Scaling of Nanocryotron Superconducting Logic

by

Reed A. Foster

S.B. Electrical Engineering and Computer Science
Massachusetts Institute of Technology, 2022

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

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2023

Authored by: Reed A. Foster
Department of Electrical Engineering and Computer Science
May 19, 2023

Certified by: Karl K. Berggren
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by: Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Scaling of Nanocryotron Superconducting Logic

by

Reed A. Foster

Submitted to the Department of Electrical Engineering and Computer Science
on May 19, 2023, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

ABSTRACT

This thesis presents the design and characterization of a superconducting shift register based on nanocryotrons. Such a shift register has applications in nanocryotron circuit testing as well as integrated readout and memory for high count rate imagers based on superconducting nanowire single photon detectors (SNSPDs). Characterization of the shift register shows that it can readily operate in large external magnetic fields that would present a challenge to Josephson-junction-based superconducting technologies. Furthermore, analysis of the input ranges which produce correct operation in a small experimental device suggest that such a circuit may be scalable to millions of nanocryotrons.

A device with a million nanocryotrons would be several orders of magnitude larger than any existing digital circuit based on superconducting nanowires. Development of circuits with more than just a few nanocryotrons has been limited in part due to the difficulty in testing and characterizing these superconducting devices. The absence of standard, well-tested nanocryotron circuits puts the burden of testing on conventional room-temperature electronics such as oscilloscopes and arbitrary waveform generators. However, limited flexibility of on-board computation for preprocessing data handicaps the ability of such systems to characterize larger scale circuits. To address this challenge, this thesis presents a design of an analog frontend for interfacing superconducting circuits with a high speed field-programmable gate array (FPGA) that could automate these tests.

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

Acknowledgments

This work would not have been possible without the assistance and support of my colleagues. I would especially like to thank:

My advisor, Professor Karl Berggren, for cultivating a wonderful environment for scientific research and his continuous engagement in and support of my work.

My mentor, Matteo Castellani, for welcoming me into the group and teaching me how to build and test superconducting nanowire devices. Without his help with fabricating my devices, over half this thesis would still just be an idea.

Torque Dandachi, from classes with Ian to the shared undergraduate office and beyond, for all of our discussions about electromagnetics, nanowire simulation, and FPGAs.

Alessandro Buzzi for our many discussions on digital logic and nanocryotrons, and for helping me to always pick the best dish on the menu in Italy.

Owen Medeiros for helping me whenever I had questions in the lab and coordinating the loan of test equipment that was crucial to this work.

Marco Colangelo for helping answer all my questions about academia, always providing insightful feedback on my writing, and setting an example of what it means to be an excellent researcher.

Dorothy Fleischer and Rinske Wijtmans Robinson for helping all of us at QNN stay organized, coordinating travel, and navigating MIT's purchasing systems.

John Simonaitis for our runs together and always rallying the troops for lunch.

Stephen Kandeh for his intense passion for FPGAs and for giving me the opportunity to practice mentorship. I always find I learn more by explaining things to others.

The rest of the members of the QNN group, for daily discussions in the hallway, group lunches, and all of the knowledge they have shared with me.

I would also like to acknowledge the major funding support for my research from DARPA and the DOE.

I would also like to thank my friends and family, without whom I would not be

the person I am today.

From 411 to Colorado and everywhere in between, Aiden Foucault, Izzi Gengaro, Tim Magoun, Abby Mauermann, Nicholas Medearis, and Ryan Wilson for the many memories over the years.

Jordan Sell for her continuous support and encouragement, especially when things weren't going well.

Mike and Charmaine Hartnett for keeping me well-fed and coming out to watch me race.

Most of all, my parents for always supporting my educational, athletic, and artistic endeavours. I'm thankful to have such thoughtful and caring role models, and I am so incredibly lucky that I can share my excitement in the nitty-gritty details of my research with you.

Contents

List of Figures

Acronyms

ADC analog-to-digital converter. 39, 40, 107, 109–111, 114, 123, 130, 132, 134, 136, 137, 139, 141–143

AWG arbitrary waveform generator. 32, 59, 73, 75, 77, 78, 81, 101

BER bit error rate. 72, 73, 75, 77, 78, 80, 81, 94

BJT bipolar junction transistor. 114

CCD charge coupled device. 45, 63

CMFB common-mode feedback. 110, 111

CMOS complementary metal-oxide-semiconductor. 31, 35, 36

CMRR common-mode rejection ratio. 141

DAC digital-to-analog converter. 104, 109, 111, 123, 126, 130, 134, 141–143

DAQ data acquisition/digitizer. 75, 77, 78

DDS direct digital synthesis. 42, 126–128, 130, 133, 134

DSO digital storage oscilloscope. 32, 73, 75, 78, 94, 101, 102

DUT device under test. 145, 146

ENIG electroless nickel immersion gold. 112

ENOB effective number of bits. 39

FDA fully-differential amplifier. 108–111, 142

FPGA field-programmable gate array. 41, 103, 126, 127, 130, 133, 134, 145, 146

FWHM full width at half-maximum. 75

HDI high density interconnect. 113

JJ Josephson junction. 35–37, 45, 93

LNA low-noise amplifier. 40, 59, 78, 105–107, 130, 134

MOSFET metal-oxide-semiconductor field effect transistor. 114, 115

NMP n-methyl-2-pyrrolidone. 70

PCB printed circuit board. 73, 104, 108, 111, 113, 114

PLL phase-locked loop. 133

PRBS pseudorandom binary sequence. 75, 78

PSRR power supply ripple rejection. 40, 114

PWM pulse width modulation. 117, 118

RFADC radio-frequency analog-to-digital converter. 103, 104, 108, 109, 124, 126, 132, 133, 137, 146

RFDAC radio-frequency digital-to-analog converter. 103, 109, 110, 128, 132–134, 137, 142, 143

RFDC radio-frequency dataconverter. 133

rms root-mean-square. 41, 124, 139, 140

SFDR spurious-free dynamic range. 42, 126–130, 134, 136, 137, 139, 140

SFQ single flux quantum. 31, 36, 37, 45

SINAD signal to noise and distortion ratio. 134, 136, 137, 139

SNR signal-to-noise ratio. 39, 42, 116, 126, 129, 130, 134

SNSPD superconducting nanowire single photon detector. 31, 36, 45, 59, 61, 63, 65, 93

SPI serial peripheral interface. 108

VGA variable-gain amplifier. 108, 109, 134, 137, 142

VNA vector network analyzer. 126, 141, 143

XOR exclusive or. 78

Glossary

FR-4 inexpensive fiberglass resin laminate used in the construction of printed circuit boards. 114

GE varnish low temperature cryogenic glue compatible with vacuum systems. 73

JTAG standard for verifying digital electronics devices with a serial debug port, named after the Joint Test Action Group which codified it. 45

LTSpice closed-source, free, circuit simulator from Linear Technology with emphasis on fast simulation. 46, 53, 55, 59

NbN niobium nitride, popular for thin-film superconductors due to its high critical temperature T_c and ease of fabrication. 67, 68, 70, 88, 90

opamp operational amplifier. 109–111

periodogram estimate of the power spectral density of a signal. 136

RFSoC series of field-programmable gate arrays manufactured by Xilinx with integrated gigasample dataconverters and multicore ARM processors. 102–104, 109, 114, 126, 130, 132–134, 136, 141, 145, 146

SerDes serializer-deserializer, a mixed signal block which facilitates the transfer of digital data between separate digital systems, for chip-to-chip, package-to-package, and board-to-board links. 45

Chapter 1

Introduction

Superconducting nanowires are an interesting candidate for low power computing [1, 2, 3] and interfacing with superconducting nanowire single photon detectors (SNSPDs) [2, 3, 4, 5]. While progress has been made in developing standards for nanocryotron logic [1], the ability to test complex circuits based on superconducting nanowire presents a challenge towards developing large scale circuits of superconducting nanowires.

Fast circuit simulation models of nanowires [6, 7, 8] have enabled rapid prototyping of electronics based on superconducting nanowires. Many of these devices are simple to fabricate, requiring only a single lithography step to pattern a single superconducting film. However, testing these fabricated devices still presents a challenge. Well-established electronics platforms, such as complementary metal-oxide-semiconductor (CMOS) and even single flux quantum (SFQ) superconducting logic, have standardized signaling schemes, logic levels, and basic circuits which are well-studied and can be expected to perform in a particular way. These basic circuits can be used to construct built-in self-test circuitry that surrounds experimental circuits to study the performance of the experimental circuit. Superconducting nanowire and nanocryotron circuits are nascent enough that these self-test circuits do not exist, and as a result, testing must be carried out by conventional electronics test equipment.

Energy-efficient operation of superconducting nanowires uses spiking behavior as opposed to latching, which poses strict requirements on the data acquisition systems used to measure the behavior of nanowire electronics. When optimized for speed,

nanowire electronics can generate spikes as short as a few nanoseconds. In order to detect such short voltage spikes, test equipment must have bandwidth at or above 1 GHz and sampling rates in the multi-GS/s range. Furthermore, it's important that the test equipment can capture the characteristic shape of the nanowire switching event to distinguish it from other events, which means that the analog voltage must be digitized to a multi-bit value. This poses a challenge for testing large scale circuits, where it may be quite helpful to observe many nanowires simultaneously, since the combined data rate the acquisition system must handle may be hundreds of gigabits per second or even multiple terabits per second.

In addition to measuring the transient response of a circuit made of superconducting nanowires, the test equipment has to provide a stimulus. This poses a challenge for testing nanowire circuits at their maximum operating frequency, since the width of stimulus pulses must be at most a few nanoseconds. Long current pulses into unshunted nanowires create a significant amount of joule-heating in the nanowire, potentially limiting the speed of a circuit that might otherwise be able to run faster if shorter pulses were used. These requirements dictate the use of high-bandwidth, high-sample-rate signal sources. Furthermore, the ability to generate randomized test patterns and arbitrary waveforms is important for fully characterizing circuits made of nanocryotrons.

Existing strategies for testing superconducting nanowire electronics utilize arbitrary waveform generators (AWGs) and digital storage oscilloscopes (DSOs) to perform stimulus generation and data acquisition. This approach works fine for probing just a few nanowires to check to see if a circuit is operating correctly, but characterization of circuits (*e.g.* calculating a bit error rate, understanding the various modes of operation for different input signal levels) is tedious and time-consuming. Limited compute ability and flexibility of DSOs require post-processing to extract useful information from the raw voltage waveforms captured by the DSOs. This cripples the potential throughput of a test setup because it involves the use of non-volatile storage and networking to save the data on the DSO and transfer it to another computer to process. Even if the DSO used the fastest solid-state storage, testing electronics

with it would be tens or even hundreds of times slower than if the processing of the waveforms could be performed in real time.

The goal of this thesis is to demonstrate a small circuit with a few nanocryotrons and to propose a data acquisition system for testing large scale nanocryotron circuits. The thesis is split into two parts. In the first part, the design and fabrication of a shift register built with nanocryotrons is discussed. In the second part, the architecture of a data acquisition system is proposed and the design and characterization of a low-noise analog frontend is discussed.

Chapter 2

Background

This chapter introduces superconducting electronics, with emphasis on superconducting nanowire and nanocryotron-based electronics. Testing of superconducting nanowires with room temperature test equipment is also discussed. A somewhat comprehensive review of relevant power-electronics background information is also presented in the context of designing low-noise power supplies for analog frontends. Design techniques for high bandwidth, low-noise analog electronics are also introduced. The chapter wraps up with a discussion of digital signal generation techniques for testing analog frontends.

2.1 Superconducting electronics

Superconductors are a class of materials, which when cooled down below a critical temperature T_c exhibit vanishing resistance and expulsion of magnetic fields [9]. Computation with electronics that exhibit zero resistance is attractive, due to the very low power consumption compared to conventional CMOS electronics.

2.1.1 Josephson junctions and single-flux-quantum logic

A Josephson junction (JJ) is formed by separating two superconducting films with a very thin oxide. When the oxide is thinner than the coherence length ξ of the

superconducting order parameter ψ, the supercurrent is able to tunnel through the oxide barrier [9]. The oxide film thickness dictates the critical current of the junction, which can be hundreds of times lower than the critical current of the superconducting wires on either side of the junction, even for high-critical-current JJs [10].

A variety of low-power logic families based on JJs exist. Of note is SFQ-based logic, which as the name suggests, represents states with single quanta of magnetic field excitations. This allows for extremely low power consumption (*e.g.* sub-attojoule dissipation per switching event) at the cost of high sensitivity to external magnetic fields. In addition to requiring extensive shielding to operate, the low critical current and normal resistance of JJs makes interfacing with conventional CMOS electronics difficult.

2.1.2 Superconducting nanowires and nanocryotrons

Superconducting nanowire-based electronics were born out of the SNSPD, a photodetector with near-unity detection efficiency, high count-rate, low jitter, and low dark-count rate that is sensitive to photons ranging from UV to mid-IR [11]. The superconducting state in a nanowire can be destroyed by exceeding the critical temperature, magnetic field, or current of the nanowire. When the superconducting state is destroyed, the current is carried by normal electrons and the nanowire becomes resistive. This hotspot can reach $100\,\Omega$ to $10\,\mathrm{k}\Omega$ in resistance, depending on the nanowire geometry and its impedance environment, meaning superconducting nanowires can readily be interfaced with conventional CMOS electronics. Several superconducting nanowire devices which utilize this behavior have been demonstrated [12, 13, 14]. This work mostly focuses on the nanocryotron (nTron) shown in Figure 2-1. The nTron is a three-terminal electrothermal switch that uses joule heating in a small constriction to suppress the critical current of a wider, superconducting channel [12]. A different variety of nTron with a much wider gate that is comparable in width to the channel is introduced in Chapter 3.

Highly disordered thin films tend to exhibit larger kinetic inductance, due to the stronger coupling between lattice phonons and the cooper pairs [15]. The kinetic

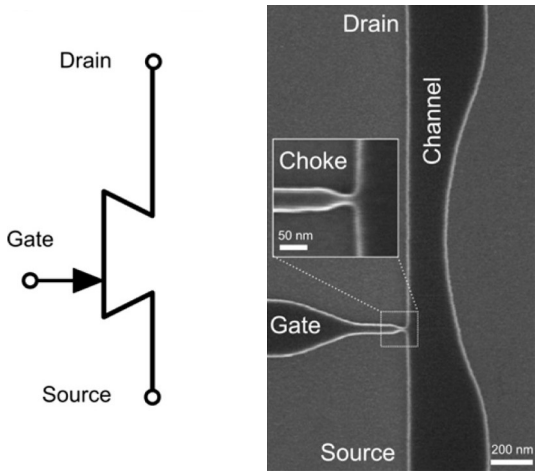

(a) (b)

Figure 2-1: Circuit schematic symbol and scanning electron micrograph of superconducting nanocryotron. Figure from [12].

inductance $L_{k,\square}$ of a superconducting film can be estimated from the measured room temperature resistance $R_{n,\square}$ and critical temperature T_c based on Equation 2.1 from [6]:

$$L_{k,\square} \approx \frac{\hbar}{\pi\Delta} R_{n,\square} \approx (1.38\,\mathrm{pH\,K/\Omega})\,\frac{R_{n,\square}}{T_c} \tag{2.1}$$

The high kinetic inductance allows for the use of nanowires in high impedance environments without latching, allowing for high output signal amplitudes.

As discussed in section 2.1.1, the low energy dissipation of SFQ logic, while being its greatest strength, is also one of its greatest weaknesses, due to the sensitivity of the circuit state to external magnetic fields. Operating in a single flux quantum regime entails dissipating a very small amount of energy per operation, but it also means a state can be flipped by the introduction of a single fluxoid. Nanowire logic enables operation in the many fluxoid regime, which reduces unwanted sensitivity to external magnetic fields. It should also be noted that superconducting nanowires can also operate in a heavily-shunted, vortex crossing regime, which is expected to have similar power dissipation to JJ-based electronics [16, 17]. However, this operating regime is expected to have the same issues as JJs with regard to magnetic field

sensitivity.

Nanowires are most energy efficient when operating in a spiking (instead of a latching) regime. In the latching regime, the bias current dissipates power continuously as long as the nanowire is in the resistive state. However, in the spiking regime, the nanowire is shunted with a low (or zero) real impedance, such as a resistor, or a resistor in series with an inductor. The shunt allows the bias current flowing through the nanowire to be diverted when the nanowire switches into the resistive state, allowing it to retrap to a superconductive state very shortly after it initially became resistive. The time constant of the nanowire and shunt inductance, as well as the DC resistance of the shunt determines how quickly the bias current returns to the nanowire after it retraps into the superconducting state.

2.1.3 Meissner effect in superconductors

In 1933, Meissner and Ochsenfeld observed that superconductors actively expel magnetic flux when cooled below their critical temperature [9]. The generation of screening currents can impact the performance of nTron circuits under the application of strong (*e.g.* 1 mT) magnetic fields by reducing or increasing the input current required to switch the nTron. This effect is important to studying the error rate, input margins, and scalability of nTron circuits in strong magnetic fields.

2.2 Testing superconducting nanowire electronics

In order to test superconducting nanowires, they must be connected to test equipment that can generate short pulses of current and measure the fast switching events of the nanowires. Ideally the current pulses are shorter than 1 ns, since longer pulses deposit more heat in the nanowire and limit thermal reset times. Fortunately, due to the small cross-sectional area of superconducting nanowires, the current required to switch the nanowire is relatively low, allowing for the use of relatively low output-power (*e.g.* 10 dBm) signal sources, so the main constraints for the signal generation are bandwidth and programmability. When a nanowire switches from superconductive

to resistive, it generates a voltage spike with a very fast (*e.g.* 10 ps) rising edge. The decay of the voltage spike as the nanowire returns to the superconductive state is longer, but can be less than 1 ns for nanowire electronics that are designed to operate in a thermally-limited regime [18]. In the thermally-limited regime, the switching rate of the nanowire is on the order of the thermal time constant required to retrap to the superconducting state, which is typically hundreds of picoseconds for thin-film superconductors on silicon substrates. For example, if one wanted to test a circuit with just ten nanocryotrons, the test equipment would have to have ten input channels, each with an analog bandwidth of several gigahertz, and be able to handle a combined datarate across all channels of nearly a terabit per second.

2.3 Low noise power supplies for analog frontends

For sensitive analog applications where measurement of an input signal with minimal degradation of signal-to-noise ratio (SNR) is desired, it is crucial that a minimal amount of noise is introduced by the power supply for the active electronics. For applications in which the analog input is digitized by an analog-to-digital converter (ADC), the noise introduced by any analog frontend electronics should be well below the input-referred noise of the ADC. The two sources of noise that limit the SNR of an ADC are real physical noise from the analog circuitry and quantization noise from the ADC's finite resolution. The quantization noise can be decreased by roughly 6 dB for each additional bit of precision. However, additional bits beyond the effective number of bits (ENOB) dictated by the analog noise and nonlinearity of the ADC are meaningless, so most ADCs will carefully design the analog circuitry so that the precision is just slightly higher than the ENOB. The noise floor set by the ADC ENOB gives us a target for minimal power supply noise. As long as the power supply noise injected on the ADC input is below the noise floor of the ADC, it will not impact the measurement. This is readily achievable with high speed ADCs which typically have ENOBs in the range of 10 bit to 14 bit.

2.3.1 Power supply ripple rejection (PSRR)

Power supply ripple rejection (PSRR) is a measure of how strongly an active device rejects noise conducted on the power supply and is defined as the ratio:

$$PSRR = \frac{(v_{\text{out}}/v_{\text{in}})_{v_{\pm}=0}}{(v_{\text{out}}/v_{\pm})_{v_{\text{in}}=0}} \tag{2.2}$$

where v_{in} and v_{out} are the small-signal input and output voltage, and $v_{\pm}$ is the small-signal voltage noise on the positive or negative supply rail.

Ideally, an active device would have no leakage of power supply noise to the output, making the PSRR infinite. In reality, channel-length modulation limits the low frequency PSRR of active circuits, and gain rolloff hurts PSRR at high frequencies. The PSRR of a multi-stage amplifier chain is dominated by the PSRR of the bias network on the input stage of the first amplifier, following a similar expression to the Friis equation for noise figure of cascaded amplifiers.

