etch time is caused by TMAH reacting with NbN, modifying the film's surface chemistry and reducing the

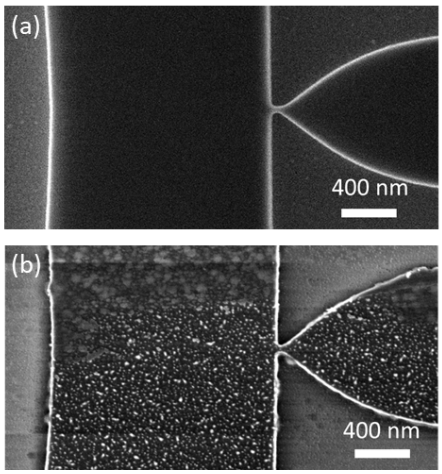

Figure 5-1: Removal of HSQ mask due to prolonged reactive ion etching. (a) SEM of a three-terminal superconducting device patterned with HSQ prior to etching. (b) SEM of the same device after the etching was complete (8 min). The rough surface suggests that the HSQ mask may have been mostly removed, exposing the NbN film of the device and leaving it vulnerable to damage.

thickess of the pure superconductor. As shown in Figure 5-2, TMAH has an observable impact on the film's appearance, increasing the surface roughness and creating clusters of material that form a barrier to reactive ion etching in CF_4. The steps taken to uncover this new understanding are presented here as two primary experiments. In the first experiment, we explored how exposure to TMAH hinders the reactive ion etching of NbN films, and demonstrated how the effect may be ameliorated. In the second experiment, we performed material analysis on both the superconducting film and the clusters formed at the surface to identify the etch contaminant species. Through these steps, it became apparent that the reaction that forms the surface contaminant deteriorates the critical parameters of the pure superconductor by consuming the film.

5.1.1 Experiment 1: impact on etch rate

To investigate the impact of TMAH on the reactive ion etching of NbN, we exposed NbN samples diced from the same wafer to the 25 % TMAH developer for various durations, and measured their sheet resistances after etching as an indicator of remaining film thickness.

Methods

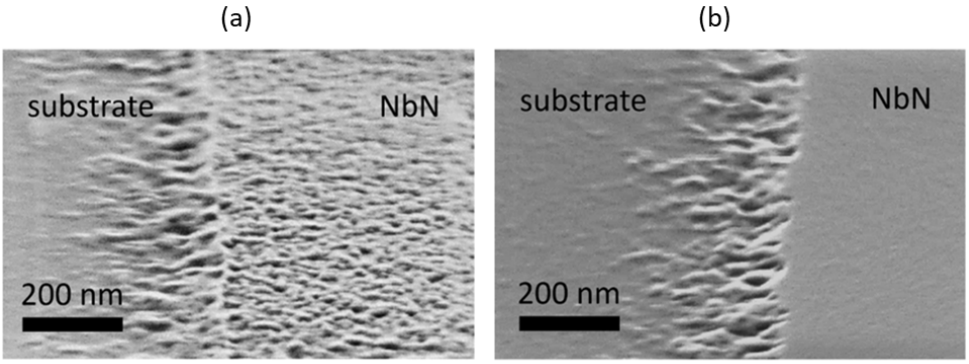

Figure 5-2: Modification of the NbN surface by TMAH. SEMs reveal the difference in surface features between (a) an NbN film that has been submerged in 25 % TMAH for 4 hours, and (b) an NbN film that has been untreated. The ridge represents the edge of a pattern made via photolithography to serve as a reference point for imaging. The treated sample was immersion rinsed in deionized water for 30 s following exposure to TMAH.

An NbN film was first bias sputtered at room temperature on a 4 inch Si wafer in an AJA sputtering system following the procedure described in Ref. [38]. The film had a sheet resistance of approximately 90 Ω/square, a critical temperature of 8.1 K, and a thickness of $\sim$20 nm. After dicing the wafer into 1 cm square samples, the samples were cleaned in acetone, methanol, and isopropyl alcohol (IPA). Once cleaned, we submerged each sample in 25% TMAH for durations ranging from 0 to 90 s. To stop exposure to TMAH, the samples were rinsed in deionized (DI) water for either the traditional 30 s used in our standard process, or for 3 min under a constantly running DI stream. Samples were then briefly rinsed in IPA, another step used in our conventional HSQ process to reduce surface tension. No HSQ was applied at any step of the process.

Following exposure to TMAH, the samples were reactively ion etched in a Plasmatherm RIE in CF_4 for 2.5 min at an RF power of 50 W and a chamber pressure of 10 mTorr. The sheet resistances of the samples were then measured using a four-point probe. The resistance of the samples were measured again after an additional minute of etching, corresponding to a total etch time of 3.5 min. To get a better sense of how the measured sheet resistances reflected the etching process, the resistances were converted into inferred film thicknesses using the conversion $d = \rho/R_s$, where d is

the remaining film thickness, R_s is the sheet resistance, and ρ is the resistivity. An estimate of ρ was obtained by using x-ray reflectometry to measure the thickness of a sister film with the same sheet resistance as the samples under investigation and calculating the resistivity using the expression above. From this approach, we obtained a value of $\rho = 2.5$ kΩ·nm, which is consistent with the results for a $\sim$5 nm film reported in prior work [81]. Given the agreement between values, we assumed a constant resistivity across all of the samples for the sake of simplicity, allowing us to compare inferred film thicknesses between samples.

Results

Figure 5-3 shows the inferred remaining NbN film thickness as a function of exposure time to the TMAH after each of the etch trials. In comparison to the undeveloped control sample, those that had been exposed to TMAH for 30–90 s and rinsed with the standard 30 s DI water procedure had noticeably greater film thicknesses after the 2.5 min etch; this difference in film thickness indicates that the processed films were less etched than the control film, as we had initially observed in the patterned chip that first motivated this study. The difference in thickness between the control and exposed samples slightly increased after an additional minute of etching, implying that a type of barrier to etching was present on the treated films.

Interestingly, subjecting the treated samples to a 3 min DI water rinse under a running stream significantly reduced the discrepancy in remaining film thickness of the treated films with respect to the control. Whereas treated sampled that were rinsed for 30 s differed in thickness from the control by about 7–8 nm, those that underwent the vigorous rinse were within 2 nm of the untreated sample film thickness. This result suggests that the material that formed the barrier to etching is water soluble. However, it should be noted that simply submerging treated samples in DI water for 3–5 min had no noticeable improvement over the 30 s rinsing process, suggesting that a constant resupply of DI water on the film surface is necessary to remove the barrier. This is consistent with prior studies comparing immersion cleaning to direct spraying, which have found that a steady water flow improves rinsing by continuously exposing the film surface to fresh solutions and allowing more efficient penetration

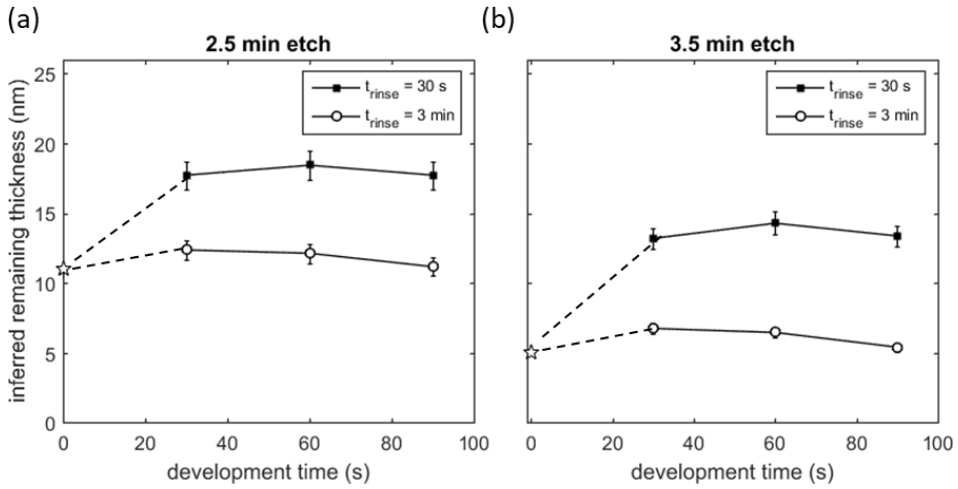

Figure 5-3: Impact of TMAH development on the reactive ion etching rate of bare NbN films in CF_4. Plots show the remaining film thickness vs. exposure time to 25% TMAH. Thicknesses (inferred from sheet resistances) were measured after (a) 2.5 min and (b) 3.5 min of RIE. The legend indicates different rinse times in DI water following exposure, while the star represents the untreated control sample. Error bars represent $\pm\sigma$, where σ is the standard deviation of the four-point resistance measurement.

into the material through shear force from the stream [82][83][84].

Additionally, the remaining thicknesses of the samples that were vigorously rinsed for 3 min hint at the possibility of the NbN film being thinned by the development process. If the etch barrier was indeed removed by the DI rinsing, the lower remaining thickness of the sample developed for 90 s in comparison to those that were developed for 30 and 60 s implies that it had a thinner initial film thickness before the etching took place, assuming the etch rate was the same between the samples. As a result, it is apparent that exposure to TMAH not only creates a barrier to the etching of NbN films but also makes a physical change on the film itself.

5.1.2 Experiment 2: material analysis

In order to explain the observed reactive ion etching trends shown in Fig. 5-3, we analyzed the composition of the etch barrier and measured the superconducting properties of NbN films exposed to TMAH.

Methods

Changes to the superconducting properties of the films were studied by measuring the R_s and T_c of the samples as a function of exposure time to 25% TMAH. To maximize the observable change with respect to the initial film thickness, we sputtered a thinner film than used previously (d = 5 nm, $R_s = 600$ Ω/square); the sheet resistances of this magnitude are similar to those of NbN films used for single photon detectors, making this study particularly relevant to the fabrication of standard superconducting devices.

Following deposition, samples of the film were submerged in 25% TMAH for durations ranging from 1 min to 25 h, followed by a vigorous 3 min DI rinse to remove the etch barrier species. Afterwards, the sheet resistance and critical temperature of each of the samples were measured. As before, R_s measurements were converted into inferred film thicknesses; for these films, ρ was estimated to be 2.9 kΩ·nm from x-ray reflectometry measurements of a sister film. A control sample with no TMAH exposure was also evaluated.

Analysis of the etch barrier composition was performed using Fourier transform infrared (FTIR) spectroscopy operated in the attenuated total reflection (ATR) mode. The sample was prepared by sputtering a thick NbN film of the same properties (R_s = 90Ω/square) as those used in Fig. 5-3 on an Si wafer. The full wafer was submerged in 25% TMAH for 4 h. To collect a film of the etch barrier, the wafer was sonicated in DI water for 20 min and rinsed in IPA. A solution of the sonicated bath was left to evaporate for 3 days, resulting in a dense liquid with visible clusters of material. Drops of this solution were then pipetted onto a clean Si water and dried on a hotplate at 90∘C for 5 min to remove excess solvents. Figure 5-4 shows a photograph of the wafer with the dried solvent. A Thermo Fisher Continuum Fourier Transform Infrared Microscope was then used to collect spectra of the material.