When designing a low-noise power supply for an amplifier chain, knowing the PSRR of an amplifier chain is helpful because it will enable the designer to quantify how much power supply output noise is acceptable. The PSRR should not be less than unity (except perhaps at extremely high frequencies, but any noise at these frequencies can easily be filtered with a bypass capacitor), so in the absence of information on component PSRR, unity PSRR can be used to inform the design of the power supply.

2.4 Receiver noise power and minimum detectable signal levels

Digitization of very small signals, such as the voltage spikes generated by a switching nanowire, requires amplification to boost the signal amplitude to a level where it is detectable by the ADC. This amplification process inherently introduces noise, although low-noise amplifiers (LNAs) are able to amplify radio-frequency signals while adding very little noise. With an amplification chain before the ADC, the lower limit

on signal detection is dominated by the input-referred noise of the amplifier. The total noise power referred to the input of the receiver P is given by

$$P = Fk_BTB \tag{2.3}$$

where F is the noise factor of the receiver, k_B is the Boltzmann constant, T is the temperature of the receiver, and B is the brick-wall bandwidth of the noise (which is slightly larger than the 3 dB bandwidth) [19].

For an amplifier at 300 K with 1 GHz brick-wall bandwidth and 3 dB noise figure, the receiver noise power is $P = -81$ dBm. The equivalent root-mean-square (rms) noise voltage across the 50 Ω input of the amplifier is 20 µV. Even if the noise figure of the amplifier were 0 dB (an ideal amplifier that introduces no noise), the noise voltage is 14 µV. This presents a challenge for measuring very small signals that have very high bandwidth. Fortunately, if the rough shape of the signal is known, strategies such as matched filtering can be used to detect signals well below the noise floor of the receiver as long as the signal power spectral density is larger than the noise spectral density [20].

2.5 Methods for digital signal generation

There are a variety of techniques that can be used to characterize the performance of the analog frontend. The most straightforward techniques are to use a signal generator and spectrum analyzer or a vector network analyzer. The results of characterization with a network analyzer will be discussed later in Chapter 7. For this application, it also makes sense to characterize the frontend with the field-programmable gate array (FPGA) that it was designed for. To do this, the FPGA must generate stimulus and measure the response of the analog frontend. By generating a pure sinusoidal tone and sweeping the frequency and power of that tone, the noise, gain, and linearity of the analog frontend can be characterized over a variety of single tone inputs. As long as the response of the amplifier is mostly linear (which is never the case, but for small

input powers is a valid approximation), the response to more complex waveforms can be inferred from a linear combination of the measured responses to individual frequencies that comprise spectrum of the complex waveform.

2.5.1 Direct digital synthesis (DDS)

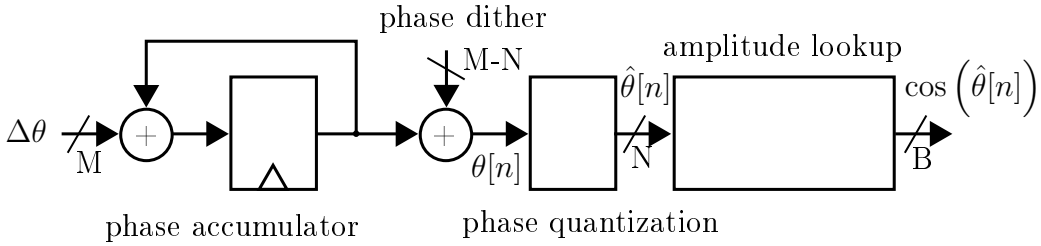

Figure 2-2: Schematic of a phase-dithering DDS. Figure adapted from [21].

Direct digital synthesis (DDS) is a technique to digitally generate sinusoidal waveforms (either as a single sinusoid or a pair of quadrature phase waveforms) [21]. Due to its digital implementation, DDS offers the ability to perform instantaneous frequency control, which is attractive in digital communication systems [21], as well as coherent imaging systems such as FMCW RADAR [22] and LiDAR. In this work, DDS is useful because it allows for the generation of very pure (*e.g.* high spurious-free dynamic range (SFDR) and SNR) sinusoidal tones which can be used for studying the noise and linearity performance of the analog frontend used for interfacing with superconducting circuits. Figure 2-2 shows a block diagram of a DDS system that employs phase dithering to reduce phase-quantization noise. An accumulator is used to track the phase of the waveform, which is quantized to index into a small precomputed lookup table that performs phase-to-amplitude conversion.

Part I

A superconducting nanowire binary shift register

Chapter 3

Design and simulation of a nanocryotron shift register

Shift registers serve a variety of functions in modern digital systems. The most common use case of a shift register is serialization and deserialization of data, used for moving data between separate digital systems, both at slow rates, such as in a 9.6 kbit/s UART serial link, and at high speeds, such as in a 56 Gbit/s SerDes transceiver [23] (although the impressive part of high speed SerDes design is most certainly not the shift register, but the analog frontend and signal conditioning which goes into the serialization and deserialization process). SerDes can also be used for the construction of debug interfaces to digital systems, such as JTAG. In the case of a superconducting shift register, the application of detector readout is also of interest. Due to the shared technology platform with SNSPDs, direct readout of large SNSPD arrays with co-integrated nanocryotron logic is quite attractive. A superconducting shift register could be used to read out an array of SNSPDs in a manner similar to that of a charge coupled device (CCD) [24]. Yet another use case of superconducting shift registers is to develop large scale superconducting nanowire systems, both by characterizing process, and by constructing serial-interface debug ports to test experimental superconducting circuits. Indeed, SFQ-based shift registers have been used to benchmark JJ processes [25] and verify digital multiplier circuits [26].

Before designing a shift register with superconducting nanowires, it is important to take a step back and look at how to leverage the strengths of superconducting nanowires in digital logic. This chapter will first discuss shift registers from a high-level or behavioral perspective, discuss how to implement them with nTrons, and discuss modeling of nTrons in LTSpice and present simulation results of a design for a nTron-based superconducting shift register.

3.1 Operating principle of the shift register

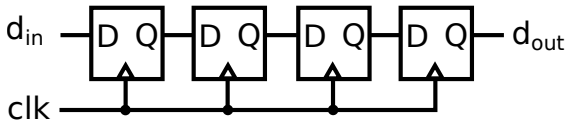

Figure 3-1: Four-stage synchronous digital shift register constructed with D flip-flops. The state in each D flip-flop is transferred to the subsequent D flip-flop on the rising edge of the clock signal.

The fundamental property of a shift register is that it transfers digital states between adjacent stages upon some stimulus. For a synchronous (clocked) shift register, this stimulus is often a clock signal. A schematic of a conventional synchronous digital shift register constructed with D flip-flops is shown in Figure 3-1. The superconducting nanowire shift register achieves the same behavior using superconducting loops instead of D flip-flops. Similar to the superconducting memory in [27], the shift register encodes states with circulating supercurrents. Each D flip-flop is replaced with two superconducting loops which can store supercurrents, and the single-phase clock signal `clk` is replaced with a two-phase clock. The loops are connected with nTrons: the channel of each nTron is shared between adjacent loops. Upon application of the clock signal, digital states can be transferred between these loops. The reason for using two loops and a two-phase clock instead of a single loop and single-phase clock will be described in more detail in section 3.1.2, but first a discussion of inductively-shunted nTrons is warranted.

3.1.1 Superconducting memory building blocks: inductively-shunted wide-gate nTrons

The nTron serves as the key building block of a superconducting shift register that is easy to fabricate and robust to external magnetic fields. In order to make a circuit out of nTrons that performs the function of a shift register, we need to consider how we will encode and store state information, and what process will be used to manipulate those states. The first important property to note is that nTrons and superconducting nanowires are well-suited for operating with currents, and due to the persistent nature of currents in a superconductor, non-volatile storage of currents. Superconducting nanowire logic can use the presence, absence, or direction of current in a loop to encode digital states. States can be manipulated by using the normal-superconducting transition of superconducting nanowires to dissipate supercurrents or divert the direction of current flow from a previously superconducting path into a different superconducting path. These two phenomena are provide sufficient functionality to build logic gates [1], digital counters [2], and shift registers.

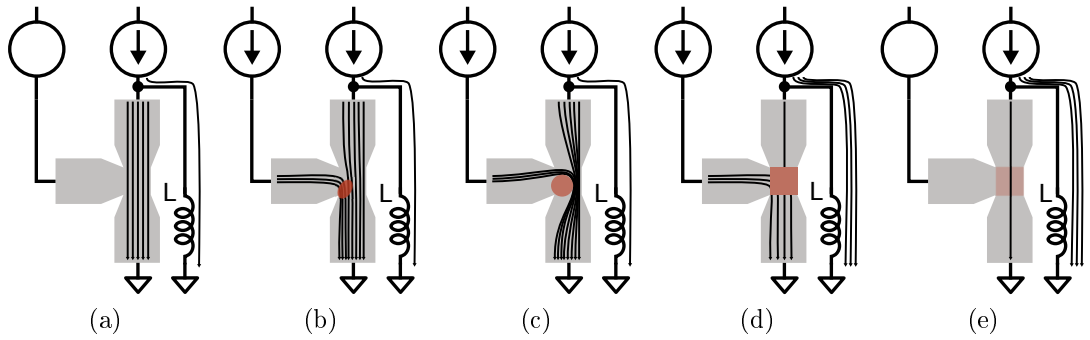

Figure 3-2: Storage of supercurrent in an inductive load. The circuit is initialized in a with a supercurrent split between the inductive load and the wide-gate nTron. In b, a current is supplied to the gate of the nTron, causing a hotspot to form in the channel of the nTron. The hotspot grows as shown in c, and eventually causes the bias current to almost entirely be diverted into the superconducting shunt inductor (d). The hotspot heals, and the gate current is removed, leaving a supercurrent stored in the inductor in e.

The inductively-shunted nTron circuit serves as a useful example for how a digital state encoded by a supercurrent can be manipulated. The desired behavior of the

inductively-shunted nTron circuit is to store a current in the superconducting inductor if and only if a gate current is present. If we assign logical values to the presence or absence of gate current and the presence or absence of current through the inductive load, then the logical or behavioral function of the circuit is to store a logical "1" conditional on the presence of a logical "1" input. This behavioral description is similar to an SR (set-reset) latch without the reset functionality.

The circuit is initialized with a current bias in Figure 3-2a. The bias current predominantly flows through the nTron due to its relatively low kinetic inductance compared to the load ($L_{\mathrm{nt}} \ll L$). When a gate current is supplied, the combined gate and drain current exceeds the local critical current density in the channel of the nTron, causing a breakdown of superconductivity, or hotspot, as shown in 3-2b. The hotspot continues to grow in 3-2c, until it covers the entire channel of the nTron and diverts the majority of the bias current into the inductive load as shown in 3-2d. Eventually, the hotspot begins to heal, and the current bias is permanently trapped in the inductive load as shown in 3-2e. On its own this configuration is a write-only-memory (*i.e.* completely useless from a practical perspective), since once a current is stored in the inductor, there is no way to move it back to the nTron. However, if the inductor is connected to the gate of a subsequent inductively-shunted nTron, then the gate current of the subsequent nTron is just the state of the first loop. Replacing the constant drain bias with a toggling clock signal thus enables the state to be transferred between loops. This idea of cascading inductively shunted nTrons forms the basis of the operating principle of the superconducting shift register, and allows synchronous transfer of digital states between superconducting loops.

3.1.2 nTron-connected flux loops and two-phase clocking

The operating principle of the superconducting shift register is illustrated in Figure 3-3. As described in section 3.1.1, the shift register is constructed by cascading inductively-shunted nTrons. This forms a series of superconducting loops consisting of superconducting kinetic inductors between nTrons. The shift register is initialized with a single logical "1" stored in the loop kinetic inductor L_{k} between nTrons U_1

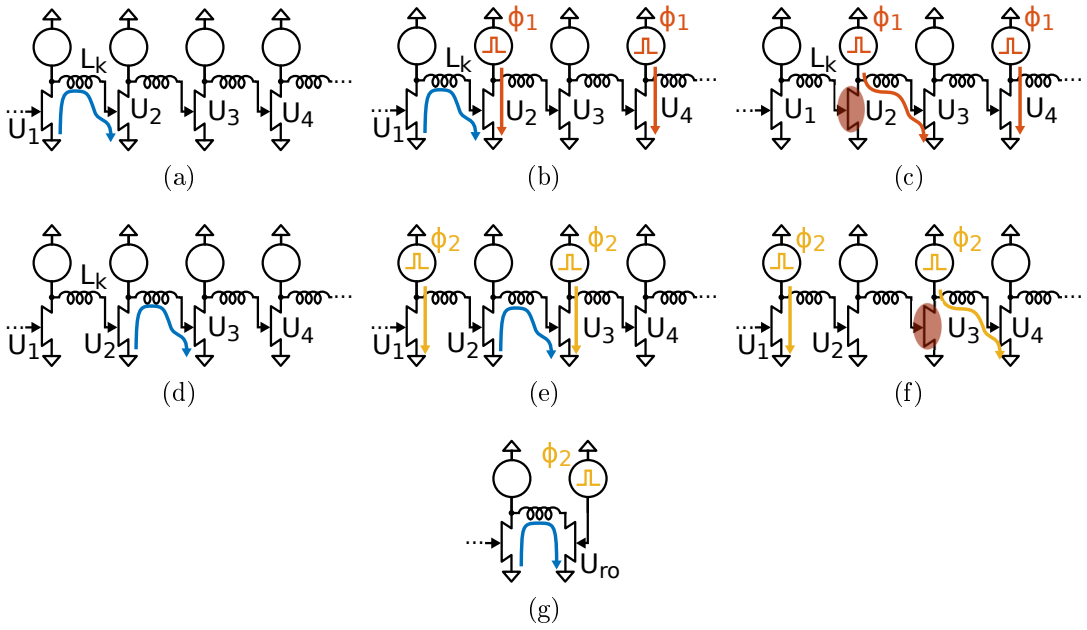

Figure 3-3: Operating principle of superconducting nanowire binary shift register. The shift register is initialized with a circulating current in the first loop formed by the kinetic inductor L_k between nTrons U_1 and U_2 as shown in a. The first clock phase ϕ_1 is applied in b, and in c, the combination of the initial circulating current and the clock signal exceed the critical current of the channel of U_2, causing it to switch and diverting the clock through the inductor between U_2 and U_3. After removing the clock, there is a circulating current left in the loop between U_2 and U_3, as shown in d. The process continues in e and f, where the second clock phase ϕ_2 is applied, causing U_3 to switch, diverting current through the inductor between U_3 and U_4. In g, a narrow-gate nTron is used to destructively read out the circulating current in the final loop of the shift register.

and U_2 as shown in 3-3a. When the first phase of the clock ϕ_1 is applied in 3-3b, the clock current combines with the circulating current in the channel of U_2, exceeding its critical current and switching the nTron into the resistive state. This diverts the clock current into the next loop as shown in Figure 3-3c. When the first phase of the clock is removed, a circulating current is left in the loop between U_2 and U_3 as shown in 3-3d. The process continues in Figure 3-3e and 3-3f, when the second clock phase ϕ_2 is applied and transfers the state yet again. An important observation to make is that every loop in the middle of the shift register resets at the end of each clock cycle. This reset is crucial for the operation of the shift register. For example, the circulating current in Figure 3-3d in the loop between U_2 and U_3 must be reset

before the next application of ϕ_1, otherwise U_2 will not switch even when a current is present in the previous loop (between U_1 and U_2). The shifting operation by ϕ_2 resets the circulating current created by the application of ϕ_1 (and vice versa). This is why the shift register uses two superconducting loops for each D flip-flop in Figure 3-1. Figure 3-3g shows the final, readout loop of the shift register, which is necessary for the exact same reason: to reset the circulating current in the final stage before the next clock phase is applied. Without it, the circulating current in the final stage would prevent the first nTron in the final loop from resetting. The readout nTron is designed as a conventional, narrow-gate nTron, since the clock is applied to the gate instead of the drain. Because it terminates the final shift register loop, there is no subsequent stage for the clock to be diverted to, so the amplitude of the clock signal does not matter. The narrow gate provides marginal power savings, since the required current to switch the readout nTron from the control signal is lower. However, the main reason for using a smaller gate is that it reduces the leakage current from the gate backwards into the final loop. This leakage behavior was predicted in simulation and observed in experiment and is discussed in section 4.4.

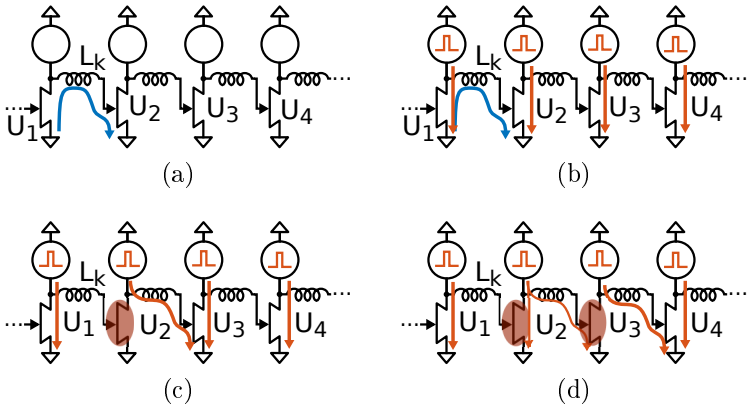

Figure 3-4: Illustration of incorrect operation when using a single phase clock. The shift register is initialized with a circulating current in the first loop. All clock signals are applied in b, which causes the nTron U_2 to switch as shown in c. In d, this switching event causes chain reaction of switching events and wipes out any logical "0"s in the shift register.

The superconducting shift register requires a two phase clock in order to ensure that the clock signal transfers a state only one loop at a time. The case of a single-

phase clock is illustrated in Figure 3-4. A logical "1" state is initialized in the loop between nTrons U_1 and U_2 as in Figure 3-3a, and the switching event of U_2 occurs upon the application of the clock signal just as in Figure 3-3c. However, the hotspot forms and diverts the clock through the loop between U_2 and U_3, but the clock is still flowing through the channel of U_3, causing it to switch as well as shown in Figure 3-4c. This process continues, and all of the logical "0" states in the shift register that proceeded the initial "1" state will be overwritten, thus rendering the circuit useless as a shift register.

3.1.3 Wide-gate nTrons: symmetric vs. asymmetric channel placement

The shift register uses wide-gate nTrons (as shown in Figure 3-5). The wide-gate design is preferred over the narrow choke constriction in [12] (shown in Figure 2-1b) because the gate critical current needs to be comparable to the channel critical current, since the output current of one shift register stage is used to drive the next shift register stage. That is, the nTron should only switch upon the application of current through the drain and channel of the nTron when there is already current present in the gate. If the gate nanowire were to switch from the diverted clock current from a previous shift register stage, then there would be no way to store the state of the previous stage in a persistent circulating current. This means that the current through the gate on its own should never be enough to cause the nTron to switch; the nTron should only switch when there are currents of (roughly) equal magnitude flowing through both the gate and drain.

Furthermore, in addition to the nTron having a wide gate, it should also be symmetric, as opposed to the asymmetric design in [12]. That is, the constriction of the channel should be centered with the middle of the gate terminal, as shown in Figure 3-5). This is due to two competing desires for the location of hotspot formation. On one hand, we would like the constriction to be on the source-side of the channel, below the point at which the gate injects current into the channel. This would increase

the selectivity of the nTron to switch only current is supplied on both the gate and drain, but not switch when only one of the two terminals is excited. If the channel constriction is on the drain-side (*i.e.* above the gate) or symmetric around the gate terminal, then the selectivity is decreased, since the input current from either the gate or drain terminal needs to be higher in order to cause the combined current to switch the source nanowire. On the other hand, a channel constriction below the gate would allow current to leak between the gate and drain terminals of the nTron, even when it is in the resistive state. This is undesirable for a shift register, as it means the clock signal needs to be larger in order to store enough loop current when the nTron switches. Therefore, the desire to isolate the gate and drain terminals when the nTron is in the resistive state while also providing large margins on input levels required to switch the nTron necessitates a symmetric nTron design.

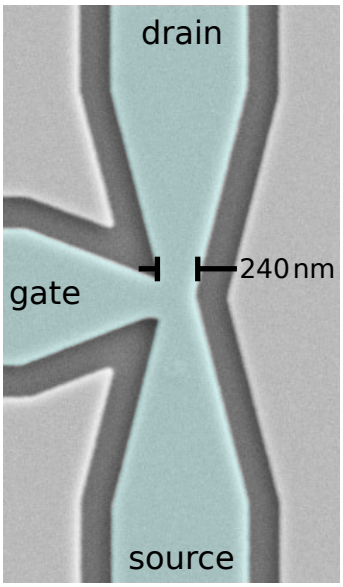

(a) False-color SEM of wide-gate symmetric nTron. SEM credit: Matteo Castellani.