Results

Figures 5-5(a) and (b) display the changes in superconducting film parameters that occurred with increased exposure to 25% TMAH. Exposing the film for as little as 60 s—a common development time—led to a 0.25 nm or 5% decrease in film thickness, and 0.2 K decrease in T_c. In the extreme case of 25 h exposure, the film thickness

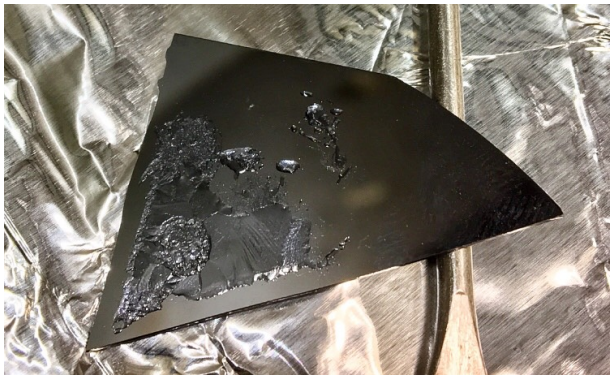

Figure 5-4: Dried solution of the etch barrier on a section of Si wafer, which was used for FTIR analysis.

decreased by nearly 3 nm or almost 57%, and the T_c decreased by roughly 1.6 K. These results suggest that the film is thinned by contact with TMAH and that even short exposure times are enough to degrade the superconducting film. Additionally, the absence of saturation in the trends of Figures 5-5(a) and (b) implies that the reaction between TMAH and the film does not only involve the thin ($\sim$1–2 nm) layer of oxide that grows on the NbN surface, which would cause the trend to plateau soon after the oxide was consumed, but rather involves the entire superconducting film.

This change may be better understood through the ATR-FTIR spectrum of the precipitated etch barrier shown in Fig. 5-5(c). In addition to the peaks characteristic of TMAH at 1490 and 951 cm^{-1}, there are low lying peaks at 838, 692, and 650 cm^{-1}. These peaks were found to be in agreement with three of the main signatures in the spectrom of a TMA-hexaniobate salt ($[(CH_3)_4N]_5[H_3Nb_6O_{19}]\cdot 20H_2O$), which has recently been synthesized by reacting hydrous niobium oxide Nb_2O_5 with TMAH in solution [85]. Remaining peaks around 1620 and 1393 cm^{-1} may correspond to the typical signatures of surface hydroxyl groups from absorbed water [86][87][88][89] and the C-H bending of methyl species [90][91], respectively, while the peak at 1000 cm^{-1} may represent the native oxide of the silicon substrate [92][93].

Identification of TMA-hexaniobate as the primary etch barrier material explains several of the phenomena reported above. First, the intrinsic water solubility of the salt accounts for the ability to remove the etch barrier through a vigorous DI water rinse. Additionally, Ref.[85] showed that TMA-hexaniobate readily forms clusters

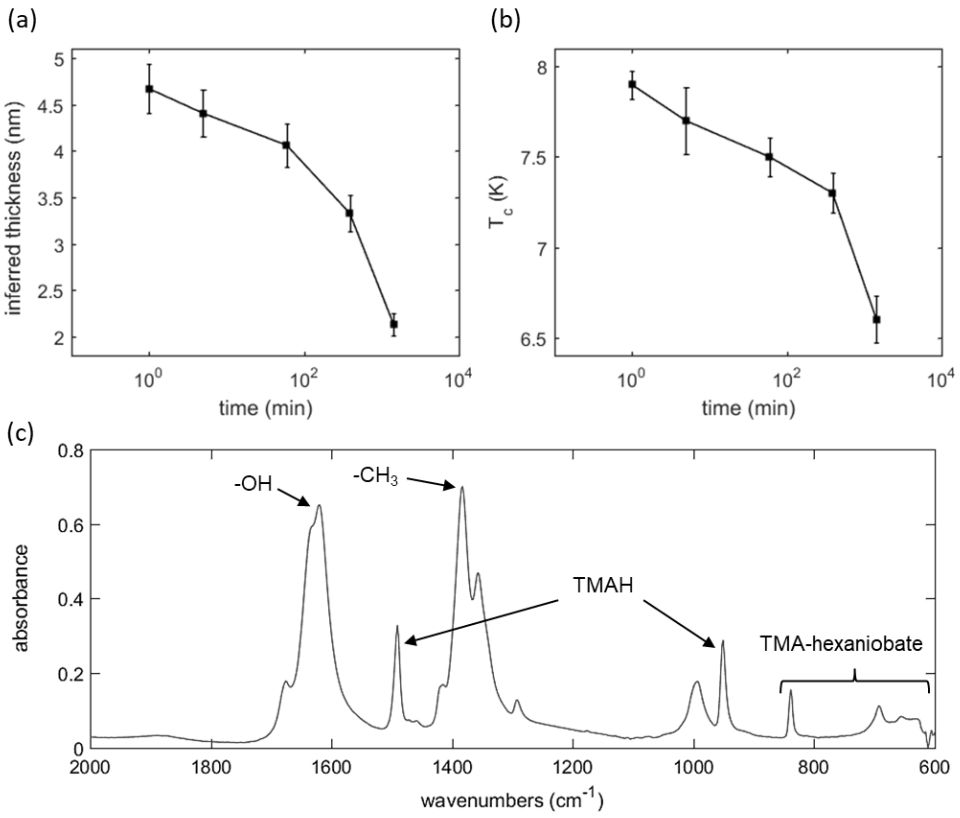

Figure 5-5: Evidence of the reaction involving NbN. (a) Inferred film thickness of a thin NbN film for various 25% TMAH development times. Error bars represent $\pm\frac{1}{2}\sigma$, where σ is the standard deviation of the four-point resistance measurement. (b) Critical temperature for the same samples shown in (a). Error bars represent $\pm\frac{1}{2}\sigma$, where σ is the standard deviation of the critical temperature measurement. (c) ATR-FTIR spectrum of precipitate from a thick NbN film left in 25% TMAH for 4 h. The spectrum reveals peaks corresponding to reported values for a TMA-hexaniobate salt.

with itself, which could manifest in the enhanced surface roughness of the treated film displayed in Fig. 5-2. Finally, an etch barrier formed from a reaction that involves Nb would account for the measured deterioration in thickness and T_c that indicated a reduction of the pure NbN film; it would also explain the similar deterioration caused by TMAH pretreatment that was reported in Ref.[94]. Thus, information from the IR spectrum revealed the mechanism behind our previous observations, which is conceptually presented in Figure 5-6.

Although there remains some debate about the precise dynamics dictating the etching of NbN in CF_4, there are certain means through which the presence of the

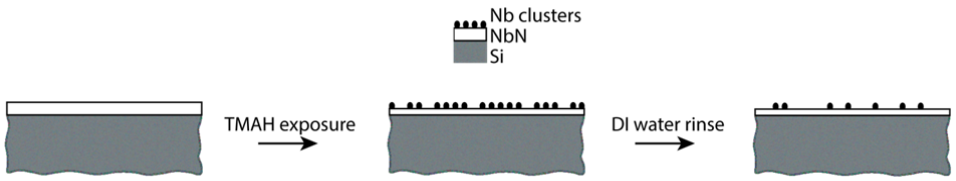

Figure 5-6: Conceptual illustration of surface modification by TMAH. Exposing the NbN film to 25% TMAH leads to a reaction that produces TMA-hexaniobate clusters, thinning the pure NbN film and creating a barrier to reactive ion etching. However, a sufficiently vigorous DI water rinse removes the clusters, leaving behind an NbN film that is thinner than the initial state. (Thanks to Marco Colangelo for the illustration.)

TMA-hexaniobate salt could interfere with the etching process. Research on reactive ion etching with CF_4 has shown that polymerization occurs as the amount of available fluorine decreases, forming unsaturated polymers that more readily adhere to surfaces [95]. This effect can be enhanced by the presence of hydrogen, producing polymer deposits such as CHF_x and blocking the underlying substrate from being etched [96]. Reduced etching via plasma polymerization has also been observed in molybdenum, which reacted with chemisorbed fluorocarbon radicals to form an etch-stop [97]. As Nb and Mo are both refractory metals, it is possible that a similar process is occuring in our samples. The TMA-hexaniobate salt may also be slowing the etch by consuming the reactant that dominates the chemical etching of NbN. Past work suggests that active atomic fluorine etches NbN by reacting to form volatile fluorides [98]; however, the etch rate was observed to decrease in the presence of CHF_3 and CH_4, which reduced the amount of available active fluorine by participating with it in a reaction [99]. Thus, it is possible that the TMA-hexaniobate salt reacts with CF_4 to form a polymer barrier that blocks the physical etching of NbN while also limiting the chemical etch by consuming the principal reactant. Further surface analysis and study of the reactive ion etching process are needed to confirm these suspicions.

5.1.3 Outlook

We observed that the 25% TMAH developer used in standard HSQ patterning adversely affects NbN films by reacting with them to form an Nb-based salt that creates

a barrier to etching in CF_4. We demonstrated that the etch barrier forms even for development times as short as 30 s, but that it can be removed through vigorous DI water rinsing. This approach offers a workable solution to the pattern transfer issues we encountered in thick NbN films; however, it does not prevent the initial film deterioration that takes place as a consequence of the reaction between TMAH and NbN. While an HSQ mask can protect the top face of a structure during development, the sides of the pattern would remain vulnerable, threatening the quality of a particularly thin or narrow feature. Our work also reveals that processes which use a 25% TMAH pretreatment before spinning HSQ in order to promote adhesion deteriorate the quality of the NbN before patterning even takes place. For particularly thin films like those used in the fabrication of single photon detectors, even minor deterioration of the surface could bear significant consequences on device performance.

Given these limitations, it would be beneficial to evaluate if more dilute TMAH-based developers such as MF CD-26 have the same effect or if they could be used as substitutes. Alternative developers of HSQ such as salty developers [100] may be used, but their influence on superconducting films has not yet been reported. It may also be worthwhile to investigate if other etch processes such as Ar ion milling of Cl_2 reactive ion etching could replace CF_4 in order to avoid the vigorous 3 min DI water rinse required to remove the etch barrier. Finally, adding a protective layer between the NbN and HSQ could prevent direct contact with the developer.

While further work is needed to create an HSQ fabrication process that poses no harm to NbN films, these experiments nonetheless explored steps that can be taken to classify and ameliorate an unknown etch contaminant in superconducting materials. By using FTIR, it was possible to obtain a spectrum of the precipitated contaminant, while measurements of film parameters, etch characteristics, and surface features provided evidence to support the contaminant's identity. Although this investigation was specific to NbN, it is possible that other materials such as tantalum undergo similar processes, motivating future studies of the same nature on a wide range of films.

5.2 Use of ma-N as a negative tone EBL resist

While the previous section revealed problems caused by the pretreatment and development of HSQ with TMAH, there are additional limitations in using HSQ with superconducting materials. First, exposed HSQ cannot be easily removed without the use of a strong reagent such as hydrofluoric acid, which damages the quality of the superconducting film. The inability to remove HSQ is particularly hindering as nanowires are gaining popularity in more complex circuits that involve integration with other technologies, such as magnetic memory elements [24]. Such complex circuit environments often demand multilayer processes, which require exposed resist to be removed between layers. Since HSQ cannot be stripped without damanging the underlying device, it is clearly unsuitable for these applications. Additionally, HSQ has a high areal dose density ($\sim$3500–4000 μC/cm^2 in a 125 kV EBL system) in comparison to many positive tone resists, which prolongs the write time for complex circuits with many components. As a result, there is a need for alternative negative-tone electron-beam resists to support the advancement of superconducting devices towards more complex applications.

Here, we report on the use of ma-N 2400 series deep ultraviolet photoresist (@microresist technology) as an electron-beam resist for patterning superconducting nanoscale devices. ma-N 2400 series resist is primarily composed of a polymeric bonding agent (phenolic resin) and a bisazide photoactive compound [101]. In contrast to HSQ, exposed ma-N can be removed with solvents like *N*-methyl-2-pyrrolidone (NMP), which is generally harmless to superconducting films like NbN, allowing for multilayer processing. Additionally, ma-N's dose is roughly four times lower than that of HSQ in 125 kV EBL systems, which could significantly reduce the write time for large patterns. While previous studies have reported minimum feature sizes of 50 nm or more [101][102], we were able to pattern repeated lines of widths down to 30 nm and individual features with minimum dimensions less than 20 nm. To demonstrate ma-N's suitability for patterning superconducting devices, we compared the performance and geometry of 36 identical devices and concluded that the reproducibility was sat-

isfactory. These efforts also enrich our understanding of superconducting nanowire uniformity in general, since there are very few systematic studies on lithographic variations in nanowire patterning.

5.2.1 Resist characterization

To discern whether ma-N was a suitable resist for patterning nanoscale features, we studied its resolution and characterized its etch resistance in a fluorine chemistry relative to that of HSQ. The process used to pattern with ma-N aligned closely with recipes reported in previous literature [102][103].

All chips were written using ma-N 2401, the most dilute of the ma-N 2400 series. Prior to exposure, ma-N 2401 was spun at 3 krpm for 60 s, yielding a thickness of about 92 nm, and then baked at 90 °C for 60 s. Afterwards, patterns were written using a 125 kV Elionix system (ELS-F125) with a 500 pA beam current, corresponding to a spot size of 2 nm. Doses ranged from 800 to 1000 μC/cm^2, depending on the pattern. Following exposure, chips were developed in Microposit MF CD-26 at room temperature for 10 s with gentle agitation by hand and immediately rinsed in DI water and blown dry using a nitrogen gun. The resolution was assessed by examining the patterns using a scanning electron microscope (Zeiss Sigma HD). Some features written only in ma-N 2401 without an underlying superconducting film were sputtered with gold-palladium prior to imaging to enhance contrast. The resist could be stripped by submerging the chip in heated NMP at 60 °C for 1 h, offering a significant advantage over HSQ. We did not observe any adverse effects of prolonged immersion in heated NMP on the quality of the quality of superconducting films.

Using the methods described above, we patterned a variety of geometries on a 1×1 cm^2 Si die with 300-nm-thick thermal oxide. Figure 5-7 displays several of these features. As shown in Fig. 5-7(a), it was possible to resolve 30-nm lines with a 120-nm pitch, a smaller resolution than what has been previously claimed [101][102]. Besides repeated lines, we also patterned isolated features. Fig. 5-7(c) shows an example of one such feature, a T-shaped pattern with the same geometry the nanocryotron (nTron) [23]. As indicated on the figure, the minimum feature size of this pattern

was ∼18 nm, measured using image analysis software (PROSEM). This result demonstrates that the resolution of ma-N resist is better than what had been previously thought and that it is suitable for patterning isolated features narrower than 20 nm.

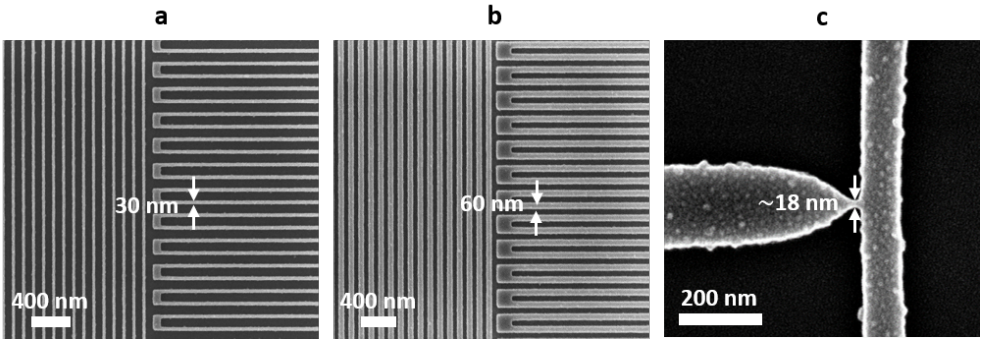

Figure 5-7: Resolution of ma-N 2401 patterned using 125 kV Elionix. (a) 30-nm lines with a 120-nm pitch. (b) 60-nm lines with a 120-nm pitch.(c) Three-terminal pattern with a minimum isolated feature size of ∼18 nm. The granularity is due to the gold-palladium sputtering that was used to enhance imaging contrast on the dielectric substrate.

To ensure that ma-N 2401 could hold up during pattern transfer to underlying superconducting films, we compared its etch resistance to that of HSQ, our traditional negative tone resist. We first spun 4% HSQ and ma-N onto two identical Si chips, both at 3 krpm for 60 s. After spinning, the ma-N chip was baked as described above, and the initial thickness of each resist was approximated using a multiwavelength ellipsometer (FilmSense FS-1) with a Cauchy model. The HSQ had an initial film thickness of about 63 nm, and the ma-N had an initial thickness of about 92 nm.

Afterwards, we reactively ion etched both chips together in CF_4 in 1 min intervals and measured their thicknesses using the ellipsometer after each step. The process was repeated until the resist with the highest etch resistance was reduced to 50% of its initial thickness. As shown in Figure 5-8, the 4% HSQ had an etch rate of roughly 14.5 nm/min, while the ma-N had an etch rate of 8.9 nm/min. It is important to note that both of these resists were unexposed, so it is possible that their etch resistances would be slightly higher following exposure [104]. However, the results shown here indicate that ma-N's etch resistance is comparable to, if not better than, that of HSQ for a fluorine-based etch, allowing us to conclude that choosing ma-N over HSQ to

pattern superconducting devices does not sacrifice the fidelity of the pattern transfer.

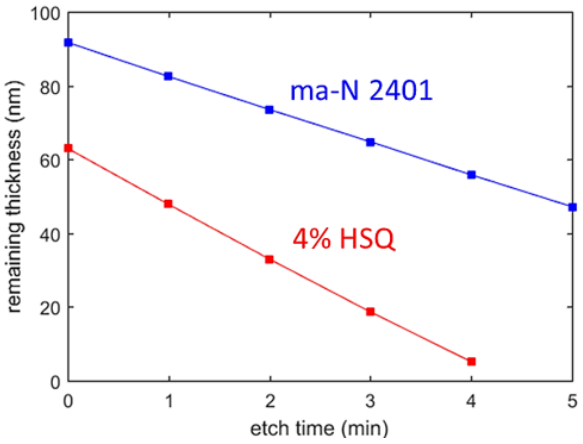

Figure 5-8: Etch rate in CF_4 of ma-n 2401 and 4% HSQ (both unexposed). The etch rate of ma-N 2401 was 8.9 nm/min, and the etch rate of 4% HSQ was 14.5 nm/min.

5.2.2 Device reproducibility

To test ma-N's reproducibility, we patterned 36 identical nTrons on the same chip and compared their performance and minimum dimensions.

To fabricate the devices, we first used an AJA sputtering system to deposit an NbN film with a sheet resistance of 150 Ω/square on a 1×1 cm^2 die cut from a 4 in. Si wafer. Prior to patterning the superconducting devices, 50-nm-thick gold pads for making electrical contact were fabricated using a bilayer photoresist process followed by lift-off. Afterwards, nanowire structures were patterned using the ma-N 2401 EBL process described above, followed by reactive ion etching in CF_4. Each device was imaged, and the resulting SEMs were processed using PROSEM software to determine the minimum dimensions.

After imaging, the current-voltage characteristics of the devices were measured in a cryogenic probe station at 4.89 K. A bias current was applied using a DC battery source (Stanford Research Systems SIM928) in series with a 100 kΩ resistor. The bias current was ramped from 0 to 30 μA, and the voltage was read out through a Keithley multimeter.

Figure 5-9(a) shows a micrograph of one of the devices. As mentioned previously,

nTrons are three-terminal devices that act similarly to a comparator, and their functionality depends heavily on their minimum dimension, the gate. To trigger an nTron, an input signal greater than the gate's switching current is applied to the gate, creating a hotspot that then causes the channel to switch from superconducting to normal [23]. An example of the current-voltage characteristics of one of the gates, indicating its switching current I_{sw} is shown in Figure 5-9(b). Since the ideal switching current of a superconducting structure is the product of the film's critical current density and the structure's width (ignoring confounding effects like nonuniformities), the width of the gate is critical in determining when the nTron fires. By comparing the width and switching current of each nTron's gate, we can therefore get a sense of how reliably ma-N resolves identical sub-50-nm features.

To compare gate widths, we used PROSEM to measure the minimum dimension of each gate's micrograph and to generate each gate's edge profile. Figure 5-9(c) displays the profiles of all 36 gates, offset from each other in the z-axis for clarity. The x-axis represents the horizontal distance across the gate, while the y-axis shows the vertical distance, which also defines the gate's width. By visually comparing the profile of each gate, we can see edge deviations from the desired pattern and observe differences between the structures in a way that is not obvious from simply looking at side-by-side micrographs.

Although the edge profiles provide qualitative evidence of variability between devices, the minimum dimensions and switching currents offer a more quantitative and sensitive comparison. Figure 5-10(a) shows a histogram of the switching currents for all 36 devices, with a mean of 15.88 μA and a standard deviation of 3.13 μA. These statistics are typical of our nanowire devices, which usually have a standard deviation of 10%–20% of the mean [94]. Figure 5-10(b) displays a histogram of the minimum gate dimensions measured using PROSEM, with an average of 33.7 nm and a standard deviation of 2.4 nm, suggesting that the patterned geometry is less variable than the switching current.

As a possible explanation of the spread in switching currents, we can plot the switching current of each device as a function of its minimum width, as shown in

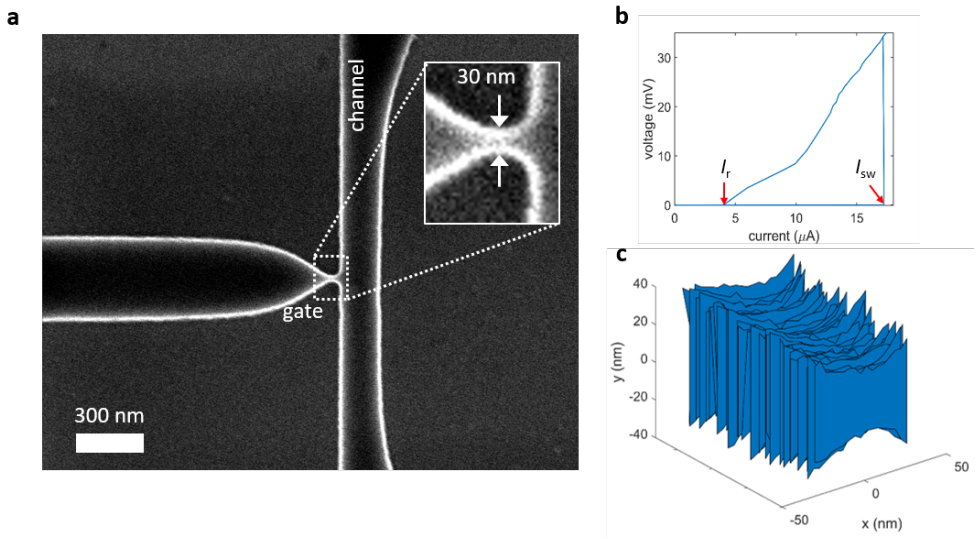

Figure 5-9: Three-terminal superconducting devices patterned using ma-N 2401. (a) SEM showing one of the 36 devices. The inset shows an enlarged view of the narrowest feature, the gate. (b) Current-voltage characteristics of the gate of one of the devices, indicating the switching current I_{sw} and the retrapping current I_r. (c) Edge profiles of the gates of all 36 patterned devices, obtained using PROSEM software. The x-axis represents the horizontal distance across the gate, while the y-axis shows the vertical distance, which also defines the gate's width.