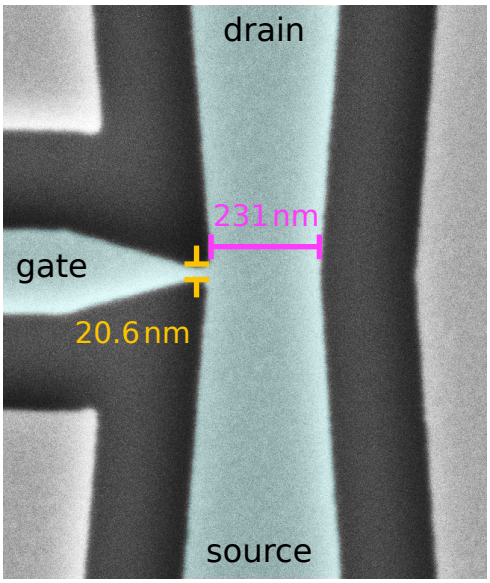

(b) False-color SEM of narrow-gate symmetric nTron. The gate constriction is 10× smaller than the channel constriction. SEM credit: Matteo Castellani.

Figure 3-5: Scanning electron micrograph of wide-gate and narrow-gate symmetric nTron. The nTron in a is part of the shift register. The width of the gate constriction is equal to the channel constriction width. The gate nanowire injects current into the channel at the center of the channel constriction. Both SEMs were captured by Matteo Castellani.

3.2 Simulation of shift register in LTSpice

Circuit simulation provides a useful tool for selecting nTron dimensions and kinetic inductor sizes due to the ability to rapidly iterate and test ideas. However, a simulation cannot hope to predict real world behavior if it doesn't accurately replicate both the circuit and the environment in which it will be tested.

3.2.1 Symmetric nTron model compatible with wide-gate nTrons in LTSpice

An existing compact model for simulating nTron in LTSpice is described in [7]. Despite being designed for asymmetric nTrons, minor modifications to the model have shown it to be an effective tool for designing circuits with narrow-gate, symmetric nTrons [1, 2]. However, a new model is needed to more accurately capture the behavior of wide-gate, symmetric nTrons.

Compact simulation models of superconducting nanowires have been demonstrated in [6] based on a phenomenological model of hotspot dynamics in superconductors from [28]. The model for a single superconducting nanowire was extended by [7] to create a model for nTrons which combined four separate nanowire models configured in a T-shape. The gate nanowire state (superconductive/resistive) and current is used to control the suppression of the critical current of the channel nanowire based on an empirical model from experimental data on nTron transfer characteristics. As mentioned previously, the model has been modified to allow the source nanowire to switch when the current through it exceeds the modulated critical current of the nTron channel $I_{\mathrm{m,ch}} I_{\mathrm{ch}}$. When used to simulate the behavior of circuits using narrow-gate symmetric nTrons, this modification is effective. However, when used to simulate wide-gate symmetric nTrons, the observed behavior does not quite line up with what would be expected in a physical device.

In this work, the nTron model of [7] was adapted to match the symmetric nTron geometry used in the shift register. Figure 3-6 shows the symmetric nTron model equivalent circuit. The expression for the channel critical current suppression due to

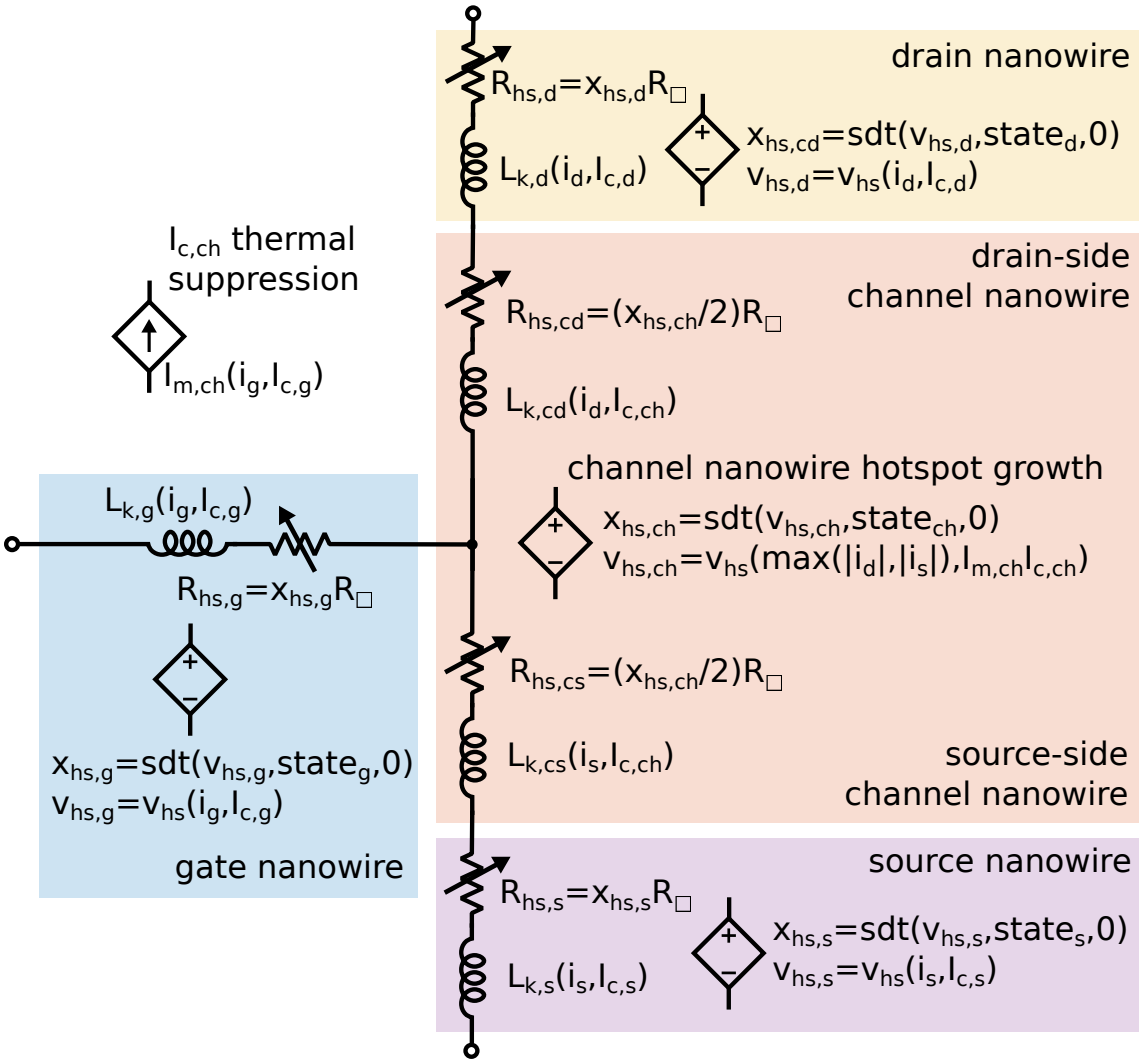

Figure 3-6: Symmetric nTron model. The gate, drain, and source nanowire are identical to the nanowire model from [6], with the hotspot integrator circuit replaced with a single behavioral source using the `sdt()` command. `sdt()` integrates the hotspot velocity when the nanowire switches to the resistive state, and resets the hotspot size to zero when the nanowire returns to the superconductive state. The hotspot velocity expression $v_{\mathrm{hs}}(i_w, I_{c,w})$ is identical to that of [6]. Each variable resistor models the hotspot of each portion of the nTron, and is set to zero resistance when its respective portion of the nTron is superconducting. The thermal suppression of the channel switching current from the gate nanowire uses the same expression as [7]. The channel nanowire is split in half, assuming symmetric hotspot growth outwards from the center of the channel. Note that the channel hotspot is shared between the two halves of the channel nanowire; the channel hotspot is assumed to grow symmetrically outwards from the center of the channel where the gate enters the channel.

heating in the gate nanowire is the same as in [7]. The expression for the hotspot velocity is the same as in [6]. In addition to the topology modification, the hotspot resistance integrator for all nanowire models was replaced with the LTSpice builtin `sdt()` as proposed in [8]. The integrated hotspot size is reset to zero when the nanowire is in a superconducting state. Using `sdt()` instead of the original capacitor-integrator in the hotspot growth model from [6] improved simulation speed of a circuit with 20 nTrons by about a factor of 2, and as discussed in [8], is expected to greatly improve reliability of simulation results. The LTSpice netlist for the symmetric nTron model is included in Appendix A.

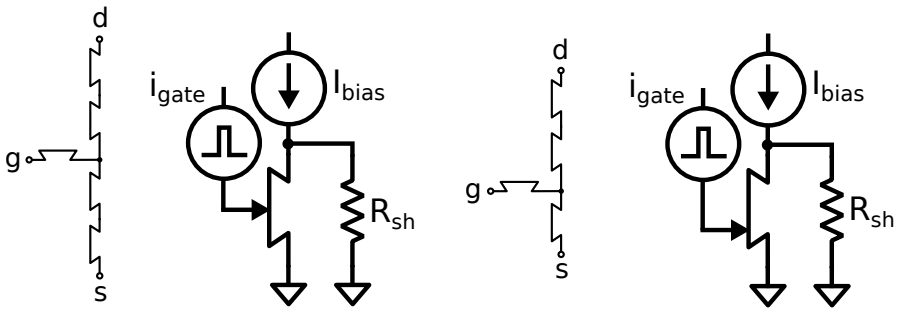

(a) Symmetric model testbench (b) Asymmetric model testbench

Figure 3-7: Simulation testbench for comparison of symmetric and asymmetric wide-gate nTron models. Both nTrons have the same geometry: `width_g` = 50 nm, `width_s` = 200 nm, `width_d`= 200 nm, `width_c` = 108 nm, `sq_g` = 22, `sq_d` = 400, `sq_s` = 22, `sq_c` = 10, `thickness` = 4 nm, `sheetRes` = 400 Ω, `Tc` = 10.5 K, `Tsub` = 4 K, `Jc` = 46 GA/m^2, `C` = 1, `A1` = 0.4. The shunt resistance is 15 Ω for both simulations.

This model was compared with the existing nTron model to confirm that it produced similar results for conventional narrow-gate nTrons. However, for wide-gate nTrons, the behavior of the two models differ slightly, as shown in Figures 3-8 and 3-9. The first difference to note is that the channel nanowire switches in the symmetric model, whereas the source nanowire switches in the asymmetric model. Given the current amplitudes and geometries in this test, an asymmetric nTron actually shouldn't switch. However, the asymmetric nTron switches due to the modification to make it behave more like a symmetric nTron. This modification only uses the channel nanowire properties to set the state of the source nanowire, however; the hotspot dynamics are still described by the width of the source nanowire, so the peak

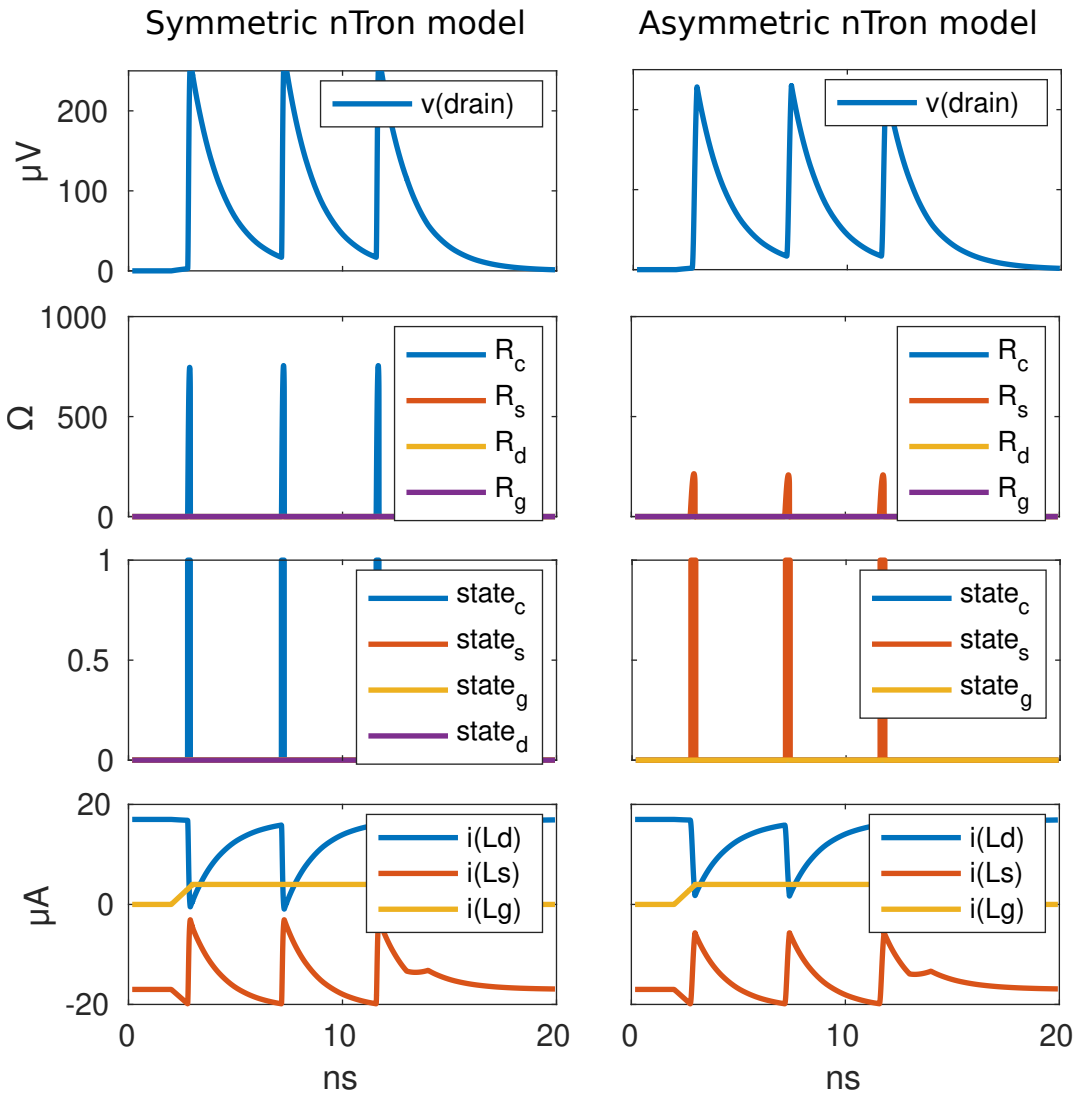

Figure 3-8: Simulation results of symmetric and asymmetric wide-gate nTron models for small gate current. Both models exhibit roughly the same behavior, although in the asymmetric model, the source nanowire switches instead of the channel, and the hotspot resistance is substantially smaller. A 17 µA bias current I_{bias} is used, and the gate current i_{gate} is a 10 ns-long, 4 µA pulse.

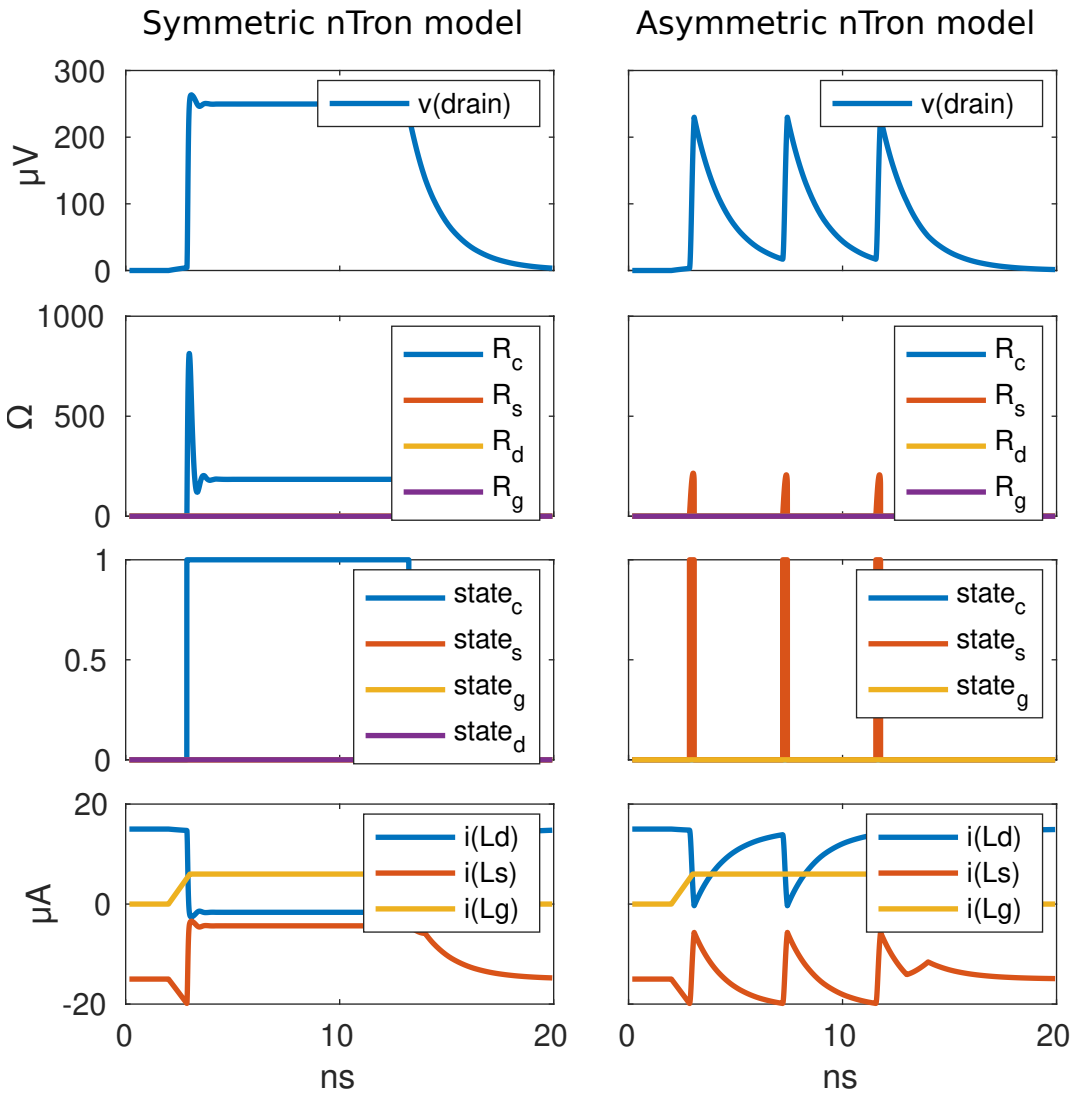

Figure 3-9: Simulation results of symmetric and asymmetric wide-gate nTron models for large gate current. The symmetric model latches, since the gate current is large enough to sustain the channel hotspot even after the drain bias is diverted to the shunt resistor. A 15 µA bias current I_{bias} is used, and the gate current i_{gate} is a 10 ns-long, 6 µA pulse.

hotspot resistance is lower as shown in Figure 3-8. This explains the slight difference in height of the voltage across the resistive load, since the hotspot doesn't grow as large in the source nanowire (which is wider than the channel nanowire). Furthermore, the asymmetric model preserves a superconducting path between the gate and drain nanowire, meaning that any gate current will continue to flow into the load. However, the result produced by the two models is qualitatively the same; something in the nTron switches and diverts the drain current into a load in both cases.

The second, more drastic difference between the two models can be observed when the gate current is increased: the symmetric model latches, while the asymmetric model does not. This is simply because the gate current is larger than the retrapping current of the channel nanowire, so after it switches, it will not return to the superconducting state until the gate current is removed. Again, this behavior is related to the asymmetric model using the critical current of the channel nanowire to decide if the source nanowire should switch.

It should be noted that these differences only pertain to wide-gate nTrons. When the width of the nTron gate is very small in comparison to the channel width (*e.g.* 10× smaller), the two models behave virtually the same.