Figure 5-10 (c). The trend looks approximately linear, which agrees with our expectations, since the switching current is the product of the current density and the minimum width. However, there are some outliers from the trend, most notably the points that have a lower switching current than expected. Such premature switching could be caused by a variety of factors, such as material inhomogeneities or grain boundaries in the polycrystalline NbN that cannot be observed by SEM. Nanowires are also susceptible to thermal fluctuations and noise, which may explain why there is a wider spread in switching current than in minimum dimension [64]. Despite these variations, the switching current statistics align with our typical device performance [94], indicating that ma-N can produce superconducting devices like nTrons with similar reproducibility to what we obtain with standard resists.

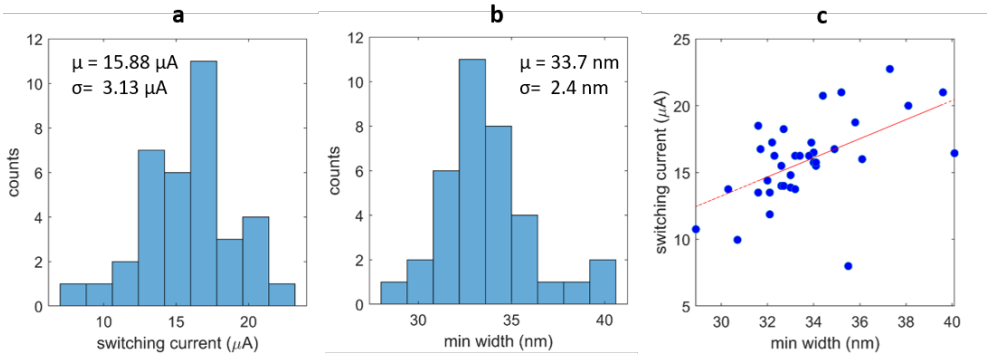

Figure 5-10: Reproducibility of the gate terminal across all 36 patterned devices. (a) Histogram of the switching currents of the gates. (b) Histogram of the minimum width across each gate. (c) Plot of the switching current of each device with respect to its minimum measured gate dimension. The line indicates that deviations from the mean switching current have a roughly linear relationship with deviations in minimum width, with a few outliers that could be due to noise or material defects.

5.2.3 Outlook

The results presented in this thesis suggest that ma-N 2401 is a suitable negative tone resist for patterning superconducting devices. In comparison to HSQ, ma-N offers advantages including a lower required dose and the ability to be removed by a solvent that does not damage superconducting films, which is critical for multilayer processes. It may also be cheaper and have a longer shelf life than HSQ, and could allow for the combination of both optical and electron-beam lithographies. To demonstrate ma-N's potential, we showed that it can reliably pattern dense 30-nm-wide lines and individual features below 20 nm, which is a better resolution than what has been previously reported. However, HSQ remains the gold standard for resolution and might have superior etch resistance to ma-N in other etch chemistries like chlorine. The advantages of HSQ in comparison to ma-N are similar to those that it holds over other polymer-based resists, and stem from HSQ's underlying structure as a small, cage-like oligomer [73]. Nevertheless, our results indicate that ma-N is a sufficient EBL resist for attaining the dimensions required by most superconducting nanowire devices.

Widespread adoption of ma-N as an electron-beam resist could spur the implementation of nanowire devices in complex circuits that use multilayer processes. Since

many of these circuits involve integration with other devices, such as standard electronic or magnetic elements, we envision these results may encourage collaboration across different platforms. Additionally, the multilayer processes supported by ma-N could inspire the creation of new nanowire devices that incorporate other materials to modify superconducting dynamics. Examples of recent nanowire devices that use normal metal to modulate superconductivity have all been patterned using positive tone resist, partially due to the limitations of HSQ [32][34]; as a result, ma-N could facilitate the development of new devices that use multilayer processes. Finally, the use of ma-N is not limited to superconducting films. We have demonstrated that it has $\sim$30 nm resolution and is processed with solvents that are compatible with many materials, suggesting that it could be used for a wide variety of applications beyond those presented here. Overall, ma-N 2401 appears to be a suitable electron-beam resist for patterning nanoscale features and has the potential to advance of complexity of superconducting nanowire-based circuits.

Chapter 6

Conclusion and outlook

As traditional computing approaches its ultimate limits, we enter an uncertain but exciting time for scientific investigation, in which hardware, algorithms, and architectures are evolving simultaneously. Inspired by the wealth of technology that grew from the material behavior of transistors, this new age of computing is an opportunity to closely study the natural dynamics of materials and devices, and harness them to create alternative computing schemes.

In this thesis, we have presented two new devices that may serve as the foundation of an alternative cryogenic computing architecture using superconducting nanowires. The first is a multilevel memory cell that could enable denser cryogenic memory, and allow superconducting logic to no longer rely on silicon-based memories at room temperature. The second is an artificial neuron using two nanowire relaxation oscillators for implementation in spiking neural networks. Both devices operate by taking advantage of characteristics that are unique to superconducting nanowires, such as nonlinear switching and high kinetic inductance. Together, these technologies have the potential to unite memory and processing under a single platform and even within the same fabrication steps, which could alleviate the bottleneck in computation speed caused by the physical distance between memory and processing units in von Neumann architectures. Additionally, a complete superconducting nanowire architecture using spiking neurons has the potential to achieve superior energy performance to CMOS while also being compatible with it. As a result, nanowires may offer a path

forward towards alternative computing without requiring complete abandonment of traditional systems.

To expand these devices into a large-scale nanowire-based architecture, it will be necessary to study critical aspects like fan-out and fan-in limitations, and the power dissipation costs of a bias distribution network. Further exploration should also study operating margins with respect to fabrication variability and noise. In order for the nanowire neuron to be useful in tasks beyond inference, it will also be crucial to develop a scheme for unsupervised learning so that synaptic weights can be adjusted on-chip in real time. Finally, it may be useful to create a more abstract, theoretical model of the nanowire neuron so that it can be more readily applied to algorithm development. While these steps remain the subject of future investigations, they will provide key insight into how a complete nanowire-based computing system may be realized.

The two neural network applications discussed in this thesis demonstrate examples where spiking dynamics offer key advantages, such as the WTA competition using stochasticity to its benefit rather than fighting against it. However, these cases only scratch the surface of how spike-based communication can be exploited. Specifically, neither of these applications or the majority of spiking neural network algorithms take advantage of the wealth of information stored in the timing of spiking signals. Instead, information is either encoded in the spiking state ("firing" versus "not firing"), or in aggregate characteristics like spiking frequency, since these are more compatible with algorithms designed for traditional hardware. The use of time domain information would not only be extremely relevant for technologies like event-based sensors in robotics, but could also improve energy efficiency and illuminate how spike timing plays a role in communication within the human brain. The emergence of new temporal logic families like race logic [105][106] offer hope that spike-based communication could soon be used to its full potential.

Besides spiking neural networks, other alternative computing schemes such as stochastic computing [107] [108], adiabatic computing [109][110], and coupled oscillator computing [111] are actively being developed and examined for their own unique

advantages. Like the nanowire devices presented here, each of these approaches offers the chance to be curious about the natural dynamics of devices and explore how they can be used to achieve what was not previously possible.

Appendix A

Neuron synapse equations

Impedance of the inductors and the oscillators:

$$Z_l = i\omega L_k \tag{A.1a}$$

$$Z_{oscillator} = \left(\frac{1}{R_s} + \frac{2}{Z_l}\right)^{-1} \tag{A.1b}$$

Impedance to the left and right of the synapses:

$$Z_{left} = \frac{Z_{oscillator}(2Z_l + Z_{oscillator})}{2Z_{oscillator} + 2Z_l} \tag{A.2a}$$

$$Z_{right} = \frac{Z_l(2Z_{oscillator} + Z_l)}{2Z_{oscillator} + 2Z_l} \tag{A.2b}$$

Capacitive synapse equations:

$$\tau_C = C\left(\frac{1}{R_1 + Z_{left}} + \frac{1}{R_2 + Z_{right}}\right)^{-1} \tag{A.3a}$$

$$Z_{in,C} = R_1 + \left(i\omega C + \frac{1}{R_2 + Z_{right}}\right)^{-1} \tag{A.3b}$$

Inductive synapse equations:

$$\tau_L = \frac{L}{\left(\frac{1}{R_{1L}} + \frac{1}{Z_{left}}\right)^{-1} + \left(\frac{1}{R_{2L}} + \frac{1}{Z_{right}}\right)^{-1}} \tag{A.4a}$$

$$Z_{in,L} = \left(\frac{1}{R_{1L}} + \frac{1}{i\omega L + \left(\frac{1}{R_{2L}} + \frac{1}{Z_{right}} \right)^{-1}} \right)^{-1} \tag{A.4b}$$

To make the input impedances the same as $\omega \to 0$, a series resistance must be added to the inductive synapse:

$$\tau_L = \frac{L}{\left(\frac{1}{R_{1L}} + \frac{1}{Z_{left}+R_{ser}} \right)^{-1} + \left(\frac{1}{R_{2L}} + \frac{1}{Z_{right}} \right)^{-1}} \tag{A.5a}$$

$$Z_{in,L} = R_{ser} + \left(\frac{1}{R_{1L}} + \frac{1}{i\omega L + \left(\frac{1}{R_{2L}} + \frac{1}{Z_{right}} \right)^{-1}} \right)^{-1} \tag{A.5b}$$

Overall, we obtain two sets of equations that we can use to solve for the approximate inductance and series resistance, if we set examine the extreme frequency limits of $\omega \to 0$ and $\omega \to \infty$.

For $\omega \to 0$:

$$R_{ser} = R_1 + R_2 \tag{A.6a}$$

$$L = (C * R_1 * R_{1L} * R_{2L})/(R_1 + R_{1L} + R_2) \tag{A.6b}$$

For $\omega \to \infty$:

$$R_{ser} = R_1 - R_{1L} \tag{A.7a}$$

$$L = C * (-R_{1L}^2 + R_1(R_{1L} + R_{2L}) + R_1 * R_s + R_{2L} * R_s) \tag{A.7b}$$

Solving for these equations with the capacitive synapse parameters in Fig. 3-10 and letting $R_{1L} = R_{2L} = 40\ \Omega$, the first set of equations gives $R_{ser} = 20\ \Omega$ and $L = 0.267\ \mu$H, while the second set of equations gives R_{ser} = -30 Ω and L = -0.300 μH. Although these are just approximations, they give an idea of the appropriate orders of magnitude for the equivalent circuit.

Appendix B

Python code for 9-pixel test

B.1 Code to create test images and calculate weights

This code is used to create the set of 30 test images and test them in the simple neural network, using code from Ref. [67]. The output is a list of weights that can later be mapped onto synaptic inductance values.

```
# -*- coding: utf-8 -*-
"""
@author: Emily
"""
## Script to compile all of the pixel data
## Note: we adopt the pixel numbering as 1-3 = top row,
   4-6 = middle row, 7-9 = bottom row

import numpy as np
import network


## Image list for Z
```

```
zim1 = np.array([1, 1, 0, 0, 1, 0, 0, 1, 1])

x = np.arange(len(zim1))

Z_total = np.copy(zim1)
## iterate to create each of the single-pixel error
for n in np.nditer(x):
    print(n)
    arr = []
    arr = np.copy(zim1)
    if zim1[n] == 0:
        arr[n] = 1
    else:
        arr[n] = 0
    Z_total = np.vstack((Z_total,arr))



## Image list for V

vim1 = np.array([1, 0, 1, 1, 0, 1, 0, 1, 0])

x = np.arange(len(vim1))

V_total = np.copy(vim1)

for n in np.nditer(x):
    print(n)
    arr = []
    arr = np.copy(vim1)
```

```
    if vim1[n] == 0:
        arr[n] = 1
    else:
        arr[n] = 0
    V_total = np.vstack((V_total,arr))


## Image list for N

nim1 = np.array([0,1,0,1,0,1,1,0,1])

x = np.arange(len(nim1))

N_total = np.copy(nim1)

for n in np.nditer(x):
    print(n)
    arr = []
    arr = np.copy(nim1)
    if nim1[n] == 0:
        arr[n] = 1
    else:
        arr[n] = 0
    N_total = np.vstack((N_total,arr))


## Correct output list-- let Z = 1, V = 2, and N = 3
zs = 0*np.ones(len(nim1)+1)
vs = 1*np.ones(len(nim1)+1)
ns = 2*np.ones(len(nim1)+1)
```

```
letters =np.concatenate((zs,vs,ns))
letters = letters.astype(int)

def vectorized_result_short(j):
    # return a 3-dimensional unit vector with 1 in the jth
        position and zeroes elsewhere. Used to convert the
        training set
    e = np.zeros((3,1))
    e[j] = 1
    return e

## Create the training and test data

inputs = np.vstack((Z_total, V_total, N_total))
total_inputs = [np.reshape(x,(9,1)) for x in inputs] #
   reshape to make into vertical

training_results = [vectorized_result_short(y) for y in
   letters]

## Zip into tuples

training_data = list(zip(total_inputs,training_results))

test_data = list(zip(total_inputs, letters))




## Run the network test
```

```
#import network
net = network.Network([9, 3])

# run program
(w_final, b_final) = net.SGD(training_data, 50,5,1,
   test_data = test_data)

# output weights
print(w_final)
```

B.2 Code to test images in LTSpice environment

This code is used to translate the 30 test images into lists that can be copied and imported into the LTSpice environment for testing.

The first blocks of code can be used to interpret each pixel color as the bias current to each pixel neuron. It is used to test one image per simulation. A list of ".param" commands is generated that can be copied and directly inserted into the Spice environment. A set of asterisks is used to denote a new image list.

The last section of code interprets pixel colors as input currents, and generates piecewise-linear (PWL) functions for each current source in order to test all 30 images in one simulation. It generates a PWL list of currents and timestamps for each pixel that can be coped and saved into a .txt file, and imported in into the Spice environment for each current source.

```
# -*- coding: utf-8 -*-
"""
@author: Emily
"""
```

```
import numpy as np


### FOR Z PATTERNS  #####
zim1 = np.array([1, 1, 0, 0, 1, 0, 0, 1, 1])

x = np.arange(len(zim1))

Z_total = np.copy(zim1)

for p in np.nditer(np.arange(len(zim1))):
        str = ".param b{} = pbias*{}".format(p,zim1[p])
        print(str)
        if p == 8:
            print("***********")

### FOR INTERPRETING PIXEL COLORS AS BIAS CURRENTS,
   RUNNING DIFFERENT LTSPICE SIM FOR EACH IMAGE###
## for printing the bias currents for the .param list
for n in np.nditer(x):
    arr = []
    arr = np.copy(zim1)
    if zim1[n] == 0:
        arr[n] = 1
    else:
        arr[n] = 0
    Z_total = np.vstack((Z_total,arr))
    for p in np.nditer(np.arange(len(arr))):
        str = ".param b{} = pbias*{}".format(p,arr[p])
        print(str)
```

```
        if p == 8:
            print("***********")


### FOR V PATTERNS  #####

vim1 = np.array([1, 0, 1, 1, 0, 1, 0, 1, 0])

x = np.arange(len(vim1))

V_total = np.copy(vim1)

for p in np.nditer(np.arange(len(vim1))):
        str = ".param b{} = pbias*{}".format(p,vim1[p])
        print(str)
        if p == 8:
            print("***********")

### FOR INTERPRETING PIXEL COLORS AS BIAS CURRENTS,
   RUNNING DIFFERENT LTSPICE SIM FOR EACH IMAGE###
## for printing the bias currents for the .param list
for n in np.nditer(x):
    arr = []
    arr = np.copy(vim1)
    if vim1[n] == 0:
        arr[n] = 1
    else:
        arr[n] = 0
    V_total = np.vstack((V_total,arr))
    for p in np.nditer(np.arange(len(arr))):
```

```
        str = ".param b{} = pbias*{}".format(p,arr[p])
        print(str)
        if p == 8:
            print("***********")



### FOR N PATTERNS  #####

nim1 = np.array([0,1,0,1,0,1,1,0,1])

x = np.arange(len(nim1))

N_total = np.copy(nim1)

for p in np.nditer(np.arange(len(nim1))):
        str = ".param b{} = pbias*{}".format(p,nim1[p])
        print(str)
        if p == 8:
            print("***********")

### FOR INTERPRETING PIXEL COLORS AS BIAS CURRENTS,
   RUNNING DIFFERENT LTSPICE SIM FOR EACH IMAGE###
## for printing the bias currents for the .param list

for n in np.nditer(x):
    arr = []
    arr = np.copy(nim1)
    if nim1[n] == 0:
        arr[n] = 1
```

```
    else:
        arr[n] = 0
    N_total = np.vstack((N_total,arr))
    for p in np.nditer(np.arange(len(arr))):
        str = ".param b{} = pbias*{}".format(p,arr[p])
        print(str)
        if p == 8:
            print("***********")

### FOR INTERPRETING PIXEL COLORS AS INPUT CURRENTS AND
   RUNNING ALL IMAGES AT ONCE ###
###  CREATE LIST OF ALL PATTERNS   ####
Lett_total = np.vstack((Z_total,V_total,N_total))

## Create PWL list for all pixels that can be copied and
   saved into a .txt file for inputting into LTspice##
for n in np.nditer(x):
    for p in np.nditer(np.arange(len(Lett_total[:,0]))):
        val = Lett_total[p,n]
        if p == 0:
            print("Pixel %d\n0 0" %n)
        multilinestr = "+5n %d\n+50n %d\n+5n 0" %(val,val)
        print(multilinestr)
        if p == len(Lett_total)-1:
            print("***********")
```

Bibliography

[1] Nihar R. Mahapatra and Balakrishna Venkatrao. "The processor-memory bottleneck: problems and solutions". In: *XRDS: Crossroads, The ACM Magazine for Students* 5.3 (Apr. 1, 1999), 2–es. ISSN: 1528-4972. DOI: 10.1145/357783.331677. URL: https://doi.org/10.1145/357783.331677 (visited on 03/16/2020).

[2] dcadmin. *Rebooting the IT Revolution: A Call to Action*. Semiconductor Industry Association. URL: https://www.semiconductors.org/resources/rebooting-the-it-revolution-a-call-to-action-2/ (visited on 01/13/2020).

[3] Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. "Towards spike-based machine intelligence with neuromorphic computing". In: *Nature* 575.7784 (Nov. 2019), pp. 607–617. ISSN: 1476-4687. DOI: 10.1038/s41586-019-1677-2. URL: https://www.nature.com/articles/s41586-019-1677-2 (visited on 01/22/2020).

[4] Ilias Sourikopoulos et al. "A 4-fJ/Spike Artificial Neuron in 65 nm CMOS Technology". In: *Frontiers in Neuroscience* 11 (2017). ISSN: 1662-453X. DOI: 10.3389/fnins.2017.00123. URL: https://www.frontiersin.org/articles/10.3389/fnins.2017.00123/full (visited on 02/27/2019).

[5] R. Matsumoto et al. "Chaos and Relaxation Oscillations in Spin-Torque Windmill Spiking Oscillators". In: *Phys. Rev. Applied* 11.4 (Apr. 2019), p. 044093. DOI: 10.1103/PhysRevApplied.11.044093. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.11.044093.

[6] Hongsik Jeong and Luping Shi. "Memristor devices for neural networks". In: *Journal of Physics D: Applied Physics* 52.2 (Oct. 2018), p. 023003. DOI: 10.1088/1361-6463/aae223. URL: https://doi.org/10.1088%2F1361-6463%2Faae223.

[7] Andy Thomas. "Memristor-based neural networks". In: *Journal of Physics D: Applied Physics* 46.9 (Feb. 5, 2013), p. 093001. ISSN: 0022-3727. DOI: 10.1088/0022-3727/46/9/093001. URL: http://dx.doi.org/10.1088/0022-3727/46/9/093001.

[8] B. D. Josephson. "Possible new effects in superconductive tunnelling". In: *Physics Letters* 1.7 (July 1, 1962), pp. 251–253. ISSN: 0031-9163. DOI: 10.1016/0031-9163(62)91369-0. URL: http://www.sciencedirect.com/science/article/pii/0031916362913690 (visited on 06/20/2017).

[9] K. K. Likharev. "Rapid Single-Flux-Quantum Logic". In: *The New Superconducting Electronics*. NATO ASI Series. Springer, Dordrecht, 1993, pp. 423–452. ISBN: 978-94-010-4848-4 978-94-011-1918-4. DOI: 10.1007/978-94-011-1918-4_14. URL: https://link.springer.com/chapter/10.1007/978-94-011-1918-4_14 (visited on 08/10/2018).

[10] Oliver T. Oberg et al. "Integrated Power Divider for Superconducting Digital Circuits". In: *IEEE Transactions on Applied Superconductivity* 21.3 (June 2011), pp. 571–574. ISSN: 2378-7074. DOI: 10.1109/TASC.2010.2086415.

[11] Naoki Takeuchi et al. "An adiabatic quantum flux parametron as an ultra-low-power logic device". In: *Superconductor Science and Technology* 26.3 (2013), p. 035010. ISSN: 0953-2048. DOI: 10.1088/0953-2048/26/3/035010. URL: http://stacks.iop.org/0953-2048/26/i=3/a=035010 (visited on 08/10/2018).

[12] D. Scott Holmes, Andrew L. Ripple, and Marc A. Manheimer. "Energy-Efficient Superconducting Computing—Power Budgets and Requirements". In: *IEEE Transactions on Applied Superconductivity* 23.3 (June 2013), pp. 1701610–1701610. ISSN: 2378-7074. DOI: 10.1109/TASC.2013.2244634.

[13] Qing-Yuan Zhao et al. "A nanocryotron comparator can connect single-flux-quantum circuits to conventional electronics". In: *Superconductor Science and Technology* 30.4 (2017), p. 044002. ISSN: 0953-2048. DOI: 10.1088/1361-6668/aa5f33. URL: http://stacks.iop.org/0953-2048/30/i=4/a=044002 (visited on 03/13/2018).

[14] Quentin P. Herr et al. "Ultra-low-power superconductor logic". In: *Journal of Applied Physics* 109.10 (May 15, 2011), p. 103903. ISSN: 0021-8979. DOI: 10.1063/1.3585849. URL: https://aip.scitation.org/doi/abs/10.1063/1.3585849 (visited on 08/10/2018).