3.2.2 Noise and parasitics

In addition to the previously-discussed nTron model, simulations of the shift register need to take into account non-idealities of experimental setups, such as noise, parasitics, and distributed effects.

Since the devices are quite small compared to their operating frequency (1 cm chip side length vs. 30 cm wavelength for 1 GHz signals), modeling of distributed effects is only necessary for the coaxial cables used to interface with room temperature electronics. Bondwire self-inductance is typically on the order of 1 nH/mm, giving a typical total inductance of roughly 5 nH for the bondwires used in this setup. This is significantly smaller than the loop kinetic inductors used in the shift register, so it does not have a substantial impact on the simulation, but for evaluating the feasibility of scaling to smaller, faster electronics, it is an important quantity to include.

Noise modeling was performed with `spice-daemon`, software that augments LTSpice with a variety of functions, including the ability to generate noise sources with arbitrary distributions and spectra which update every time a simulation is run [8]. The noise from the room-temperature electronics is dominated by the AWG and LNA, so additive noise was included on the inputs and outputs of the device. The thermal noise from on-chip bias resistors (which would be at a temperature of 4 K) was not considered. The noise was modeled as $1/f$ noise with the `colorednoise` python library, which is based on [29]. The standard deviation of the noise distribution was initially chosen arbitrarily, and gradually increased to roughly 10 % of the maximum amplitude of the input signals, and roughly 50 % of the output signal. This helped guide the sizing of the nTrons and kinetic inductors towards a noise-resilient design.

The final design used 100 nH kinetic inductors for the loops, 270 nm-wide wide-gate nTrons for the inter-loop nTrons, and a readout nTron with an 240 nm-wide channel and 40 nm-wide gate. The shift clock bias resistors were chosen to be 1.96 kΩ, the readout clock bias resistor was 11 kΩ, and the shunt resistors were 50 Ω (for a combined shunt resistance of 25 Ω when including the input impedance of the coax with oscilloscope inputs set to 50 Ω). The bias resistors must be sufficiently large so that the current flowing through them does not dip significantly when the nTron they are biasing switches. If there is significant current droop during a switching event, the amount of current left circulating in the loop inductor will be lower, which may degrade the operating margins of the shift register (*i.e.* the range of clock amplitudes that result in correct operation). Bias resistors as small as 100 Ω were used successfully in simulation, however 1.96 kΩ was selected for the experimental device to increase the operating margins of the shift register.

3.2.3 Simulation results

A four-loop shift register was simulated in LTSpice to verify that the proposed operating principle was viable, and to gain an understanding for how the nTrons and kinetic inductors should be sized. A follow-on simulation was performed to explore various techniques for integrating the shift register with SNSPD arrays. A diagram

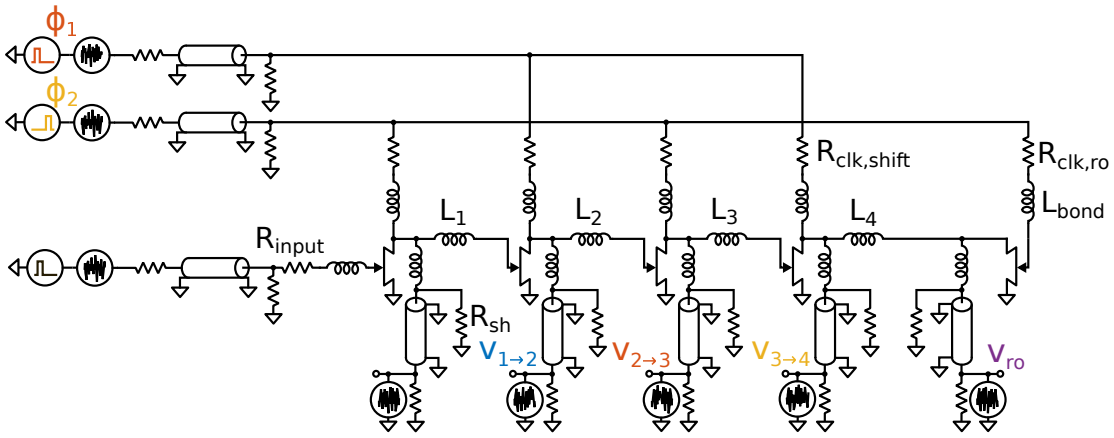

(a) Simulation setup for four-loop serial operation shift register. Noise sources, distributed effects from coaxial cables, and bondwire inductance are included. The shift clock bias resistors are 100 Ω, the readout clock resistor is 215 Ω, and the loop inductors are 15 nH. All nTrons have a channel width of 140 nm, and the readout nTron has a gate width of 35 nm. The shunt resistors, which were included in the fabricated shift register out of fear of the nTrons latching, are set to 10 kΩ (making the effective shunt resistance 50 Ω in total due to the oscilloscope and coaxial cables).

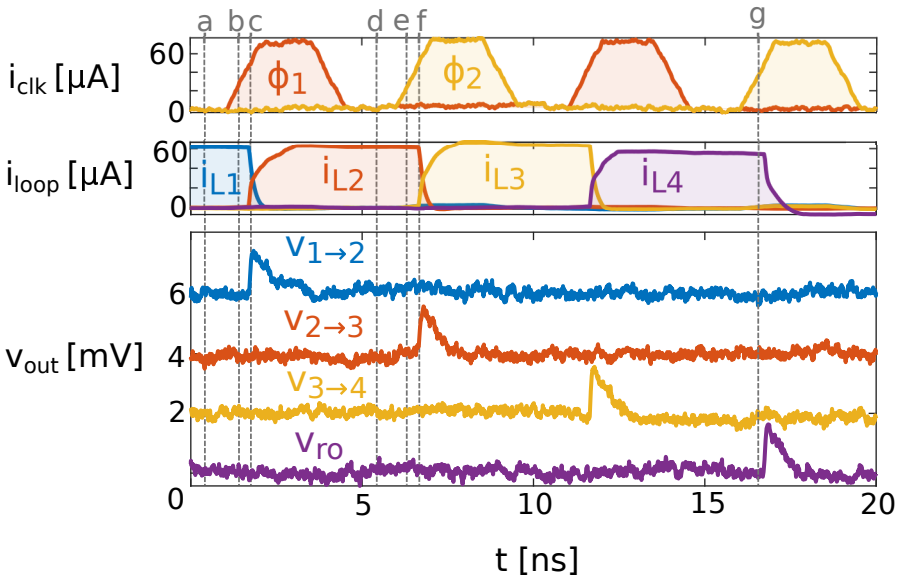

(b) Transient simulation results of the above circuit. Time labels a,b,c,d,e,f,g correspond to the subfigures of Figure 3-3. The voltage waveforms in the lower pane are offset vertically for clarity.

Figure 3-10: Circuit schematic and results of LTSpice simulation of shift register. The results in b show the shift register initialized with a circulating current in the first loop (the initialization uses the third voltage source in a). The state is transferred from loop to loop following the procedure described in section 3.1.2, and the state transitions are indicated with voltage pulses when each nTron switches.

of the simulation setup with just the shift register (no SNSPDs) is shown in Figure 3-10a, and the results of the simulation in 3-10b. The results presented in Figure 3-10 are for a shift register design with slightly different nTron dimensions and loop inductances than the fabricated shift register, but the behavior is qualitatively the same.

Just as in Figure 3-3a, the shift register is initialized with a circulating current in the first loop. When the first clock phase is applied, the current in the first loop is destroyed, and the clock is diverted, forming a current in the second loop. The normal resistance of the nTron hotspot creates a voltage spike that decays exponentially (due to the L/R time constant of the loop kinetic inductance combined with the $25\,\Omega$ shunt resistance). Each spike in the bottom pane of 3-10b indicates a switching event in one of the nTrons: $v_{1\to2}$ indicates the transfer of state between the first and second loop, $v_{2\to3}$ between the second and third, and so on. In simulation, we are able to directly observe the loop currents (the shift register state) as well as the state and hotspot resistance of each nanowire in the nTron models. However when characterizing the shift register in experiment, we can only see the voltage spikes (transitions between shift register states), which just indicate that an nTron has switched, so it is useful to get comfortable with inferring the state of the shift register from the state transitions.

3.3 Shift-register-based SNSPD array readout

One of the challenges with SNSPD array readout is the scaling of cable count with imager size. Each SNSPD needs to be biased near its switching current to maintain high detection efficiency, and all of the information of photon arrivals at each pixel needs to be collected somehow. A variety of readout techniques have been demonstrated [30, 4, 31, 32], all of which are designed preserve the timing information of incident photon arrivals with as high of precision as possible. The motivation for this design decision is to preserve the fine timing resolution of SNSPDs that beats out other detectors. However, SNSPDs are attractive over other kinds of photon detectors for applications such as spectroscopy where timing information is non-critical,

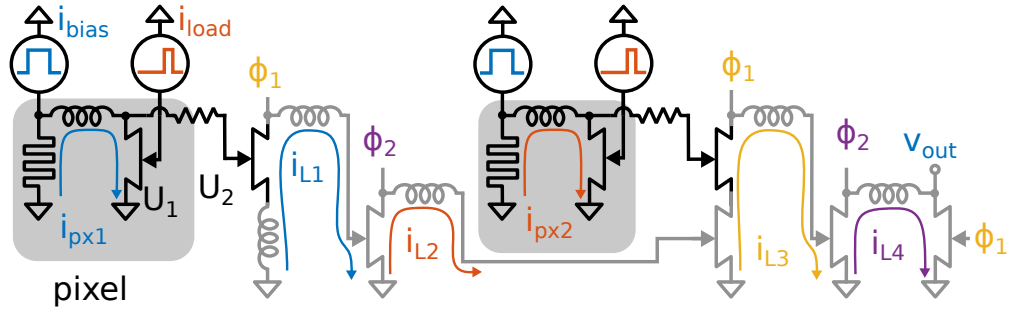

(a) 2-pixel SNSPD array readout with a shift register. Each SNSPD is part of a destructive readout binary memory which stores a circulating current when a photon hits the SNSPD. Upon the application of a current pulse i_{load}, the memory state of all pixels are simultaneously read out in parallel and loaded into the shift register with nTrons U_1 and U_2. The shift register then proceeds to operate as described in section 3.1.2 to read out the SNSPD data serially. This process is pipelined, so as soon as the pixels are read into the shift register, they are sensitive to photons again.

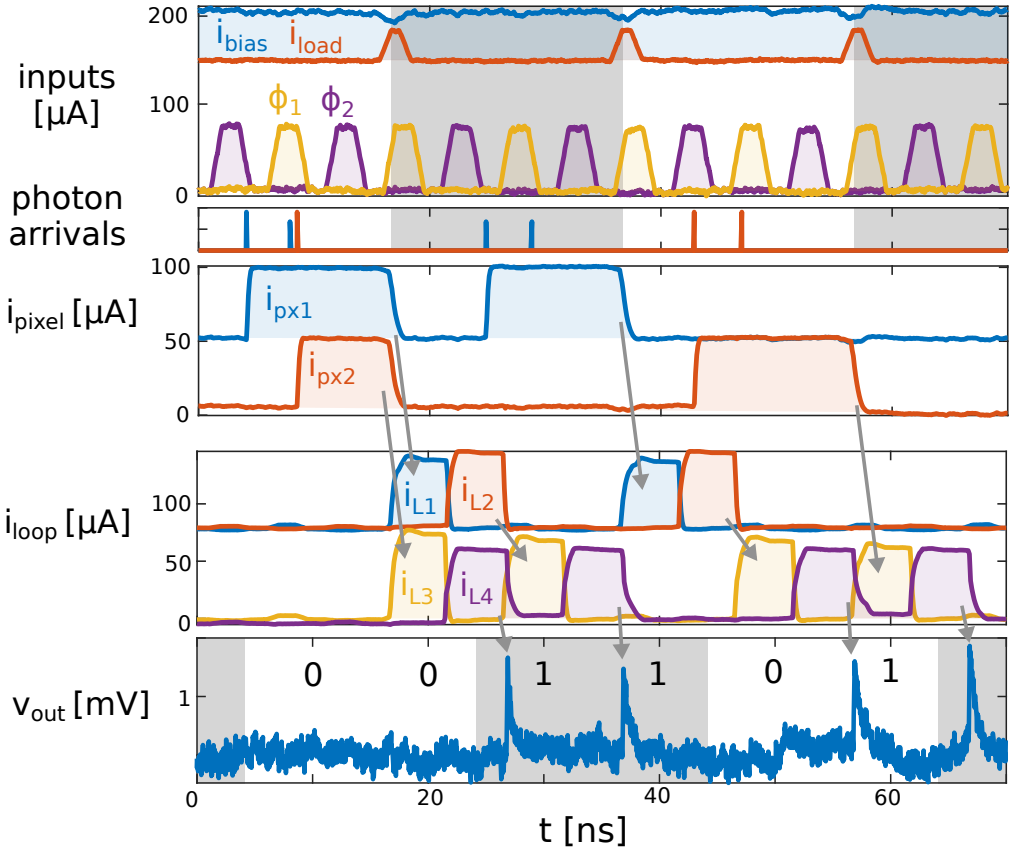

(b) Simulation results of 2-pixel SNSPD array readout with nTron shift register. The alternating grey/white regions represent subsequent image frames. In the first interval, a photons hit both SNSPDs. In the next interval, that result ("11", where the most significant bit corresponds to the rightmost pixel) is read out with the shift register, while the SNSPD bias is still enabled. Two photons hit the first pixel. In the third interval, the result "01" is read out, and two photons hit the second pixel. In the final interval, the result "10" is read out.

Figure 3-11: Schematic and simulation of SNSPD array readout technique.

but high detection efficiency, mid-infrared sensitivity, and low dark count rates are still critical factors [33].

If timing precision below a few microseconds is not critical, then the variety of readout techniques that can be used is increased. Shift register readout of an SNSPD array could follow a similar approach to CCDs, in which the imager integrates photons for some period of time, and then serially shuffles out the integrated photon counts one pixel at a time. Connecting the superconducting binary shift register presented in this work to an array of SNSPDs would provide a straightforward path to scaling SNSPD-based imagers to the megapixel scale and beyond.

An example of how this might be constructed at an array scale is shown in figure 3-11. Each SNSPD is placed in a destructive-readout memory cell, which stores a binary state based on the branch that the bias current i_{bias} flows through. When a photon hits the SNSPD, it becomes resistive, diverting the bias current into the right branch of the pixel. The nTron U_1 is used to read out the memory with the application of the current pulse i_{load}. If the bias current is flowing through the right branch of the pixel, the hotspot in the gate of U_1 will cause it to switch, generating a pulse of current that flows through the gate of U_2 while also resetting the branch current back into the left branch through the SNSPD. The pulse of current through the gate of U_2 will cause it to switch, storing a circulating current i_{L1} in the first loop of the shift register. The shift register state is then read out just as described in section 3.1.2.

Figure 3-12 shows a block diagram of a proposed 1024-pixel SNSPD array with shift register read out on a single wire. The maximum photon count rate (across the imager) is limited by the clock rate of the shift register. The imager can detect at most 1024 photons between frames, and the frame rate is set by the imager size and shift register clock rate. For example, if the shift register is clocked at 100 MHz, then it would take roughly 10 µs to read out a single frame of the 1024-pixel imager, setting a maximum photon count rate at 100 Mcps. However, the implementation of the readout is flexible: depending on the count rate requirements and number of cables available, the imager can be split into banks to operate at higher count rates.

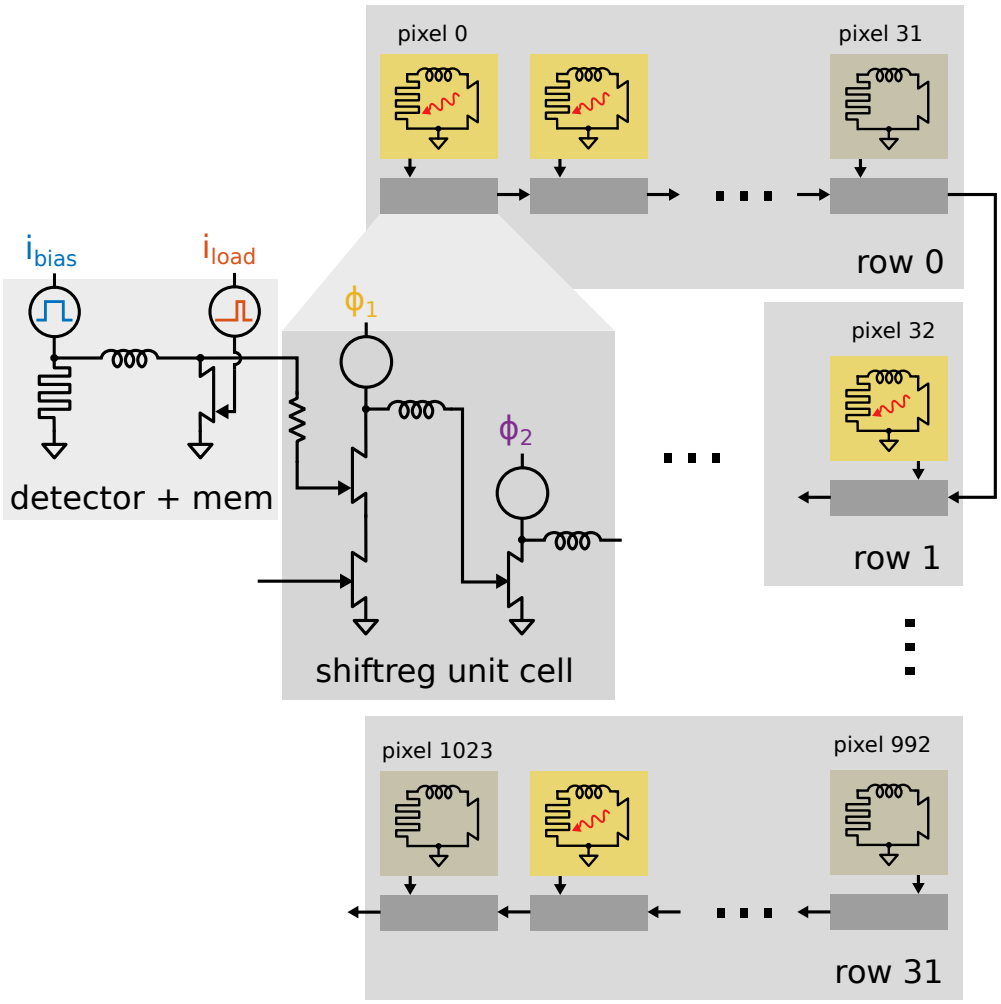

Figure 3-12: Example 1024-pixel SNSPD array utilizing shift register readout. All 1024 pixels are destructively read out into the shift register at once. The shift register then serially outputs the pixel data, while the SNSPDs in each pixel detect more incident photons.

If read out on 16 wires at the same shift register clock rate as before, the maximum count rate is increased to 1.6 Gcps.

If extremely high count rates are needed, but only for a short period of time, it is possible to slightly modify the architecture of the imager to act like a high speed camera, where each pixel has a multi-frame in-pixel memory implemented with a small shift register. This breaks the imager speed-size tradeoff, since the local, per-pixel memory makes the maximum count rate per pixel independent of imager size. In the continuous output architecture described by Figure 3-12, larger imagers require longer inter-frame periods, since more data has to be clocked out, which hurts the

frame rate of the imager. The in-pixel memories could be read out at a much slower rate after the event of interest is over, making Tcps imagers possible (*e.g.* a 100×100 pixel array with in-pixel memories clocked at 100 MHz could achieve a peak count rate of 1 Tcps). Frame depth is a concern, since the introduction of very large in-pixel memory may hurt the fill-factor of the imager. Furthermore, the per-pixel frame rate is limited by the speed of the SNSPD. This is not a huge concern, since large active area SNSPDs with count rates as high as 300 MHz have been demonstrated [34]. This type of imager could have interesting applications in particle physics experiments, where high luminosity beams are expected to generate a lot of particle detection events in a very short period of time [35].

3.4 Summary

While the information in this chapter focuses on design and simulation of shift registers and their application in SNSPD array readout, the tools and techniques presented are applicable for any superconducting nanowire electronics. More broadly, the general approach presented in this chapter for working with nTrons is useful for anyone working on electronics with novel devices.

Chapter 4

Fabrication and characterization of the nanocryotron shift register

This chapter discusses the fabrication of the shift register, the experimental setup used to characterize it, and the characterization results. The characterization measures a bit-error-rate of the shift register across various input signal levels, frequencies, and under the application of an external magnetic field. The results of the characterization are presented as bias-margin plots, which provide a graphical way of understanding the behavior of the shift register under specific operating conditions (*e.g.* clock frequency, strength of an external magnetic field), as well as a technique to evaluate scalability of the circuit.

Matteo Castellani performed the NbN deposition and lithography, and also took scanning electron micrographs of the fabricated device.

4.1 Fabrication

The shift register was fabricated with a single layer of 20 nm-thick NbN. The fabrication process is shown in Figure 4-1. The layout of the shift register was designed programmatically using the python package `phidl` [36, 37]. Figure 4-2 shows the layout of the shift register that was tested and characterized in this work.

Eight shift registers with various designs were fabricated on a single 1 cm×1 cm

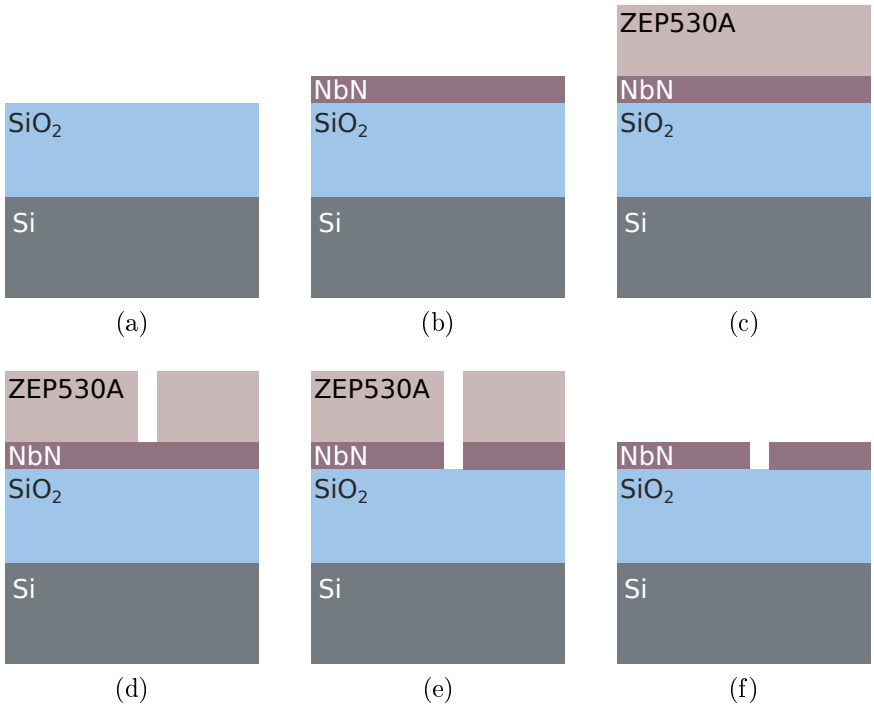

Figure 4-1: Cross sections of the shift register chip at various stages in the single-layer NbN nanowire fabrication process. (a): starting Si chip with 300 nm SiO_2. (b): deposition of 16 nm of NbN. (c): spin-coating of ZEP 530A electron-beam resist. (d): development of resist. (e): pattern transfer with reactive ion etching. (f): final resist removal.

silicon chip with 300 nm thermal SiO_2 oxide. The 16 nm-thick NbN film was deposited in an AJA sputtering system with reactive magnetron sputtering of a niobium target in a nitrogen atmosphere at 153 W peak power (see Figure 4-1b). The deposition pressure was 2.5 mtorr, and the nitrogen and argon flow rates were 6 sccm and 26.5 sccm respectively. The deposition rate was roughly 2 nm/min.

The film was then patterned with electron beam lithography and reactive ion etching pattern transfer. Electron-beam lithography is desired due to the shorter wavelength of keV electrons as compared to visible and ultraviolet photons used in photolithography processes, therefore allowing for a smaller diffraction-limited features (the minimum feature size of the shift register is 30 nm, which is unresolvable with near-ultraviolet lithography). ZEP530A positive-tone resist was spun on at 5000 rpm for 60 s and baked at 180 °C for 120 s. Electron-beam lithography was used to write the resist in an Elionix ELS-F125 with a dose of 550 $\mu C/cm^2$ (4-1c). The

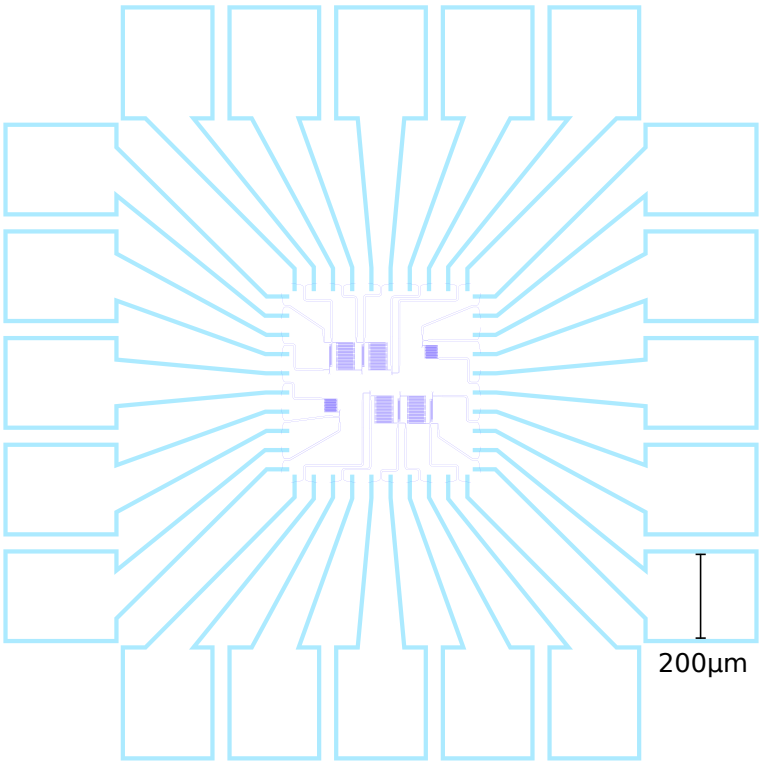

(a) Overview of shift register layout including pads. The shift register circuit is duplicated to ease layout and provide a second circuit if there is an issue with fabrication that renders the first inoperable.

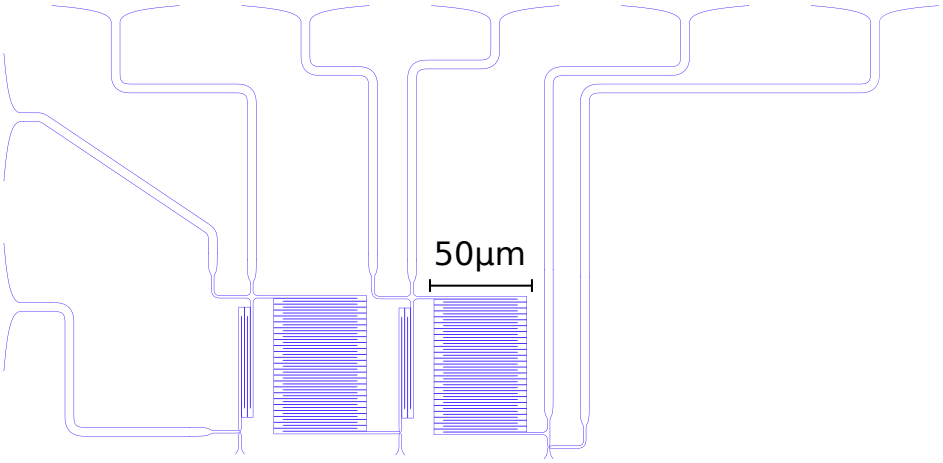

(b) Detail of the shift register layout.

Figure 4-2: Layouts generated with the help of `phidl`. ZEP530A is a positive-tone resist, so the pattern shown will be etched out of the NbN (*i.e.* in order to make a wire, the pattern uses two lines, resulting in two trenches in the NbN on either side of the wire). (a) shows the GDS pattern of the NbN etch for a pair of identical shift registers and pads to allow for wirebonding to the shift registers. The pads (cyan in (a)) are wider and placed on a different layer so they can be written with higher beam current to speed up fabrication. (b) shows the low-current layer in more detail.

exposed resist was developed in o-xylene at 5 °C for 60 s and rinsed with IPA at room temperature (4-1d). Reactive ion etching with CF_4 at 50 W of power was used to remove the exposed NbN layer (4-1e). The remaining resist was stripped in a heated bath of n-methyl-2-pyrrolidone (NMP) at 70 °C for 1 hour (4-1f).

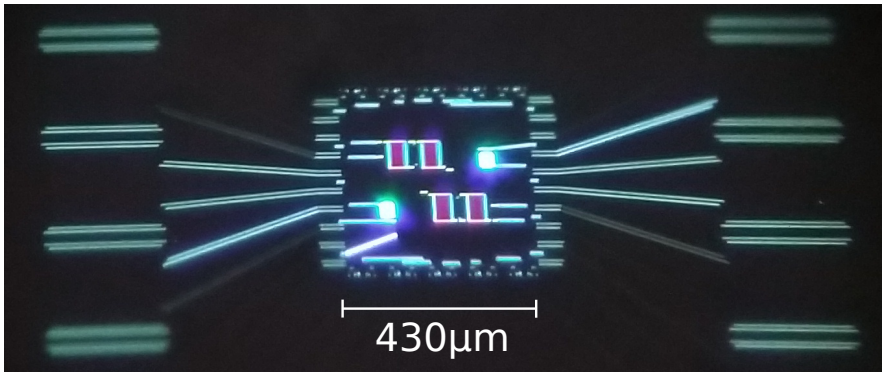

(a) Optical micrograph of shift register. The kinetic inductors form diffraction gratings that appear pink under illumination.

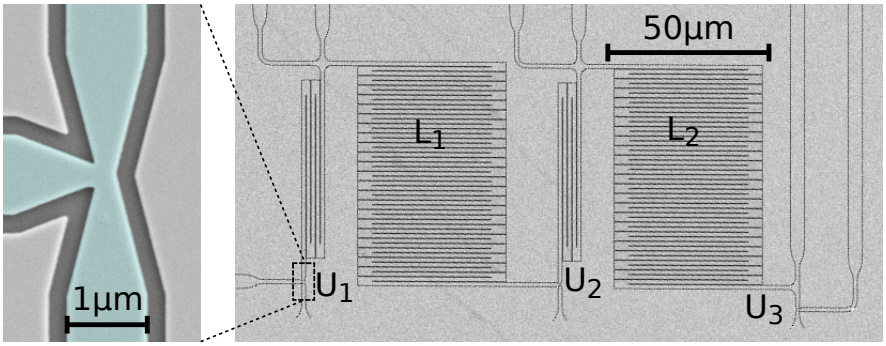

(b) False-color scanning electron micrograph of shift register and closeup of wide-gate nTron. The nTron is colored light blue and the surrounding ground plane is light grey. The dark grey substrate is visible at the bottom of the trenches etched in the NbN around the nTron. The large kinetic inductors are labeled L_1 and L_2. The two wide-gate nTrons each have a 200 square inductor in series with their drain. SEM credit: Matteo Castellani.

Figure 4-3: Optical and electron micrographs of shift register chip.

After deposition, the sheet resistance of the film $R_{n,\square}$ was measured to be 169 Ω per square. After lithography, the resistance to ground for each input of the shift register was measured and compared with the expected total resistance based on the sheet resistance to verify that there weren't any significant errors in the lithography process. The T_c one of the wires was measured to be 9.4 K. Based on the measured sheet resistance and T_c, the kinetic inductance $L_{k,\square}$ was estimated to be 25 pH per

square using Equation 2.1.

Of the eight shift registers that were fabricated, two were tested and only one was characterized thoroughly. The first shift register that was tested used a different switch geometry than the wide-gate symmetric nTrons. It did not work because the switch geometry did not block current from flowing between loops even when it was switched. The second design that was tested and characterized was based on wide-gate symmetric nTrons. Figure 4-3 shows an optical and scanning electron micrograph of the second design. The results of this shift register will be discussed in more detail in the following sections.

In the scanning electron micrograph shown in Figure 4-3b, there are two large kinetic inductors L_1 and L_2. The left inset shows a closeup of one of the wide-gate symmetric nTrons used to transfer state between the superconducting loops. Note that there are two unlabeled kinetic inductors; each one is in series with the drain a wide-gate nTrons. These were included under the erroneous assumption that the nTrons would otherwise latch in a resistive state without a small load impedance. While it is the case that a large real impedance will cause an nTron to latch, an inductive load with zero real impedance will allow the nTron to reset. In the shift register, L_1 (L_2) and U_2 (U_3) provide a superconducting path to ground when U_1 (U_2) switches, ensuring that the nTrons do not latch.

The extra kinetic inductance should not significantly hurt the performance of the shift register, although it is expected to limit the maximum operating frequency of the shift register to about 500 MHz since the LR circuit formed with the shunt resistance acts as a low-pass filter on the current passing through the nTron. Another disadvantage this inductor introduces can be seen in Figures 4-7 and 4-8, where the LR circuit acts as a differentiator on the voltage across the shunt, producing small output voltage pulses on the rising and falling edges of the clock signal. This is disadvantageous because it makes it harder to distinguish between a switching event and a fast-edge clock signal charging/discharging a resistively-shunted inductor.

4.2 Measurement and characterization

In order to test the shift register, it was cooled down in a liquid helium dewar and connected to room temperature electronics that measured the bit error rate (BER) of the shift register over a range of operating conditions and signal levels. The BER is a measure of errors in the shift register output signal, normalized to the total number of bits sent into the shift register. Testing the BER of a shift register is relatively straightforward, since it should just output every bit it was sent (in order with no changes or bit flips).

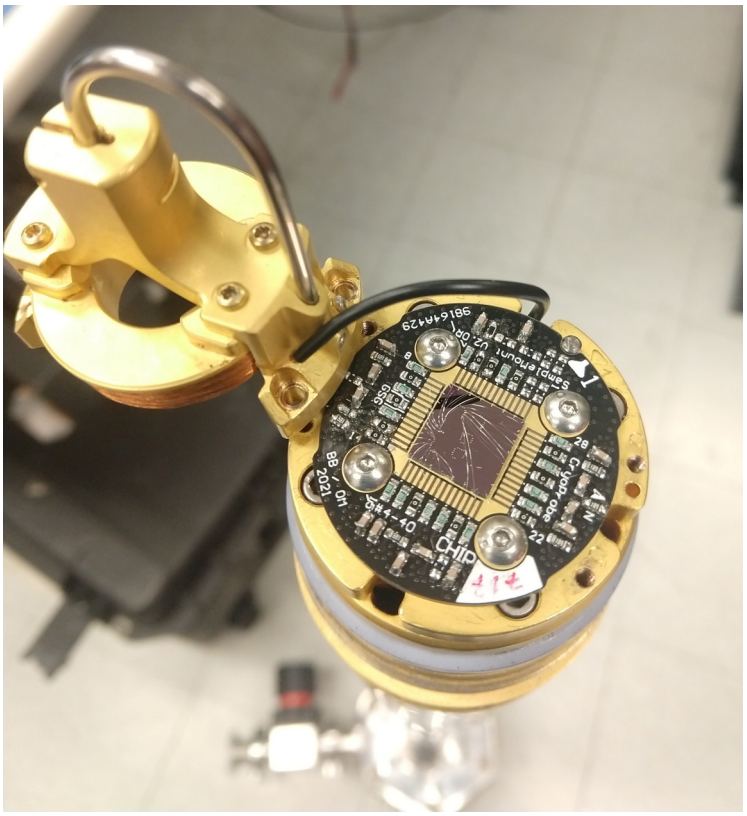

(a) Shift register chip and sample PCB mounted to cryogenic dip probe. The wirebonds are visible as thin hairlike wires across the surface of the chip. The coil of the superconducting electromagnet used for applying out-of-plane magnetic fields can be seen at the bottom of the mounting bracket to the left of the PCB. After installation of the PCB, the bracket is spun around to cover and hold the PCB securely in place.

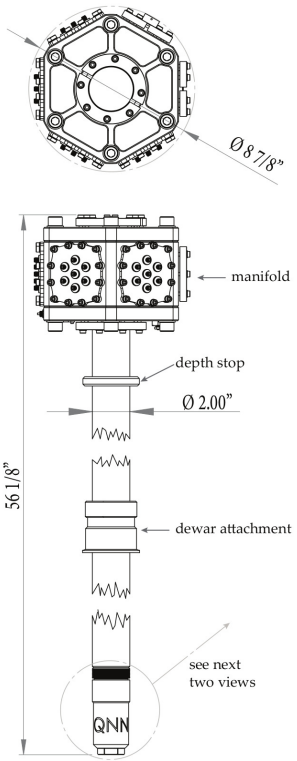

(b) Cryogenic dip probe mechanical diagram. The endcap labeled "QNN" covers the sample and forms a seal to allow the probe to pump down to vacuum. Figure from [38].

Figure 4-4: Cryogenic dip probe setup used to test the shift register at 4.2 K in a liquid helium dewar.

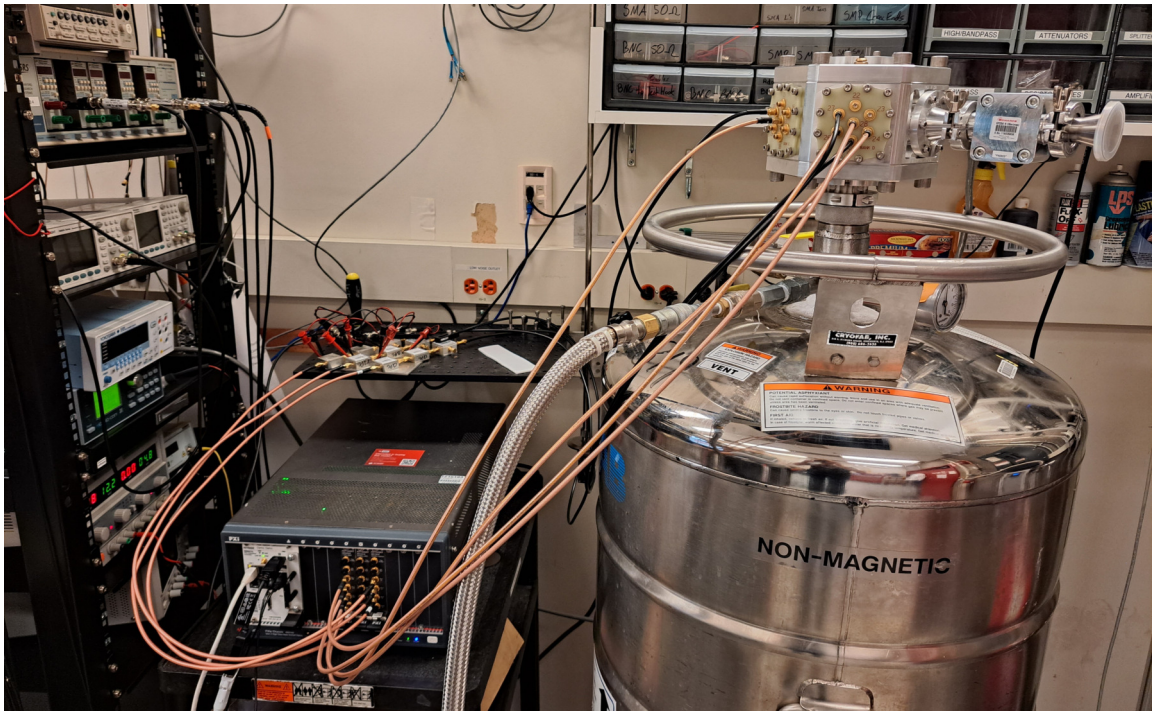

Figure 4-5: Shift register characterization experimental setup. The Keysight PXIe arbitrary waveform generator and digitizer are connected to cryogenic dip probe and low noise amplifiers. An equivalent circuit of the setup is shown in Figure 4-6.

The shift register chip was glued with GE varnish to a printed circuit board (PCB) to provide good thermal heatsinking of the chip. Wirebonds were used to make electrical connections to the circuit. The PCB was mounted at the end of a cryogenic dip probe [38] to make electrical connections to the experimental circuit. Figure 4-4 shows the chip and PCB mounted to the dip probe, and a mechanical diagram of the dip probe. Figure 4-5 shows the dip probe lowered into the dewar and connected to room temperature test equipment used to characterize the shift register.