[15] M. H. Volkmann et al. "Implementation of energy efficient single flux quantum digital circuits with sub-aJ/bit operation". In: *Superconductor Science and Technology* 26.1 (Nov. 2012). Publisher: IOP Publishing, p. 015002. ISSN: 0953-2048. DOI: 10.1088/0953-2048/26/1/015002. URL: https://doi.org/10.1088%2F0953-2048%2F26%2F1%2F015002 (visited on 04/28/2020).

[16] H. Terai et al. "The effects of DC bias current in large-scale SFQ circuits". In: *IEEE Transactions on Applied Superconductivity* 13.2 (June 2003). Conference Name: IEEE Transactions on Applied Superconductivity, pp. 502–506. ISSN: 1558-2515. DOI: 10.1109/TASC.2003.813917.

[17] D. A. Buck. "The Cryotron-A Superconductive Computer Component". In: *Proceedings of the IRE* 44.4 (Apr. 1956), pp. 482–493. ISSN: 2162-6634. DOI: 10.1109/JRPROC.1956.274927.

[18] D. R. Young. "Recent developments in high speed superconducting devices". In: *British Journal of Applied Physics* 12.8 (Aug. 1961), pp. 359–362. ISSN: 0508-3443. DOI: 10.1088/0508-3443/12/8/302. URL: https://doi.org/10.1088%2F0508-3443%2F12%2F8%2F302 (visited on 03/10/2020).

[19] A.E. Brennemann. "The in-line cryotron". In: *Proceedings of the IEEE* 51.3 (Mar. 1963), pp. 442–451. ISSN: 1558-2256. DOI: 10.1109/PROC.1963.1847.

[20] R. F. Broom and E. H. Rhoderick. "Thermal propagation of a normal region in a thin superconducting film and its application to a new type of bistable element". In: *British Journal of Applied Physics* 11.7 (July 1960), pp. 292–296. ISSN: 0508-3443. DOI: 10.1088/0508-3443/11/7/310. URL: https://doi.org/10.1088%2F0508-3443%2F11%2F7%2F310 (visited on 03/03/2020).

[21] G. N. Gol'tsman et al. "Picosecond superconducting single-photon optical detector". In: *Applied Physics Letters* 79.6 (Aug. 1, 2001), pp. 705–707. ISSN: 0003-6951. DOI: 10.1063/1.1388868. URL: http://aip.scitation.org/doi/abs/10.1063/1.1388868 (visited on 07/30/2017).

[22] Francesco Marsili et al. "Single-Photon Detectors Based on Ultranarrow Superconducting Nanowires". In: *Nano Letters* 11.5 (May 11, 2011), pp. 2048–2053. ISSN: 1530-6984. DOI: 10.1021/nl2005143. URL: http://dx.doi.org/10.1021/nl2005143 (visited on 10/13/2015).

[23] Adam N. McCaughan and Karl K. Berggren. "A Superconducting-Nanowire Three-Terminal Electrothermal Device". In: *Nano Letters* 14.10 (Oct. 8, 2014), pp. 5748–5753. ISSN: 1530-6984. DOI: 10.1021/nl502629x. URL: http://dx.doi.org/10.1021/nl502629x (visited on 09/07/2016).

[24] Minh-Hai Nguyen et al. "Cryogenic Memory Architecture Integrating Spin Hall Effect based Magnetic Memory and Superconductive Cryotron Devices". In: *Scientific Reports* 10.1 (Jan. 14, 2020), pp. 1–11. ISSN: 2045-2322. DOI: 10.1038/s41598-019-57137-9. URL: https://www.nature.com/articles/s41598-019-57137-9 (visited on 03/10/2020).

[25] Adam N. McCaughan et al. "Using Geometry To Sense Current". In: *Nano Letters* 16.12 (Dec. 14, 2016), pp. 7626–7631. ISSN: 1530-6984. DOI: 10.1021/acs.nanolett.6b03593. URL: http://dx.doi.org/10.1021/acs.nanolett.6b03593 (visited on 02/03/2017).

[26] Qing-Yuan Zhao et al. "A nanocryotron comparator can connect single-flux-quantum circuits to conventional electronics". In: *Superconductor Science and Technology* 30.4 (2017), p. 044002. ISSN: 0953-2048. DOI: 10.1088/1361-6668/aa5f33. URL: http://stacks.iop.org/0953-2048/30/i=4/a=044002 (visited on 03/08/2017).

[27] Andrew Murphy, Dmitri V. Averin, and Alexey Bezryadin. "Nanoscale superconducting memory based on the kinetic inductance of asymmetric nanowire loops". In: *New Journal of Physics* 19.6 (2017), p. 063015. ISSN: 1367-2630. DOI: 10.1088/1367-2630/aa7331. URL: http://stacks.iop.org/1367-2630/19/i=6/a=063015 (visited on 08/29/2018).

[28] Emily Toomey et al. "Frequency Pulling and Mixing of Relaxation Oscillations in Superconducting Nanowires". In: *Physical Review Applied* 9.6 (June 14, 2018), p. 064021. DOI: 10.1103/PhysRevApplied.9.064021. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.9.064021 (visited on 08/07/2018).

[29] Shi-Zeng Lin and Dibyendu Roy. "Role of kinetic inductance in transport properties of shunted superconducting nanowires". In: *Journal of Physics: Condensed Matter* 25.32 (2013), p. 325701. ISSN: 0953-8984. DOI: 10.1088/0953-8984/25/32/325701. URL: http://stacks.iop.org/0953-8984/25/i=32/a=325701 (visited on 11/16/2016).

[30] Qing-Yuan Zhao et al. "A compact superconducting nanowire memory element operated by nanowire cryotrons". In: *Superconductor Science and Technology* 31.3 (2018), p. 035009. ISSN: 0953-2048. DOI: 10.1088/1361-6668/aaa820. URL: http://stacks.iop.org/0953-2048/31/i=3/a=035009 (visited on 03/20/2018).

[31] Andrew J. Kerman et al. "Electrothermal feedback in superconducting nanowire single-photon detectors". In: *Physical Review B* 79.10 (Mar. 26, 2009). ISSN: 1098-0121, 1550-235X. DOI: 10.1103/PhysRevB.79.100509. arXiv: 0812.0290. URL: http://arxiv.org/abs/0812.0290 (visited on 09/07/2016).

[32] E. Toomey et al. "Bridging the Gap Between Nanowires and Josephson Junctions: A Superconducting Device Based on Controlled Fluxon Transfer". In: *Physical Review Applied* 11.3 (Mar. 4, 2019), p. 034006. DOI: 10.1103/PhysRevApplied.11.034006. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.11.034006 (visited on 04/12/2019).

[33] Adam N. McCaughan et al. "A kinetic-inductance-based superconducting memory element with shunting and sub-nanosecond write times". In: *Superconductor Science and Technology* (2018). ISSN: 1361-6668. DOI: 10.1088/1361-6668/aae50d. URL: http://iopscience.iop.org/10.1088/1361-6668/aae50d (visited on 10/18/2018).

[34] Brenden A. Butters. "Development of a scalable superconducting memory". Thesis. Massachusetts Institute of Technology, 2018. URL: https://dspace.mit.edu/handle/1721.1/120406 (visited on 03/10/2020).

[35] M. Tinkham et al. "Hysteretic I-V curves of superconducting nanowires". In: *Physical Review B* 68.13 (Oct. 17, 2003), p. 134515. DOI: 10.1103/PhysRevB.68.134515. URL: http://link.aps.org/doi/10.1103/PhysRevB.68.134515 (visited on 11/16/2016).

[36] H. L. Hortensius et al. "Critical-current reduction in thin superconducting wires due to current crowding". In: *Applied Physics Letters* 100.18 (Apr. 30, 2012), p. 182602. ISSN: 0003-6951. DOI: 10.1063/1.4711217. URL: http://aip.scitation.org/doi/abs/10.1063/1.4711217 (visited on 07/30/2017).

[37] Nikhil Kumar et al. "Controlling hysteresis in superconducting constrictions with a resistive shunt". In: *Superconductor Science and Technology* 28.7 (2015), p. 072003. ISSN: 0953-2048. DOI: 10.1088/0953-2048/28/7/072003. URL: http://stacks.iop.org/0953-2048/28/i=7/a=072003 (visited on 12/20/2016).

[38] Andrew E. Dane et al. "Bias sputtered NbN and superconducting nanowire devices". In: *Applied Physics Letters* 111.12 (Sept. 18, 2017), p. 122601. ISSN: 0003-6951. DOI: 10.1063/1.4990066. URL: http://aip.scitation.org/doi/abs/10.1063/1.4990066.

[39] O. Murat(Oguzhan Murat) Onen. "Construction of a superconducting circuit simulator and its applications in cryogenic computing". Thesis. Massachusetts Institute of Technology, 2019. URL: https://dspace.mit.edu/handle/1721.1/121740 (visited on 03/10/2020).

[40] Karl K. Berggren et al. "A Superconducting Nanowire can be Modeled by Using SPICE". In: *Superconductor Science and Technology* (2018). ISSN: 1361-6668. DOI: 10.1088/1361-6668/aab149. URL: http://iopscience.iop.org/10.1088/1361-6668/aab149 (visited on 03/14/2018).

[41] Murat Onen et al. "Design and characterization of superconducting nanowire-based processors for acceleration of deep neural network training". In: *Nanotechnology* 31.2 (Oct. 2019), p. 025204. ISSN: 0957-4484. DOI: 10.1088/1361-6528/ab47bc. URL: https://doi.org/10.1088%2F1361-6528%2Fab47bc (visited on 03/10/2020).

[42] Murat Onen et al. "Single-Photon Single-Flux Coupled Detectors". In: *Nano Letters* 20.1 (Jan. 8, 2020), pp. 664–668. ISSN: 1530-6984. DOI: 10.1021/acs.nanolett.9b04440. URL: https://doi.org/10.1021/acs.nanolett.9b04440 (visited on 03/10/2020).

[43] Emily Toomey, Ken Segall, and Karl K. Berggren. "Design of a Power Efficient Artificial Neuron Using Superconducting Nanowires". In: *Frontiers in Neuroscience* 13 (2019). ISSN: 1662-453X. DOI: 10.3389/fnins.2019.00933. URL: https://www.frontiersin.org/articles/10.3389/fnins.2019.00933/full (visited on 01/14/2020).

[44] Patrick Crotty, Dan Schult, and Ken Segall. "Josephson junction simulation of neurons". In: *Physical Review E* 82.1 (July 19, 2010), p. 011914. DOI: 10.1103/PhysRevE.82.011914. URL: https://link.aps.org/doi/10.1103/PhysRevE.82.011914 (visited on 06/15/2018).

[45] G. Bard Ermentrout and David H. Terman. "The Hodgkin–Huxley Equations". In: *Mathematical Foundations of Neuroscience*. Vol. 35. New York, NY: Springer New York, 2010, pp. 1–28. ISBN: 978-0-387-87707-5 978-0-387-87708-2. DOI: 10.1007/978-0-387-87708-2_1. URL: http://link.springer.com/10.1007/978-0-387-87708-2_1 (visited on 01/24/2019).

[46] E. M. Izhikevich. "Which model to use for cortical spiking neurons?" In: *IEEE Transactions on Neural Networks* 15.5 (Sept. 2004), pp. 1063–1070. ISSN: 1045-9227. DOI: 10.1109/TNN.2004.832719.