Three separate characterizations of the shift register BER were performed. The first set of experiments were performed with an Agilent 33600A AWG and LeCroy Waverunner 620Zi DSO. The second and third set of experiments leveraged a Keysight data acquisition setup on loan. This chapter primarily discusses the second and third set of experiments. Furthermore, only Figure 4-11 contains data from the first two experiments; the rest of the data presented in this chapter was collected in the third set of experiments. The BER of the shift register was measured with a 1 GS/s Keysight

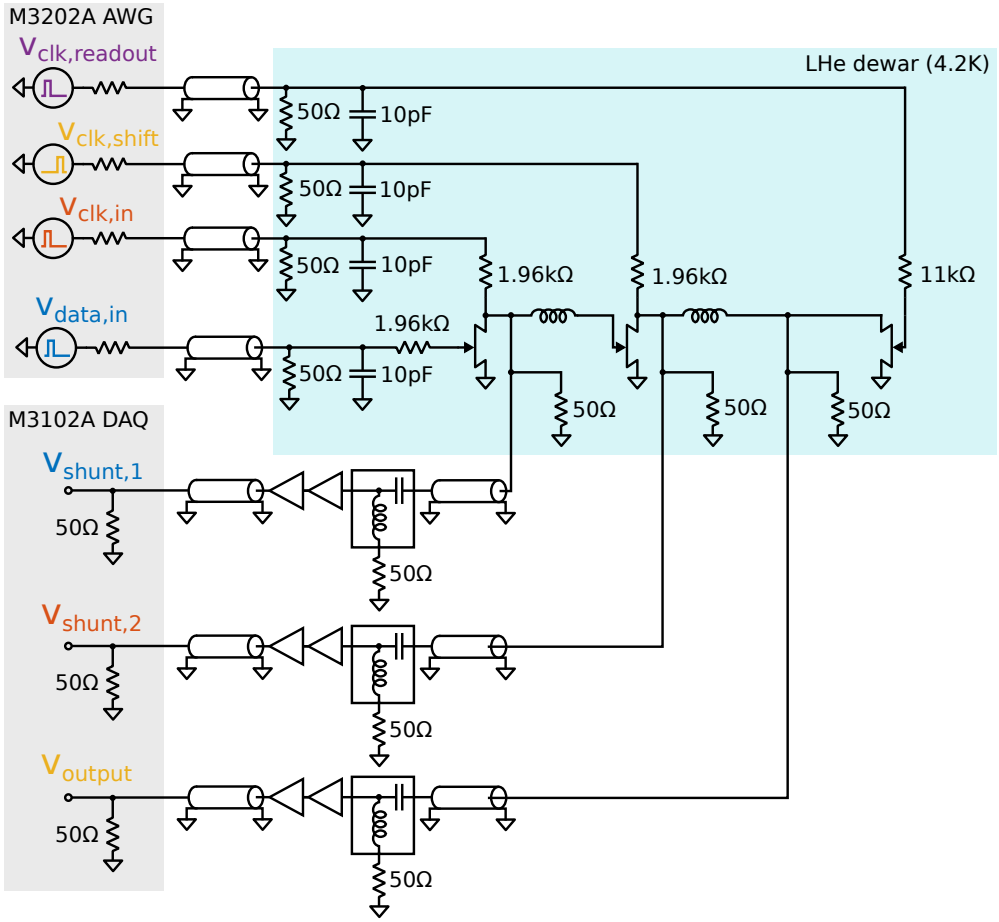

Figure 4-6: Equivalent circuit of experimental shift register circuit and readout electronics. The results presented in Figure 4-11 used 67 pF capacitors instead of the 10 pF capacitors shown here. In addition, the all bias resistors were 11 kΩ instead of 1.96 kΩ (note that the readout nTron bias resistor was unchanged). Each channel has two LNA-2000s from RF-Bay for signal amplification. The results in Figure 4-11a also used an Agilent 33600A AWG and LeCroy Waverunner 620Zi DSO in place of the Keysight AWG and DAQ.

M3202A PXIe AWG and 500 MS/s M3102A data acquisition/digitizer (DAQ). The AWG was used to generate the clock and data signals, which were produced by encoding binary symbols with 2.1 ns full width at half-maximum (FWHM) voltage pulses. Figures 4-7 and 4-8 show the input stimulus to the shift register, as well as its response. The data signal is a pseudorandom binary sequence (PRBS) of binary symbols, where a binary "1" is encoded with the presence of a voltage pulse, and a "0" is encoded with the absence of a voltage pulse as shown in Figure 4-7. The data and input clock signals ($v_{\text{data,in}}$ and $v_{\text{clk,in}}$) are in phase with each other, and the readout clock ($v_{\text{clk,readout}}$). These signals are part of the first clock phase ϕ_1 as discussed in section 3.1.2. The second clock phase ϕ_2 is comprised of the shift clock ($v_{\text{clk,shift}}$), which is 180° out-of-phase with the other signals.

Instead of using a PRBS stimulus, characterization with the AWG and DSO used a sequence of "1", "0", "1", "0" in order to avoid tedious programming of waveforms to the AWG. Based on the operation of the shift register, this sequence should give roughly the same BER as a random sequence, since any stray circulating currents left over after a shift operation would eventually accumulate and cause errors. However, it is plausible that in some circumstance, the particular pattern "1", "0", "1", "0" would enable the shift register to produce the correct output without actually performing a shift operation. Therefore, one of the big attractors of using the automated Keysight equipment was the ability to easily and rapidly upload long, arbitrary sequences of pulses to the AWG.

4.3 Bit error rate calculation

Calculation of the BER of a continuously-clocked shift register is quite straightforward, since the output should identically match the input, delayed by one clock period per logical stage. Therefore, for the two-loop shift register, we would expect the output to match the input exactly, delayed by one clock period. The delay is only one clock period and not two because of the use of a two-phase clock.

The calculations were performed with custom python software [39] that controlled

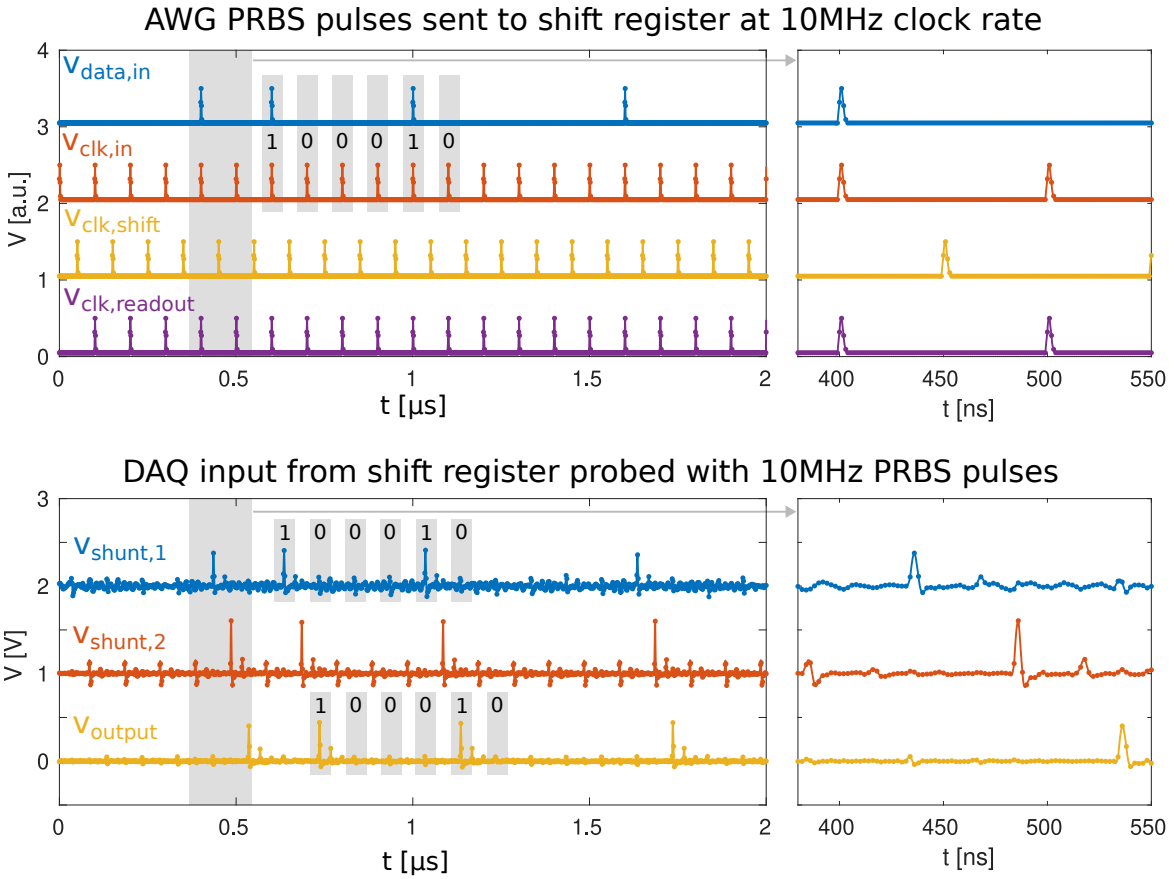

Figure 4-7: Transient response of shift register with 10 MHz clock rate. The upper plots show the input signal sent to the shift register. The lower plots show the signal measured by the digitizer. The actual voltage across the nTrons is about 300× lower, since the output of the shift register is passed through two LNAs with a total gain of 50 dB. Note the small, symmetric (positive and negative) pulses on the $v_{\text{shunt},1}$ and $v_{\text{shunt},2}$ signals that coincide with the clock signal from the AWG. These are due to the extra 200 square kinetic inductor in series with the drain of the first two nTrons. There are also noticeable reflections from the impedance mismatch between the LNA output and the DAQ input. These are visible as small pulses delayed from nanowire switching pulses by about 30 ns.

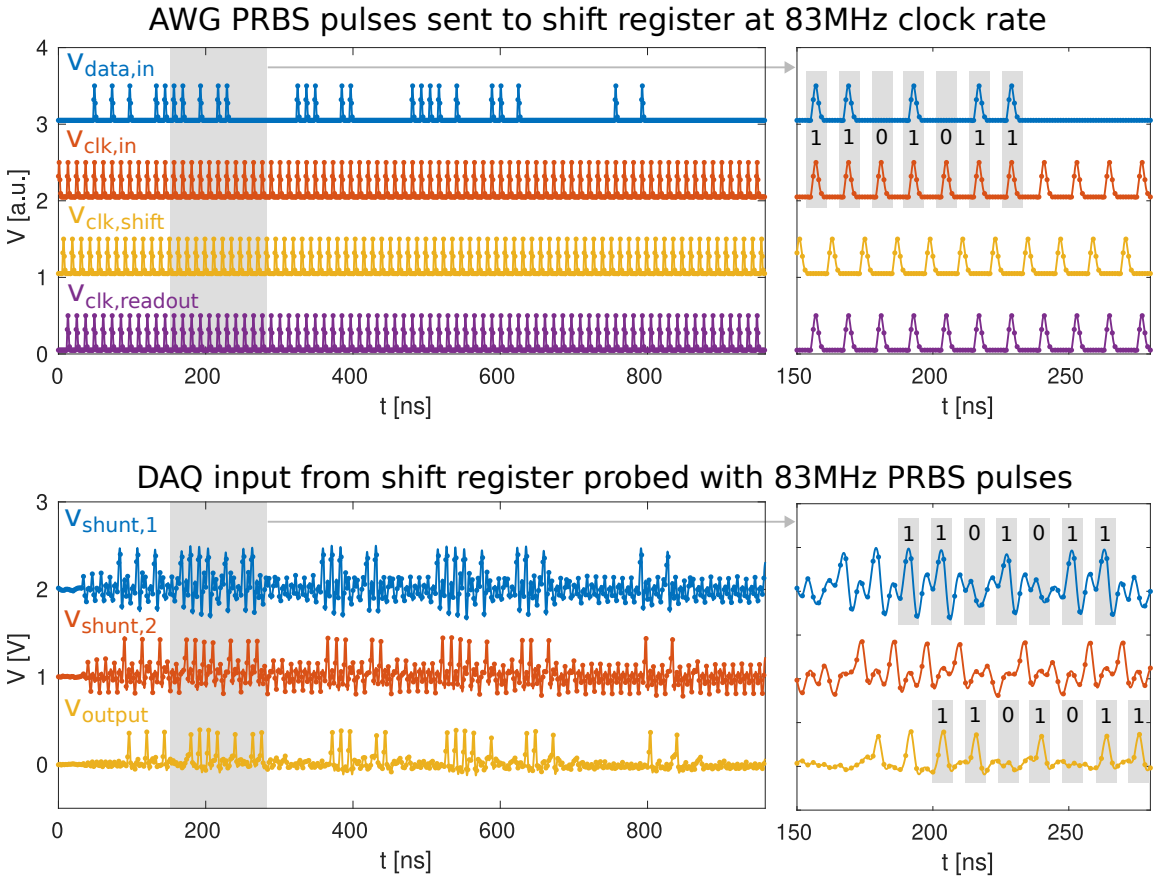

Figure 4-8: Transient response of shift register with 83 MHz clock rate. The two phase clock is very close to overlapping. Also of note are the symmetric pulses from charging and discharging the 200 square kinetic inductors, which are more apparent here than in the 10 MHz test. The reflections are buried beneath other pulses and cannot be seen easily here.

sample generation and collection as well as triggering of the AWG and DAQ in the PXIe chassis. First, before calculating the BER, the digitized waveform was deskewed to account for cable delays. This was performed once at the beginning of the experiment as a calibration. The shift register was supplied a set of clock and data signal amplitudes that were large enough to cause the first nTron in the shift register to switch when the data signal was a "1" but not when it was a "0". The cable delay was estimated by taking the cross-correlation between a downsampled copy of the input data signal and the digitized waveform of the voltage across the first nTron. A sample cross-correlation plot is shown in Figure 4-9. The lag at which the cross-correlation was maximal was saved as the cable delay. This delay was later used

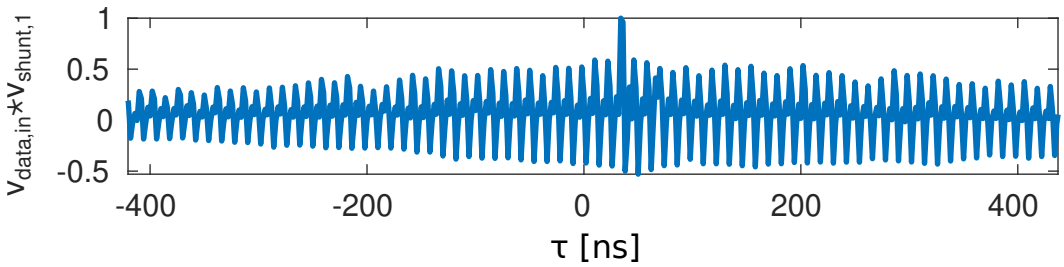

Figure 4-9: Cross correlation of AWG data and digitized waveform of the voltage across the first nTron in the shift register. The peak lag τ_{max} is 34 ns, which is consistent with the expected cable delay for about 5 m of coaxial cable.

to deskew digitized waveforms of the voltage across the output nTron, as shown in Figure 4-10b.

The BER was calculated by shifting the full-sample-rate digitized waveform of the shift register output by one clock period, then using `scipy`'s peak-finding algorithm to find the voltage spikes in the digitized waveform. The peaks were required to be at least 60 % of the maximum voltage in the entire waveform, and at least 7/8 of a clock period apart. This helped mitigate the effect of noise and reflections caused by impedance mismatches between the DAQ and LNA output impedance. The full-rate sample indices of the peaks were divided by the symbol period set by the shift register clock frequency (*e.g.* a clock frequency of 10 MHz would have a symbol period of 50 samples at the DAQ's 500 MS/s sample rate) to obtain a symbol index into the sequence of input symbols from the PRBS. An empty bit vector that represents the digitized symbols was filled with zeros, except where a peak was found, in which case the vector was set to "1" at that index. A bitwise exclusive or (XOR) was then performed between bit vectors of the PRBS input symbols and the digitized output. The total number of nonzero bits in the result is exactly the number of flipped bits, and was divided by the total number of symbols sent to calculate a bit error rate. This procedure was performed in near-real time on the Keysight PXIe setup, allowing for the generation of high resolution plots like the ones in Figure 4-11b and 4-11c in roughly 2 minutes.

BER measurements of the shift register with the AWG and DSO were performed by saving waveform data to a computer and postprocessing it offline using the same

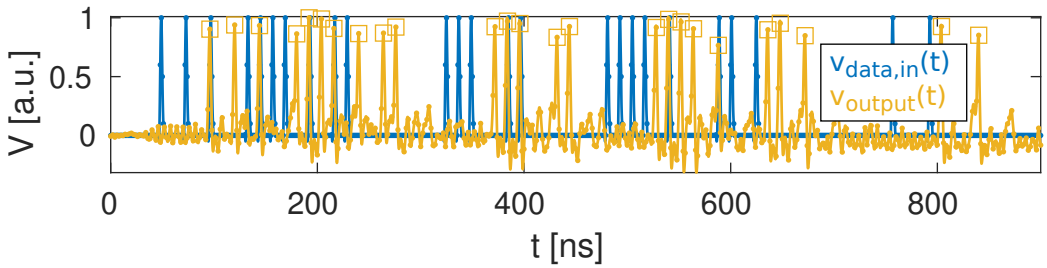

(a) A peak-finding algorithm detects switching events in the digitized output nTron voltage.

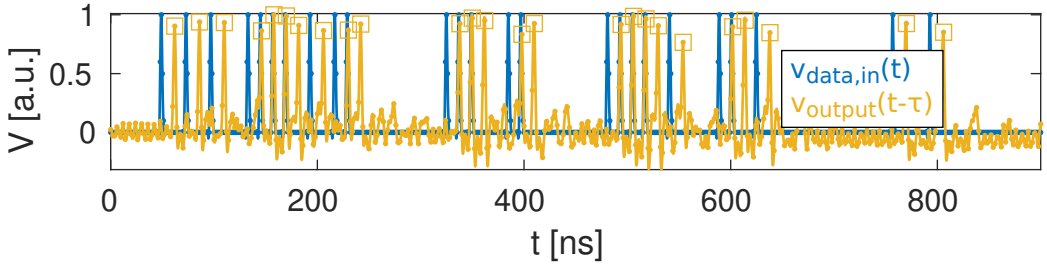

(b) The digitized waveform is deskewed by the cable delay measured with the cross-correlation method.

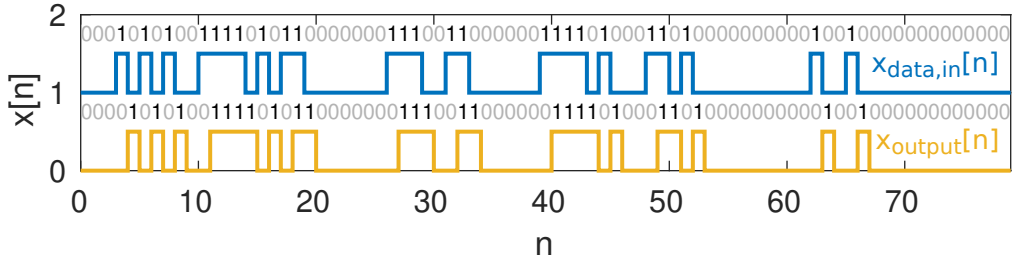

(c) The peaks in the shift register output are classified as a "1" or "0" for the time interval of each sample, and the output is compared with a digitized copy of the input signal.

Figure 4-10: Sample BER calculation. First, the peaks in the digitized waveform of the voltage across the output nTron were found. Then, the digitized waveform was deskewed by the cable delay. After digitization in (c), the output signal was shifted by one clock cycle (to account for the expected delay of the shift register) and XOR'd with the symbols transmitted by the AWG.

procedure, as opposed to calculating it in real time. Saving data and postprocessing was more time consuming, and took about 30 minutes to run, even for lower resolution sweeps such as the one in Figure 4-11a. Furthermore, the amount of storage space required for saving full-sample-rate digitized waveforms makes measurement of BERs substantially below 10^{-3} prohibitive; the raw data used to generate Figure 4-11a takes up over 10 GB uncompressed, whereas the figure itself can be rendered with a few kB of data once postprocessing is performed.