[47] Richard Bertram et al. "Topological and phenomenological classification of bursting oscillations". In: *Bulletin of Mathematical Biology* 57.3 (Jan. 1, 1995), pp. 413–439. ISSN: 0092-8240. DOI: 10.1016/S0092-8240(05)81776-8. URL: http://www.sciencedirect.com/science/article/pii/S0092824005817768 (visited on 02/19/2019).

[48] John Rinzel. "Bursting oscillations in an excitable membrane model". In: *Ordinary and Partial Differential Equations*. Ed. by Brian D. Sleeman and Richard J. Jarvis. Lecture Notes in Mathematics. Springer Berlin Heidelberg, 1985, pp. 304–316. ISBN: 978-3-540-39640-6.

[49] Takuya Nanami, Kazuyuki Aihara, and Takashi Kohno. "Elliptic and parabolic bursting in a digital silicon neuron model". In: 2016 International Symposium on Nonlinear Theory and Its Applications. Yugawara, Japan, 2016, p. 4.

[50] S. K. Dana, D. C. Sengupta, and C. Hu. "Spiking and Bursting in Josephson Junction". In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 53.10 (Oct. 2006), pp. 1031–1034. ISSN: 1549-7747. DOI: 10.1109/TCSII.2006.882183.

[51] C. Soto-Treviño, N. Kopell, and D. Watson. "Parabolic bursting revisited". In: *Journal of Mathematical Biology* 35.1 (Nov. 1, 1996), pp. 114–128. ISSN: 1432-1416. DOI: 10.1007/s002850050046. URL: https://doi.org/10.1007/s002850050046 (visited on 03/12/2019).

[52] F. Strumwasser. "The demonstration and manipulation of a circadian rhythm in a single neuron". In: *Circadian clocks*. Amsterdam: North-Holland Publishing Co., 1965, pp. 442–462.

[53] Barbara O. Alving. "Spontaneous Activity in Isolated Somata of Aplysia Pacemaker Neurons". In: *The Journal of General Physiology* 51.1 (Jan. 1, 1968), p. 29. DOI: 10.1085/jgp.51.1.29. URL: http://jgp.rupress.org/content/51/1/29.abstract.

[54] Steve Furber and Steve Temple. "Neural systems engineering". In: *Journal of the Royal Society Interface* 4.13 (Apr. 22, 2007), pp. 193–206. ISSN: 1742-5689. DOI: 10.1098/rsif.2006.0177. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2359843/ (visited on 04/17/2019).

[55] Qing-Yuan Zhao et al. "Single-photon imager based on a superconducting nanowire delay line". In: *Nature Photonics* 11.4 (Apr. 2017), pp. 247–251. ISSN: 1749-4893. DOI: 10.1038/nphoton.2017.35. URL: https://www.nature.com/articles/nphoton.2017.35 (visited on 04/10/2019).

[56] Stefano Ferraina, Martin Paré, and Robert H. Wurtz. "Comparison of Cortico-Cortical and Cortico-Collicular Signals for the Generation of Saccadic Eye Movements". In: *Journal of Neurophysiology* 87.2 (Feb. 1, 2002), pp. 845–858. ISSN: 0022-3077. DOI: 10.1152/jn.00317.2001. URL: https://doi.org/10.1152/jn.00317.2001 (visited on 06/06/2019).

[57] Bruce P. Bean. "The action potential in mammalian central neurons". In: *Nature Reviews Neuroscience* 8 (June 1, 2007), p. 451. URL: https://doi.org/10.1038/nrn2148.

[58] Anthony J. Annunziata et al. "Tunable superconducting nanoinductors". In: *Nanotechnology* 21.44 (Oct. 2010), p. 445202. ISSN: 0957-4484. DOI: 10.1088/0957-4484/21/44/445202. URL: https://doi.org/10.1088%2F0957-4484%2F21%2F44%2F445202 (visited on 04/10/2019).

[59] Paul A. Merolla et al. "A million spiking-neuron integrated circuit with a scalable communication network and interface". In: *Science* 345.6197 (Aug. 8, 2014), pp. 668–673. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1254642. URL: https://science.sciencemag.org/content/345/6197/668 (visited on 06/23/2019).

[60] Steve Furber. "Large-scale neuromorphic computing systems". In: *Journal of Neural Engineering* 13.5 (Aug. 2016), p. 051001. ISSN: 1741-2552. DOI: 10.1088/1741-2560/13/5/051001. URL: https://doi.org/10.1088%2F1741-2560%2F13%2F5%2F051001 (visited on 02/19/2019).

[61] Jeffrey M. Shainline et al. "Circuit designs for superconducting optoelectronic loop neurons". In: *Journal of Applied Physics* 124.15 (Oct. 12, 2018), p. 152130. ISSN: 0021-8979. DOI: 10.1063/1.5038031. URL: https://aip.scitation.org/doi/full/10.1063/1.5038031 (visited on 04/13/2019).

[62] Jeffrey M. Shainline. "Fluxonic Processing of Photonic Synapse Events". In: *arXiv:1904.02807 [cs]* (Apr. 4, 2019). arXiv: 1904.02807. URL: http://arxiv.org/abs/1904.02807 (visited on 04/13/2019).

[63] R. S. Person and L. P. Kudina. "Discharge frequency and discharge pattern of human motor units during voluntary contraction of muscle". In: *Electroencephalography and Clinical Neurophysiology* 32.5 (May 1972), pp. 471–483. ISSN: 0013-4694. DOI: 10.1016/0013-4694(72)90058-2.

[64] A. N. McCaughan, D. M. Oh, and S. W. Nam. "A Stochastic SPICE Model for Superconducting Nanowire Single Photon Detectors and Other Nanowire Devices". In: *IEEE Transactions on Applied Superconductivity* 29.5 (Aug. 2019), pp. 1–4. ISSN: 1051-8223. DOI: 10.1109/TASC.2019.2892049.

[65] M. Prezioso et al. "Training and operation of an integrated neuromorphic network based on metal-oxide memristors". In: *Nature* 521.7550 (May 2015), pp. 61–64. ISSN: 1476-4687. DOI: 10.1038/nature14441. URL: https://www.nature.com/articles/nature14441 (visited on 03/10/2020).

[66] Michael L. Schneider et al. "Energy-efficient single-flux-quantum based neuromorphic computing". In: *2017 IEEE International Conference on Rebooting Computing (ICRC)*. 2017 IEEE International Conference on Rebooting Computing (ICRC). ISSN: null. Nov. 2017, pp. 1–4. DOI: 10.1109/ICRC.2017.8123634.

[67] Michael Nielsen. "Neural Networks and Deep Learning". In: (), p. 224.

[68] Yanqing Chen. "Mechanisms of Winner-Take-All and Group Selection in Neuronal Spiking Networks". In: *Frontiers in Computational Neuroscience* 11 (2017). ISSN: 1662-5188. DOI: 10.3389/fncom.2017.00020. URL: https://www.frontiersin.org/articles/10.3389/fncom.2017.00020/full (visited on 02/20/2020).

[69] Nancy Lynch, Cameron Musco, and Merav Parter. "Winner-Take-All Computation in Spiking Neural Networks". In: *arXiv:1904.12591 [cond-mat, q-bio]* (Apr. 25, 2019). arXiv: 1904.12591. URL: http://arxiv.org/abs/1904.12591 (visited on 02/20/2020).

[70] Jonathan Binas et al. "Learning and stabilization of winner-take-all dynamics through interacting excitatory and inhibitory plasticity". In: *Frontiers in Computational Neuroscience* 8 (2014). ISSN: 1662-5188. DOI: 10.3389/fncom.2014.00068. URL: https://www.frontiersin.org/articles/10.3389/fncom.2014.00068/full (visited on 02/25/2020).

[71] Emily Toomey et al. "Influence of tetramethylammonium hydroxide on niobium nitride thin films". In: *Journal of Vacuum Science & Technology B* 36.6 (Oct. 15, 2018), 06JC01. ISSN: 2166-2746. DOI: 10.1116/1.5047427. URL: https://avs.scitation.org/doi/abs/10.1116/1.5047427 (visited on 12/19/2018).

[72] Emily Toomey, Marco Colangelo, and Karl K. Berggren. "Investigation of ma-N 2400 series photoresist as an electron-beam resist for superconducting nanoscale devices". In: *Journal of Vacuum Science & Technology B* EIPBN2019.1 (Sept. 25, 2019), p. 051207. ISSN: 2166-2746. DOI: 10.1116/1.5119516@jvb.2020.EIPBN2019.issue-1. URL: https://avs.scitation.org/doi/full/10.1116/1.5119516%40jvb.2020.EIPBN2019.issue-1 (visited on 03/10/2020).

[73] A. E. Grigorescu and C. W. Hagen. "Resists for sub-20-nm electron beam lithography with a focus on HSQ: state of the art". In: *Nanotechnology* 20.29 (2009), p. 292001. ISSN: 0957-4484. DOI: 10.1088/0957-4484/20/29/292001. URL: http://stacks.iop.org/0957-4484/20/i=29/a=292001 (visited on 04/25/2017).

[74] J. K. W. Yang et al. "Fabrication development for nanowire GHz-counting-rate single-photon detectors". In: *IEEE Transactions on Applied Superconductivity* 15.2 (June 2005), pp. 626–630. ISSN: 1051-8223. DOI: 10.1109/TASC.2005.849971.

[75] W. Henschel, Y. M. Georgiev, and H. Kurz. "Study of a high contrast process for hydrogen silsesquioxane as a negative tone electron beam resist". In: *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena* 21.5 (Sept. 1, 2003), pp. 2018–2025. ISSN: 1071-1023. DOI: 10.1116/1.1603284. URL: https://avs.scitation.org/doi/abs/10.1116/1.1603284 (visited on 04/04/2018).

[76] Joel K. W. Yang et al. "Understanding of hydrogen silsesquioxane electron resist for sub-5-nm-half-pitch lithography". In: *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena* 27.6 (Nov. 1, 2009), pp. 2622–2627. ISSN: 1071-1023. DOI: 10.1116/1.3253652. URL: https://avs.scitation.org/doi/abs/10.1116/1.3253652 (visited on 04/04/2018).

[77] Xiang-Lei Han, Guilhem Larrieu, and Emmanuel Dubois. "Realization of Vertical Silicon Nanowire Networks with an Ultra High Density Using a Top–Down Approach". In: *Journal of Nanoscience and Nanotechnology* 10.11 (Nov. 1, 2010), pp. 7423–7427. DOI: 10.1166/jnn.2010.2841.

[78] Khan Mohammad Moniruzzaman. *Development of a single-step fabrication process for nanoimprint stamps*. 2011. URL: http://urn.kb.se/resolve?urn=urn:nbn:se:hh:diva-14388 (visited on 04/04/2018).

[79] F. Najafi et al. "Fabrication Process Yielding Saturated Nanowire Single-Photon Detectors With 24-ps Jitter". In: *IEEE Journal of Selected Topics in Quantum Electronics* 21.2 (Mar. 2015), pp. 1–7. ISSN: 1077-260X. DOI: 10.1109/JSTQE.2014.2372054.

[80] Daniel Fan and Yasin Ekinci. "Photolithography reaches 6 nm half-pitch using EUV light". In: *Extreme Ultraviolet (EUV) Lithography VII*. Extreme Ultraviolet (EUV) Lithography VII. Vol. 9776. International Society for Optics and Photonics, Mar. 18, 2016, p. 97761V. DOI: 10.1117/12.2219737. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/9776/97761V/Photolithography-reaches-6-nm-half-pitch-using-EUV-light/10.1117/12.2219737.short (visited on 04/04/2018).