4.4 Bias margins

The *bias margins* of the shift register are defined as the ranges over which the input signal amplitudes can be varied without negatively impacting the BER of the shift register. The maximum and minimum *allowable amplitude* of the input signals define the upper and lower amplitudes at which the BER begins to degrade. A *bias point* is defined as the set of input signal amplitudes used on the shift register. A bias margin analysis was performed by sweeping the shift clock and readout clock amplitudes ($v_{\mathrm{clk,shift}}$ and $v_{\mathrm{clk,readout}}$) and measuring the BER of the shift register at each bias point. The other two free parameters (the input clock and input data amplitudes: $v_{\mathrm{clk,in}}$ and $v_{\mathrm{data,in}}$), were held fixed for each bias margin analysis of shift and readout clock amplitudes. A series of low-resolution bias margin analyses of shift and readout clocks were performed, iteratively changing the input clock and data amplitude between each analysis to rapidly find a local optimum where the width of the shift and readout clock bias margins were widest.

Using the same clock signal for many stages (like in Figure 3-10a) rather than a separate clock signal for each stage is necessary for scaling up to shift registers with more than one or two loops. However, since nanofabrication is not perfect, each stage will have slightly different optimal operating points and margins for input signal amplitudes. Furthermore, the bias resistors will not be identical, so each stage will receive a slightly different current. If we make reasonable assumptions about how the bias margins may vary from nTron to nTron and how resistor uniformity may vary

over a chip, the bias margin analysis of the two-loop shift register tells us a lot about whether the design will be scalable to the thousand- or even million-device level.

Figure 4-11 shows several plots of BER versus clock amplitude. The dark regions of the plots represent correct operation of the shift register. In 4-11a, the shift register was tested with 800 binary symbols at each bias point. In 4-11b, 100,000 symbols were sent to the shift register for each bias point, allowing measurement of BERs down to 10^{-5}. The bias margin analysis plotted in Figure 4-11c used 1,000 samples at each bias point, however the fine sweep over shift and readout clock amplitudes gives a much cleaner picture of the bias margin landscape.

For the third set of experiments, the 67 pF capacitors on the clock and data signals were replaced with 10 pF capacitors in an attempt to operate the device at higher frequencies. This caused a slight shift in the optimal bias regions between Figures 4-11 and 4-12. The shift happens because the 67 pF filter removed a substantial amount of energy from the pulses sent to the shift register that was no longer filtered by the 10 pF filter. Increasing the cutoff frequency of the filter reduced the attenuation of the high frequency content, meaning the AWG did not need to generate as strong of a pulse to switch the nTrons in the shift register.

As can be seen in Figure 4-12a, the bias margins have a characteristic shape. The optimal bias region (in black) slopes down and to the right in the lower half of the plot (*i.e.* for $i_{\text{readout,pp}} < 80\,\mu\text{A}$). This is due to the transfer characteristics of the readout nTron (which is a narrow-gate device). If a larger readout clock is used, it would generate more heat in the choke of the readout nTron, meaning less current through the channel of the readout nTron would be required to switch it. Since the amount of current flowing through the channel of the readout nTron depends on the shift clock current, a lower shift clock amplitude would be required if the readout clock amplitude is increased. This trend only holds for the range $30\,\mu\text{A} < i_{\text{readout,pp}} < 80\,\mu\text{A}$.

For readout clock amplitudes below 30 µA, the current through the gate of the readout nTron was insufficient to cause it to switch. Since the readout nTron is a narrow-gate device, the gate must switch and substantially suppress the switching current of the channel for the shift register to work properly. Simply increasing the

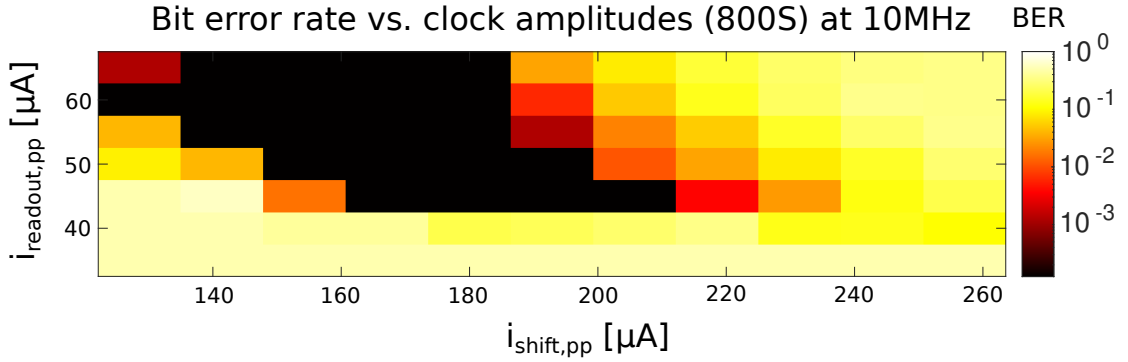

(a) Plot of bit error rate at various clock amplitudes. Bias margins were generated by postprocessing data measured with conventional AWG and DSO setup. The shift register was tested over a sequence of 800 samples. This plot took approximately 30 minutes to generate.

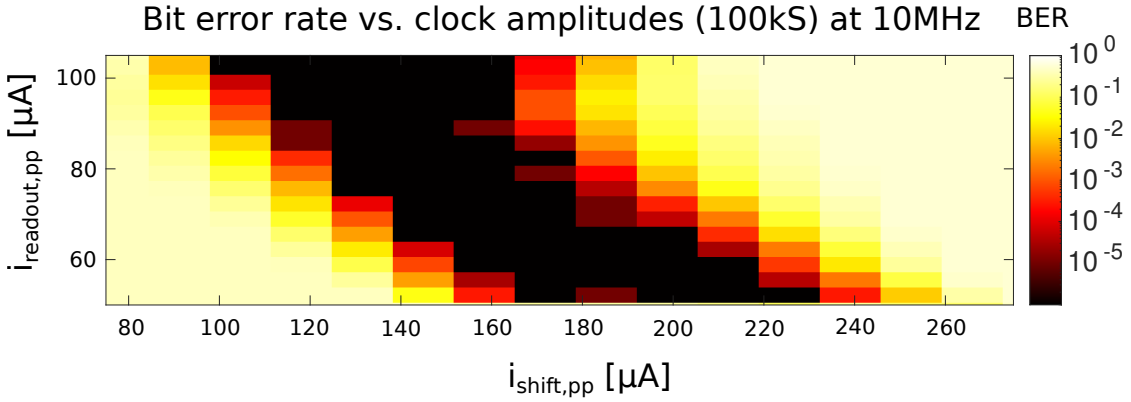

(b) High-bit-error-rate-resolution bias margin plot. The shift register was tested over a sequence of 100,000 samples. Bias margins were generated with PXIe controller with DAQ and AWG cards calculating the bit error rate in near-real-time. This plot took approximately 2 minutes to generate.

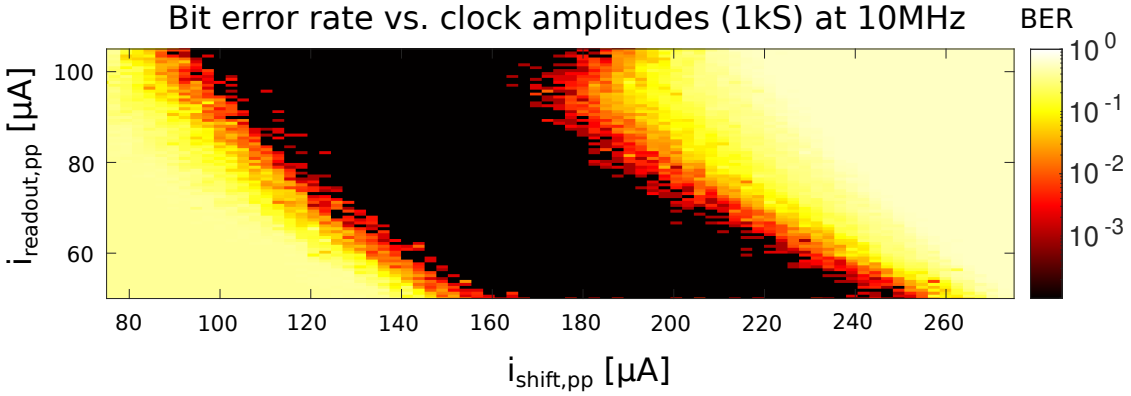

(c) High-bias-current-resolution bias margin plot tested over a sequence of 1000 samples. Bias margins were generated with PXIe controller with DAQ and AWG cards calculating the bit error rate in near-real-time. This plot took approximately 2 minutes to generate.

Figure 4-11: Bit error rate bias margins on various test setups. The automated setup that processes waveforms in real time is able to produce substantially higher resolution plots.

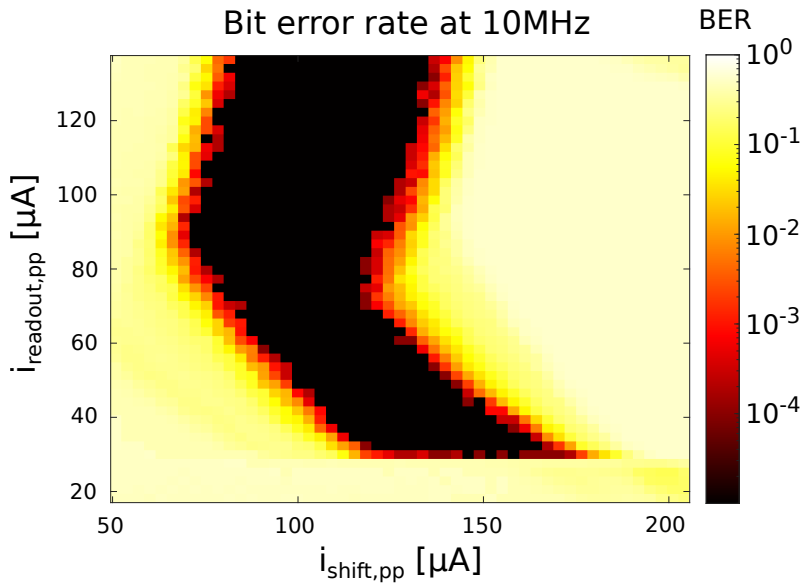

(a) Bit error rate at a clock frequency of 10 MHz. The maximum width of the shift clock bias margin is 46 µA (±24 %), and the readout clock bias margins have a width of over 102 µA (±59 %)

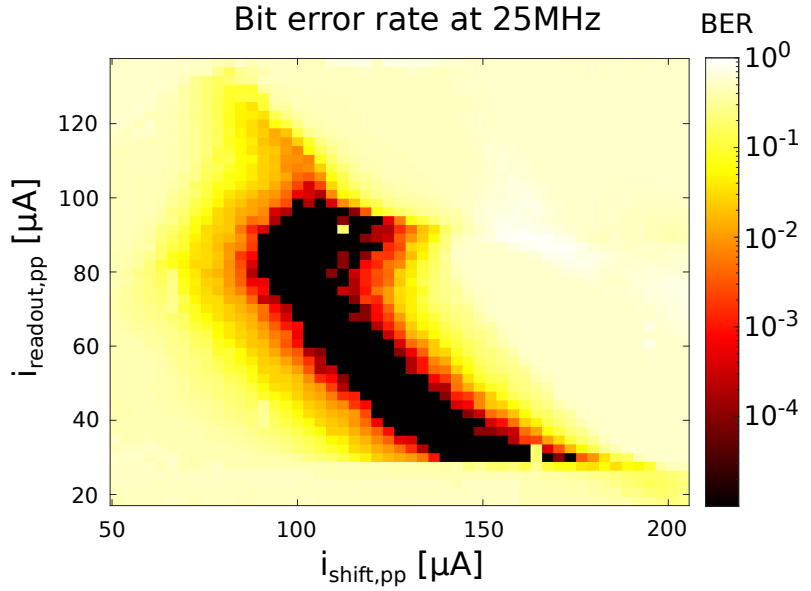

(b) Bit error rate at a clock frequency of 25 MHz. The width of the shift clock bias margin is 30 µA (±14 %), and the readout clock bias margin width is 35 µA (±23 %).

Figure 4-12: Bit error rate bias margins on the Keysight PXIe setup for 10 MHz and 25 MHz clock frequencies. The bias margins for both the readout and shift clocks are smaller at higher frequencies.

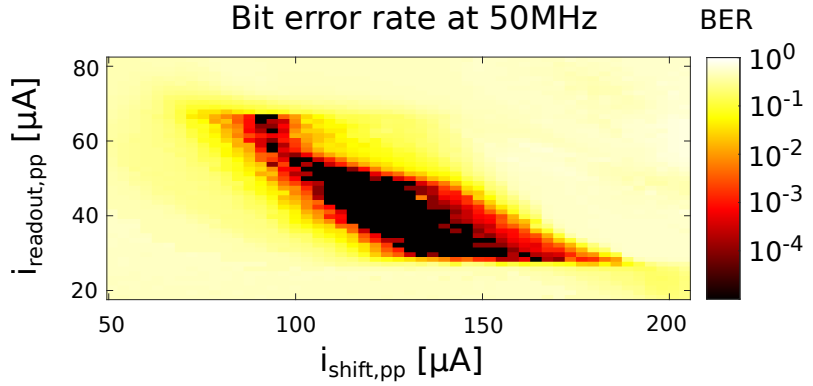

(c) Bit error rate at a clock frequency of 50 MHz. The width of the shift clock bias margin is 24 µA (±9.6 %), and the readout clock bias margin width is 15 µA (±20 %).

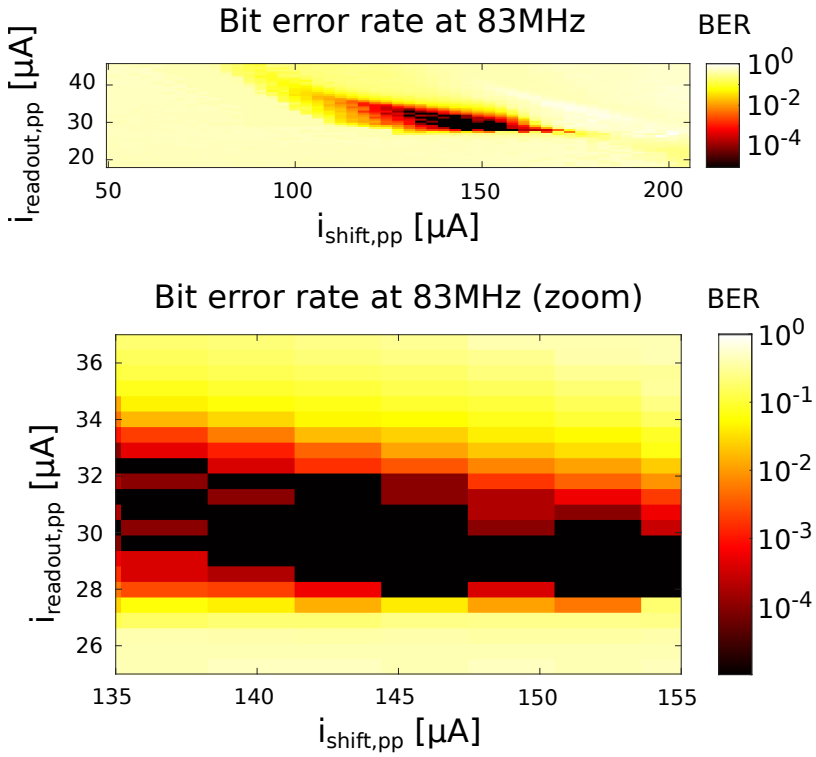

(d) Bit error rate at a clock frequency of 83 MHz. The width of the shift clock bias margin is 18 µA (±6.3 %), and the width of the readout clock bias margin is 3.3 µA (±5.4 %).

Figure 4-12: Bit error rate bias margins on the Keysight PXIe setup for 50 MHz and 83 MHz clock frequencies. The shift register barely works at 83 MHz, with shift clock bias margins of ±6.3 % and readout clock bias margins of ±5.4 %.

shift clock current further did not produce correct operation, since the shift clock would need to be roughly doubled to switch the readout nTron from current summation. This would violate the required current range for the middle nTron.

For readout clock amplitudes above 80 µA, something interesting happened: the optimal shift clock amplitude started to increase. This behavior was predicted in simulation as well, and mentioned briefly at the end of section 3.1.2. If the output clock amplitude is large enough, a non-negligible amount of current is injected backwards into the final loop of the shift register. The gate current has two paths it can take after entering the channel of the readout nanowire. It will split between the final loop of the nTron (since there is a path to ground back through the loop kinetic inductor and the middle nTron) and the source terminal of the readout nTron. If only the reactive impedance of each path is considered, the current should split in a ratio of roughly 1:100 (or less), where the majority of the current is dumped to ground through the source nanowire of the readout nTron. However, the hotspot growth in the channel will contribute a non-negligible amount of real impedance (on the order of 100 Ω) to both paths. Therefore, the current splitting will be closer to 50:50 than what would be predicted by the ratio of reactive impedances from the loop and source nanowire kinetic inductances. The voltage pulse has order of magnitude 400 MHz bandwidth so the total impedance of the two paths would be roughly $100 + 100j\Omega$ back through the loop and $100 + 1j\Omega$ to ground. We would therefore expect that roughly 1/3 of the current applied at the gate of the readout nTron would end up in the final loop, and the other 2/3 would flow to ground through the source terminal of the readout nTron. As observed in Figure 4-12a, the slope of the optimal bias region was about 1/3, which is consistent with the prediction of this hypothesis.

4.4.1 Bias margins as a function of frequency

As the shift register clock frequency increased, the widths of the bias margins decreased, as seen in Figure 4-12. The increase in the minimum shift clock amplitude may be explained by the L/R time constant of the shift register loops. In order to operate at a clock rate of $f = 1/T$, after a switching event, each loop current needs

to settle "close enough" to its final value within half of the clock period $T/2$, since a two-phase clock is used. Based on the margins in 4-12a, we can roughly quantify how many time constants constitutes "close enough". The minimum shift clock amplitude at which the shift register works properly is about 24 % below the optimal shift clock amplitude (if the center of the correct operating region is taken to be optimal). This means that the half-clock period $T/2$ needs to be at least 1.5 time-constants. Given that the time constant of the loop kinetic inductor and nTron shunt resistance is $L/R \approx 2.5\,\text{ns}$, we would expect a maximum clock rate on the order of 130 MHz.

However, this explanation does not account for the reduction in maximum allowable readout and shift clock amplitude. Taking a closer look at the different error types in Figure 4-13, we can see the reduction in maximum allowable clock amplitudes exhibit two types of behavior.

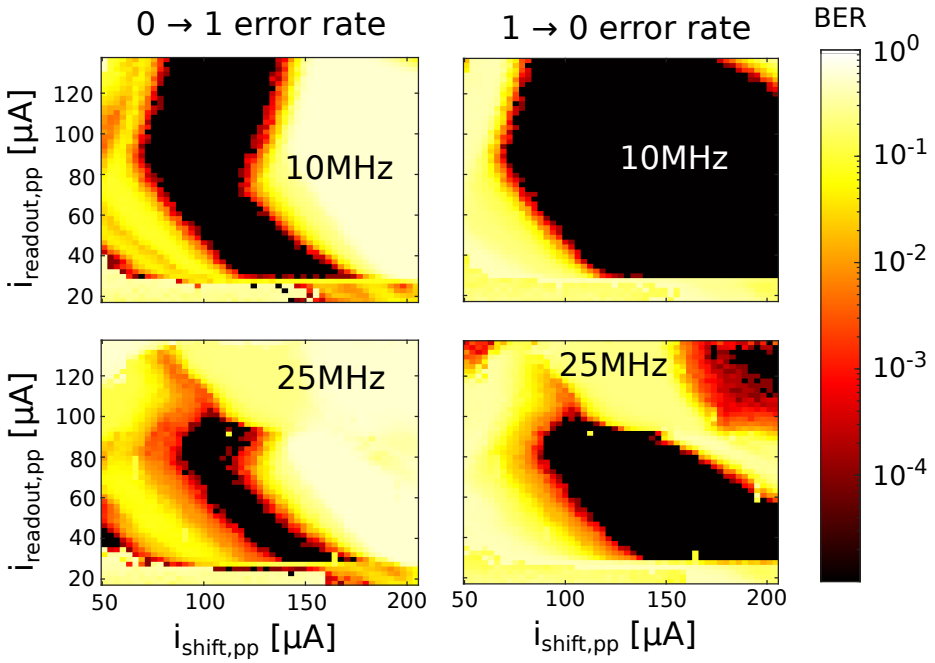

Figure 4-13: Bit error rate for 10 MHz and 25 MHz operation, separated by error type (both "1" → "0" and "0" → "1" errors). The rate of "1" → "0" errors at high clock amplitudes increases substantially when the clock frequency is increased to 25 MHz.