[81] Owen Medeiros et al. "Measuring thickness in thin NbN films for superconducting devices". In: *Journal of Vacuum Science & Technology A* 37.4 (May 21, 2019), p. 041501. ISSN: 0734-2101. DOI: 10.1116/1.5088061. URL: https://avs.scitation.org/doi/abs/10.1116/1.5088061 (visited on 08/29/2019).

[82] Werner Kern. "The Evolution of Silicon Wafer Cleaning Technology". In: *Journal of The Electrochemical Society* 137.6 (June 1, 1990), pp. 1887–1892. DOI: 10.1149/1.2086825. URL: http://jes.ecsdl.org/content/137/6/1887.abstract.

[83] Juan Bardina. "Methods For Surface Particle Removal: A Comparative Study". In: *Particulate Science and Technology* 6.2 (Jan. 1, 1988), pp. 121–131. ISSN: 0272-6351. DOI: 10.1080/02726358808906491. URL: https://doi.org/10.1080/02726358808906491.

[84] Ahmed A. BUsnaina and Fen Dai. "The Removal of Submicron Particles in Liquid-Based Cleaning". In: *The Journal of Adhesion* 67.1 (May 1, 1998), pp. 181–193. ISSN: 0021-8464. DOI: 10.1080/00218469808011107. URL: https://doi.org/10.1080/00218469808011107.

[85] Lauren B. Fullmer et al. "Nb2O5 and Ta2O5 Thin Films from Polyoxometalate Precursors: A Single Proton Makes a Difference". In: *Crystal Growth & Design* 15.8 (Aug. 5, 2015), pp. 3885–3892. ISSN: 1528-7483. DOI: 10.1021/acs.cgd.5b00508. URL: https://doi.org/10.1021/acs.cgd.5b00508 (visited on 04/10/2018).

[86] Mathias Van Thiel, Edwin D. Becker, and George C. Pimentel. "Infrared Studies of Hydrogen Bonding of Water by the Matrix Isolation Technique". In: *The Journal of Chemical Physics* 27.2 (Aug. 1, 1957), pp. 486–490. ISSN: 0021-9606. DOI: 10.1063/1.1743753. URL: https://doi.org/10.1063/1.1743753 (visited on 09/06/2018).

[87] Anton Litke et al. "Role of Adsorbed Water on Charge Carrier Dynamics in Photoexcited TiO2". In: *The Journal of Physical Chemistry C* 121.13 (Apr. 6, 2017), pp. 7514–7524. ISSN: 1932-7447. DOI: 10.1021/acs.jpcc.7b00472. URL: https://doi.org/10.1021/acs.jpcc.7b00472.

[88] Teruaki Fuchigami and Ken-ichi Kakimoto. "Spiky niobium oxide nanoparticles through hydrothermal synthesis". In: *Journal of Materials Research* 32.17 (Sept. 2017), pp. 3326–3332. ISSN: 0884-2914, 2044-5326. DOI: 10.1557/jmr.2017.200. URL: https://www.cambridge.org/core/journals/journal-of-materials-research/article/spiky-niobium-oxide-nanoparticles-through-hydrothermal-synthesis/730FF20E178E89969EB22C7CB047E9AC/core-reader (visited on 06/04/2018).

[89] Fabiane de Oliveira Cantão et al. "Utilization of Sn/Nb2O5 composite for the removal of methylene blue". In: *Química Nova* 33.3 (2010), pp. 528–531. ISSN: 0100-4042. DOI: 10.1590/S0100-40422010000300007. URL: http://www.scielo.br/scielo.php?script=sci_abstract&pid=S0100-40422010000300007&lng=en&nrm=iso&tlng=en (visited on 06/04/2018).

[90] Lester Andrews and George C. Pimentel. "Infrared Spectrum of the Methyl Radical in Solid Argon". In: *The Journal of Chemical Physics* 47.9 (Nov. 1, 1967), pp. 3637–3644. ISSN: 0021-9606. DOI: 10.1063/1.1712434. URL: https://doi.org/10.1063/1.1712434 (visited on 09/06/2018).

[91] B. Nolin and R. Norman Jones. "THE INFRARED ABSORPTION SPECTRA OF DEUTERATED ESTERS: I. METHYL ACETATE". In: *Canadian Journal of Chemistry* 34.10 (Oct. 1, 1956), pp. 1382–1391. ISSN: 0008-4042. DOI: 10.1139/v56-177. URL: https://doi.org/10.1139/v56-177 (visited on 09/06/2018).

[92] Robert T. Conley. "Typical Absorption Characteristics of Organo-Silicon Compounds". In: *Infrared Spectroscopy*. 1st ed. Boston, MA: Allyn and Bacon, Inc., 1966, pp. 176–178.

[93] P Gupta et al. "FTIR studies of H2O and D2O decomposition on porous silicon surfaces". In: *Surface Science* 245.3 (Apr. 2, 1991), pp. 360–372. ISSN: 0039-6028. DOI: 10.1016/0039-6028(91)90038-T. URL: http://www.sciencedirect.com/science/article/pii/003960289190038T (visited on 09/10/2018).

[94] F. Najafi et al. "Fabrication Process Yielding Saturated Nanowire Single-Photon Detectors With 24-ps Jitter". In: *IEEE Journal of Selected Topics in Quantum Electronics* 21.2 (Mar. 2015), pp. 1–7. ISSN: 1077-260X. DOI: 10.1109/JSTQE.2014.2372054.

[95] W. W. Stoffels, E. Stoffels, and K. Tachibana. "Polymerization of fluorocarbons in reactive ion etching plasmas". In: *Journal of Vacuum Science & Technology A* 16.1 (Jan. 1, 1998), pp. 87–95. ISSN: 0734-2101. DOI: 10.1116/1.581016. URL: https://doi.org/10.1116/1.581016 (visited on 09/06/2018).

[96] J H Thomas, X Chun Mu, and S J Fonash. "An X-Ray Photoelectron SpectroscopyStudy of CF4/H2 Reactive Ion Etching Residue on Silicon". In: 134.12 (), p. 4.

[97] Seong-Ju Park et al. "Reactive Ion Etching of Molybdenum In CF4/O2 Plasma". In: *MRS Proceedings* 68 (1986), p. 65. ISSN: 0272-9172. DOI: 10.1557/PROC-68-65. URL: https://www.cambridge.org/core/article/reactive-ion-etching-of-molybdenum-in-cf4o2-plasma/C64787B3BA11031784E7D3B54639B7C8.

[98] A. C. Seabra et al. "Dry Etching Of Niobium Oxyde Thin Films". In: *Dry Processing for Submicrometer Lithography*. Dry Processing for Submicrometer Lithography. Vol. 1185. International Society for Optics and Photonics, Jan. 30, 1990, pp. 80–86. DOI: 10.1117/12.978048. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/1185/0000/Dry-Etching-Of-Niobium-Oxyde-Thin-Films/10.1117/12.978048.short (visited on 05/09/2018).

[99] X. F. Meng et al. "NbN edge junction fabrication: edge profile control by reactive ion etching". In: *IEEE Transactions on Magnetics* 25.2 (Mar. 1989), pp. 1239–1242. ISSN: 0018-9464. DOI: 10.1109/20.92520.

[100] Joel K. W. Yang and Karl K. Berggren. "Using high-contrast salty development of hydrogen silsesquioxane for sub-10-nm half-pitch lithography". In: *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena* 25.6 (Nov. 1, 2007), pp. 2025–2029. ISSN: 1071-1023. DOI: 10.1116/1.2801881. URL: https://avs.scitation.org/doi/abs/10.1116/1.2801881 (visited on 04/04/2018).

[101] H. Elsner et al. "Evaluation of ma-N 2400 series DUV photoresist for electron beam exposure". In: *Microelectronic Engineering*. International Conference on Micro- and Nanofabrication 46.1 (May 1, 1999), pp. 389–392. ISSN: 0167-9317. DOI: 10.1016/S0167-9317(99)00113-6. URL: http://www.sciencedirect.com/science/article/pii/S0167931799001136 (visited on 01/17/2019).

[102] B. Bilenberg et al. "Comparison of high resolution negative electron beam resists". In: *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures* 24.4 (2006), p. 1776. ISSN: 10711023. DOI: 10.1116/1.2210002. URL: http://scitation.aip.org/content/avs/journal/jvstb/24/4/10.1116/1.2210002 (visited on 01/17/2019).

[103] M. Nuzaihan M. N et al. "Top-Down Nanofabrication and Characterization of 20 nm Silicon Nanowires for Biosensing Applications". In: *PLoS ONE* 11.3 (Mar. 29, 2016). ISSN: 1932-6203. DOI: 10.1371/journal.pone.0152318. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4811568/ (visited on 08/28/2019).

[104] Joel K. W. Yang, Vikas Anant, and Karl K. Berggren. "Enhancing etch resistance of hydrogen silsesquioxane via postdevelop electron curing". In: *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena* 24.6 (Nov. 1, 2006). Publisher: American Institute of Physics, pp. 3157–3161. ISSN: 1071-1023. DOI: 10.1116/1.2395949. URL: https://avs.scitation.org/doi/abs/10.1116/1.2395949 (visited on 03/16/2020).

[105] Advait Madhavan, Timothy Sherwood, and Dmitri Strukov. "Race Logic: A hardware acceleration for dynamic programming algorithms". In: *2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)*. 2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA). ISSN: 1063-6897. June 2014, pp. 517–528. DOI: 10.1109/ISCA.2014.6853226.

[106] Advait Madhavan, Timothy Sherwood, and Dmitri Strukov. "Race Logic: Abusing Hardware Race Conditions to Perform Useful Computation". In: *IEEE Micro* 35.3 (May 2015), pp. 48–57. ISSN: 1937-4143. DOI: 10.1109/MM.2015.43.

[107] B. R. Gaines. "Stochastic Computing Systems". In: *Advances in Information Systems Science: Volume 2*. Ed. by Julius T. Tou. Advances in Information Systems Science. Boston, MA: Springer US, 1969, pp. 37–172. ISBN: 978-1-4899-5841-9. DOI: 10.1007/978-1-4899-5841-9_2. URL: https://doi.org/10.1007/978-1-4899-5841-9_2 (visited on 03/10/2020).

[108] Armin Alaghi and John P. Hayes. "Survey of Stochastic Computing". In: *ACM Transactions on Embedded Computing Systems* 12.2 (May 1, 2013), 92:1–92:19. ISSN: 1539-9087. DOI: 10.1145/2465787.2465794. URL: https://doi.org/10.1145/2465787.2465794 (visited on 03/10/2020).

[109] J.S. Denker. "A review of adiabatic computing". In: *Proceedings of 1994 IEEE Symposium on Low Power Electronics*. Proceedings of 1994 IEEE Symposium on Low Power Electronics. ISSN: null. Oct. 1994, pp. 94–97. DOI: 10.1109/LPE.1994.573218.

[110] R. Barends et al. "Digitized adiabatic quantum computing with a superconducting circuit". In: *Nature* 534.7606 (June 2016), pp. 222–226. ISSN: 1476-4687. DOI: 10.1038/nature17658. URL: https://www.nature.com/articles/nature17658 (visited on 03/10/2020).

[111] Damir Vodenicarevic et al. “A Nanotechnology-Ready Computing Scheme based on a Weakly Coupled Oscillator Network”. In: *Scientific Reports* 7.1 (Mar. 21, 2017), pp. 1–13. ISSN: 2045-2322. DOI: 10.1038/srep44772. URL: https://www.nature.com/articles/srep44772 (visited on 03/10/2020).