The first type of behavior observed is an increase in "0" → "1" errors in the absence of "1" → "0" errors. In this regime, the shift register output was the same sequence of symbols it was sent with the "0"s flipped to "1"s. This phenomenon presents as a

dark black region in the "1" → "0" error rate plot that overlaps with a white or yellow region in the "0" → "1" error rate plot. An increase in "0" → "1" errors with negligibly few "1" → "0" errors is the expected behavior when increasing the clock amplitude: for example, a large shift clock would be expected to switch the middle nTron even in the absence of a circulating current in the first loop of the shift register.

However, the second behavior observed for substantially larger readout and shift clocks (towards the upper right corner of the bias margin plots) is somewhat paradoxical: there is an increase in both "0" → "1" and "1" → "0" errors. That is, for high shift and readout clocks amplitudes, the shift register was more likely to output a "0" one clock cycle after receiving a "1" than for a slightly lower clock amplitude. This is paradoxical because we would expect that, for higher clock amplitudes, the nTrons would be more likely to just switch all the time, always outputting a "1", and increasing the clock amplitude would only increase the probability of outputting a "1". The most likely explanation for this type of error is that high clock amplitudes cause frequent misalignment of the output data in time due early switching of the readout nTron.When the shift clock signal is applied to the middle nTron, the readout nTron may still be cooling down from the previous clock phase, so a sufficiently large shift clock can switch both the middle nTron and readout nTron. If the shift register were fed a triplet of bits "010" and the shift clock signal was so large that it switched both the middle nTron and output nTron (*i.e.* the middle nTron switches first, and the diverted current is large enough that it switches the readout nTron as well), then the output will be delayed by only a half clock cycle instead of a full clock cycle. This would result in a "0" → "1" error (from the early switching of the readout nTron), followed by a "1" → "0" error (since the shift clock switched both the middle and readout nTrons, there would be no circulating current left). At higher frequencies, the readout nTron would have less time to cool down, which could explain the observed reduction in maximum allowable amplitudes for the readout and shift clocks.

4.4.2 Shift register sensitivity to magnetic fields

As discussed in 2.1, the ability of superconducting nanowire electronics to operate in a many fluxoid regime reduces sensitivity to external magnetic fields. This robustness to external fields is shown in Figure 4-14, where the shift register is demonstrated to operate in fields as high as 6 mT.

The shift register is designed to store a current of approximately 100 µA. Each shift register has a geometric self-inductance of roughly 4 nH, so when storing a logical "1", the loop contains about 200 single flux quanta (Φ_0). This makes the loop state (*i.e.* the presence or absence of current) relatively insensitive to fluctuations. For example, thermally activated phase slips that change the stored flux in each loop by Φ_0 would not impact the state stored in the shift register.

When a small field was applied (*e.g.* 1 mT), the bias margins shifted left or right (corresponding to a decrease or increase in the optimal shift clock amplitude), depending on the sign of the applied field. The overall shape of the bias margins remained the same. However, for higher fields, the maximum allowable shift clock amplitude decreased for both positive and negative fields, causing the width of the bias margins to shrink asymmetrically with applied field direction. This behavior is most likely due to a combination of screening currents formed by the Meissner effect and current crowding, the former dominating for small fields, and the latter giving rise to the asymmetry in the margin shapes at higher field strength. Note that due to the positive tone resist process, the shift register is completely surrounded by large planes of superconducting NbN. The etch pattern is topologically equivalent to a sheet with three holes in it, one for each loop of the shift register, and a third larger hole for the outline of the pads and input terminals as shown in Figure 4-15, so even when a switching event occurs in the shift register, there is not a break in the superconductor that would allow flux from an externally-applied magnetic field to penetrate any of the loops formed by the trenches around the shift register.

When an external magnetic field is applied, a screening current will form around the edges of the chip due to the Meissner effect. The area of the internal loop of

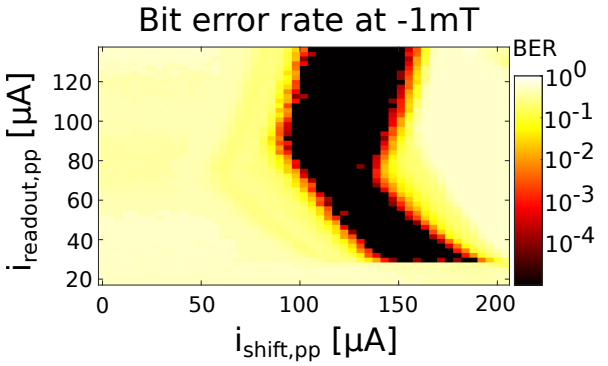

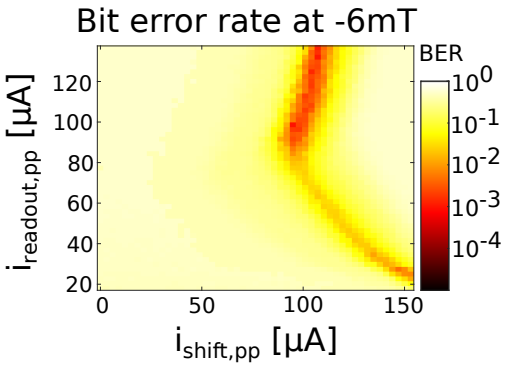

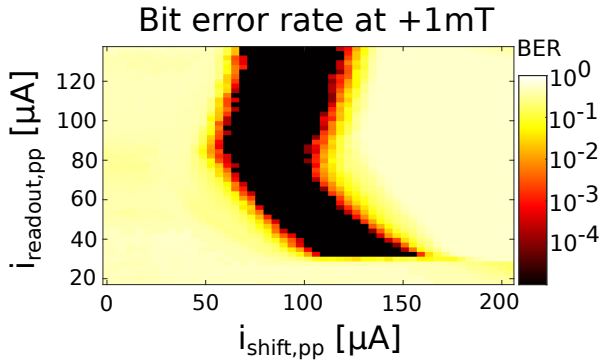

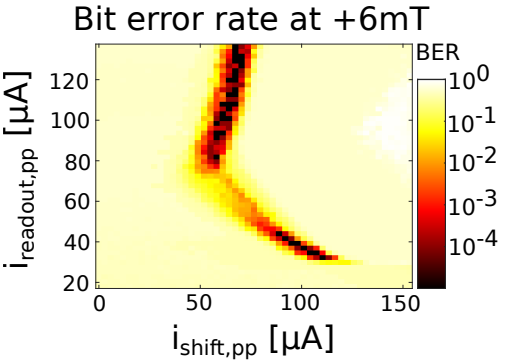

(a) Bias margins of the shift register under the application of a ±1 mT field. A positive field (into the plane of the chip) reduced the required shift clock amplitude, and a negative field (out of the plane) increased the required shift clock amplitude. The width of the shift clock bias margins decreased slightly from the ambient-field case to 41 µA for ±1 mT (±25 % for +1 mT and ±20 % for −1 mT).

(b) Bias margins of the shift register under a larger ±6 mT field. The −6 mT field prevented the shift register from operating with a BER below 10^{-3}, while there are still bias points at which the shift register could operate with a BER below 10^{-4} under the application of a +6 mT field. The shift clock bias margin is roughly 6 µA (±4 %) for the +6 mT field.

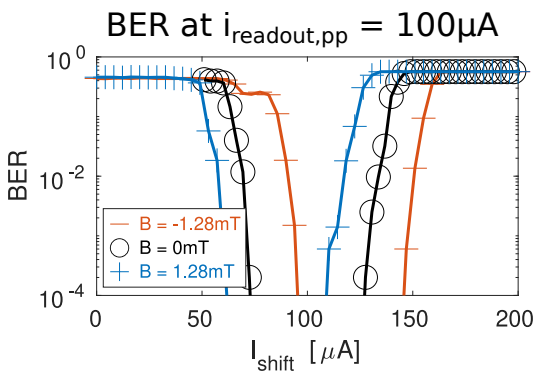

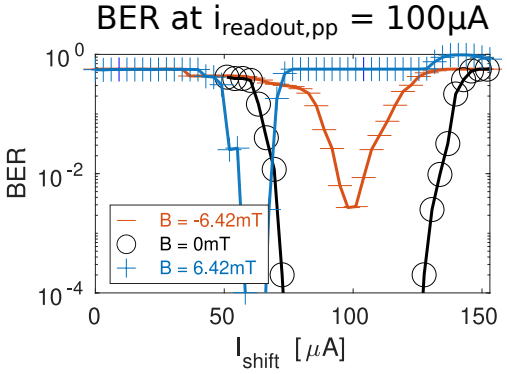

(c) Cross-section of bias margins under the application of a ±1 mT field. The cross-section is taken at a readout clock amplitude of 100 µA.

(d) Cross-section of bias margins under the application of a ±6 mT field. The change in shape of bias margins was asymmetric for positive and negative fields, in stark contrast with the results under a ±1 mT field.

Figure 4-14: Bias margins as a function of magnetic field for a 10 MHz clock frequency. The width of the bias margins remains roughly the same for small fields, but for large fields, the width of the bias margins drops drastically.

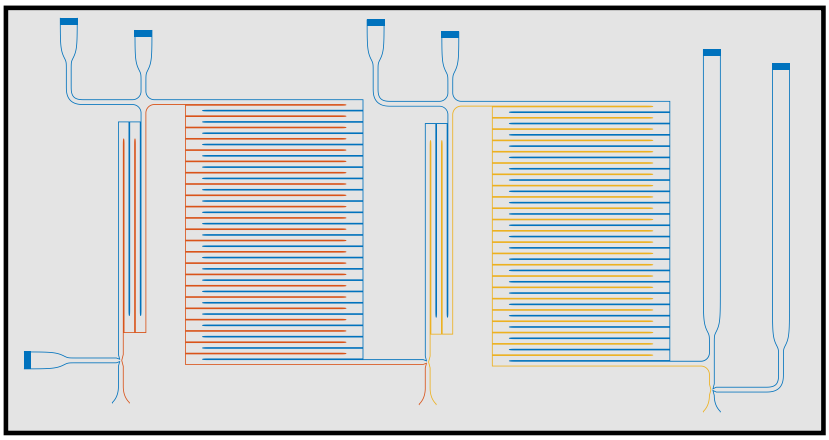

(a) Two loop shift register layout with contiguous etch slots color coded. The pads are not shown, but short trenches that connect the trench on either side of each wire that goes to the pads are added.

(b) Fully simplified topology of the slots in the NbN layer.

Figure 4-15: Topology of shift register etch pattern at varying degrees of simplification. Each etch region (first loop slot, second loop slot, pad and input/output wire slot) is color coded. (a) shows a minimally modified diagram of the layout used to pattern the NbN to make the shift register. As can be seen in the topologically equivalent, simplified etch pattern in (b), the NbN layer is fully connected around the shift register. Therefore, no external magnetic flux can penetrate the slots in the NbN around the shift register unless a large swath of NbN becomes resistive from the corner of the chip to an edge of one of the slots around the shift register.

the shift register was roughly $450\,\mu m^2$, so the induced current would have to screen roughly $220\Phi_0$ of flux per millitesla of external magnetic field (applied out-of-plane). Most of this screening should come from large currents that flow around the edges of the chip in the NbN ground plane, but some current would be induced on either edge of the nanowire that forms the shift register. This screening current would superimpose with the circulating currents and clock signals, and if large enough would perturb the operation of the shift register by either increasing or decreasing the total current density around the corners in the nTron constriction. The geometric self-inductance of the loop is approximately 4 nH, so if the field were screened entirely by the nanowire (*e.g.* in the case of no ground plane), then the screening current would be $110\,\mu A/mT$. This is unsurprisingly much larger than the observed shift in bias margins of $15\,\mu A/mT$. However, the screening current in the ground plane should cancel most of the applied field, leaving only a small screening current that actually

flows in the nanowire. Therefore, the only flux that threads the loops of the shift register is the flux generated by the states stored in the shift register.

The addition or subtraction of screening current from the circulating current explains the shape of the bias margins under the application of small magnetic fields. However, at higher fields, this explanation seems to break down, as can be seen in Figure 4-14d, where the upper limit for the shift clock *decreases* under the application of a negative field (whereas it increased under the application of a positive field). This is most likely due to a current crowding effect. The left side of the nTron is more prone to switching due to the two sharp corners on either side of where the gate nanowire connects to the channel, as compared to the more gradual corner on the right side of the channel. Under the application of a magnetic field, the screening current decreases the total current on one side of the nTron, it increases the total current on the other side of the nTron. Due to the direction of flow of the loop current and clock currents, the bottom-left corner of the nTron is most prone to switching, so for small magnetic fields, the effect of screening current mostly impacts the switching at this corner of the nTron. However, when applying a large enough magnetic field, the other corners begin to play a role. For the application of a negative fields, the screening current increases the total current passing around these corners, thus reducing the maximum allowable shift clock amplitude. When applying a positive field (into the plane of the device), the screening current increases the current density on the gate-source (lower left) corner of the nTron, which should decrease the maximum allowable shift clock amplitude.

4.4.3 Two-loop shift register bias margins as a tool for evaluating scalability

As we have seen in previous sections, the properties of the nTron play the dominant role in the bias margins of the shift register. Therefore, when evaluating the practicality of scaling this two-loop shift register (or a similar nTron-based circuit), we focus our attention on the reproducibility of nTrons with identical characteristics.

In a perfect world, if we fabricated a million nTrons with the same designed dimensions, each nTron would have identical switching currents and transfer characteristics. However, we can expect the realities of fabrication to limit the degree to which these devices are identical. To understand if it is possible to fabricate a working million-stage shift register, we need to understand the distribution of the bias margins of each nTron. Only the shift clock is relevant to this analysis of scalability, since only a single readout nTron is used in a shift register. Therefore, we can just consider the width and center of a slice of the bias margin plot at a fixed readout clock amplitude, similar to those plotted in 4-14c and 4-14d. We make a simplifying assumption that the shape of the region for which the shift clock is optimal will be identical for all nTrons, with shifted center depending on variations in the nTron width. We also assume the variations are independent and identically distributed, furthermore, that they are normally distributed. Based on these assumptions, the probability that for a collection of N nTrons with standard deviation in width σ_w, the minimum width and maximum width will be within W of each other is given by Equation 4.1:

$$\left(\int_{-W/2}^{W/2} \frac{1}{\sigma_w \sqrt{2\pi}} \exp\left(\frac{-x^2}{2\sigma_w^2} \right) dx \right)^N \tag{4.1}$$

A study of nTron choke-width reproducibility using a similar spin-on polymer negative tone resist demonstrated a gate width standard deviation of 2.4 nm for 30 nm-wide chokes [40]. Considering operation at 83 MHz, the width of the optimal shift clock region is 18 µA, so the maximum allowable spread in widths for the nTrons in a shift register is 24 nm (assuming a critical current density of 46 GA/m^2 and nanowire thickness of 16 nm).

Therefore, if we consider a shift register with a million stages and nTrons with widths of (270.0 ± 2.4) nm, then we would expect 56 % of the fabricated shift registers would operate at 83 MHz. If the operating frequency were decreased to 50 MHz, then even shift registers with a billion stages would be expected to operate, with an expected yield of 99 %. Realistically, this is an optimistic upper bound for the yield based on the assumptions that were made, since this work used a different resist

from [40], and more crucially, the shape and distribution of optimal bias margins of the nTrons in a circuit may vary, which will alter the shape and distribution of the optimal bias margin centers and widths. Reasons for these variations could include fabrication defects or geometry variations such as sharpness of the nTron corners, or when considering circuits with billions of nTrons, uniformity of deposited films over a large area. However, this analysis provides the impetus for developing larger scale nTron circuits, since without building and testing large scale circuits, any second-order effects that would reduce yield are difficult to characterize.

4.5 Energy analysis

Due to the low operating current of the shift register, the energy consumption is estimated to be roughly 80 fJ/bit. This energy is dominated by the clocking, which dissipates 100 µA through 1.96 kΩ for 2 ns twice per cycle of the clock. The intrinsic energy of information stored in each loop of the shift register is much lower: roughly 260 aJ/bit (100 µA in a 52 nH loop). The total energy dissipation including clocking limits operation of kilopixel arrays to roughly 50 MHz (for a frame rate of 50 kfps) if housed in a lightweight cryocooler [41]. However, reduction of the clock impedance and moderate scaling of the shift register are expected to improve power consumption by several orders of magnitude, allowing for larger, faster imagers with similar power dissipation.

4.6 Scaling and improvements

The experimental shift register demonstrates very broad bias margins, suggesting it is readily scalable to millions of stages. Furthermore, the ability to withstand external magnetic fields up to 6 mT out-of-plane would allow a shift register to operate in harsh environments unshielded, which gives it a leg up against other superconducting technologies that operate with JJs. However, the energy consumption must be decreased and speed increased if the shift register is to be used in megapixel SNSPD

array readout.

The operating frequency can be improved by decreasing the size of the loop inductors and nTrons. Smaller loop inductors will have a lower electrical L/R time constant, and smaller nTrons can operate with lower currents, reducing the amount of excess heat after a switching event. At high frequencies, the lower electrical time constant is expected to reduce the minimum allowable shift and readout clock amplitudes at high frequencies, and the lower operating current is expected to increase the allowable maximum clock amplitudes.

Furthermore, decreasing the size of the shift register, particularly the loop area will reduce its sensitivity to external magnetic fields. The loop-based logic gates in [1] use smaller loop areas than this work, and exhibit improved bias margins over the shift register.

In addition to decreasing the operating current of the shift register by reducing the size of the nTrons, decreasing the clock bias resistor will enable lower energy dissipation. The operating current has a more drastic effect due to the quadratic dependence. By reducing the operating current by a factor of 10 and clock impedance by a factor of 20, megapixel arrays operating at clock rates up to 100 MHz (frame rate of 100 fps when read out over a single wire) are feasible from a power perspective, consuming 4 mW.

Aside from these intrinsic challenges that must be addressed for circuits with thousands to millions of nTrons to become reality, the ability to effectively instrument and characterize these electronics poses a big hurdle for scaling up. As discussed in section 4.3, the amount of data required to store and postprocess raw waveforms makes bias margin analysis tedious and time-consuming. Preprocessing the data with math and measurement features common to most DSOs is a must; however, the variety and flexibility of such functions for performing more than the simplest analysis of raw waveform data is either impractical or impossible, still leaving a huge amount of data that is slow to transfer from the DSO to computer and time-consuming to postprocess. An automated setup that is able to characterize circuits in real time or faster would enable measurements of BER below 10^{-6}, and would allow for rapid

exploration of the behavior of nTron circuits in a variety of environments and test conditions.

Chapter 5

Nanocryotron shift register applications and outlook

This part of the thesis has presented an advancement in superconducting logic that may help enable the development of larger scale superconducting circuits and high count rate photon imagers. The experimental demonstration of synchronous state transfer in the fabricated two-loop shift register opens the door for serialization and deserialization of digital data. The two most pertinent applications for this superconducting shift register are digital readout of high count rate imagers and low temperature electronics for testing of large scale superconducting circuits.

The detailed discussion of the approach towards designing such a circuit as well as specific design decisions should provide a starting point for future development of other nTron-based circuits. In addition, although presented under the lens of the shift register, the interpretation of the results may be useful for circuit designers in the future. Specifically, hypotheses presented about thermal relaxation of the readout nTron, loop current injection, and magnetic field sensitivity are all relevant to any superconducting circuit based on circulating currents, such as [1, 27]. Further studies should be performed to understand the limits of the hypothesized effects. To address the hypothesis on thermal relaxation, a two pulse test with variable inter-pulse delay (similar to pump probe measurements in ultrafast optics) could be used to probe the temporal dependence of the channel switching current suppression in nTrons when

driven with large gate currents. Such a study could help understand the thermal speed limit of nTron circuits. A negative-tone process could be used to fabricate an identical shift register without large NbN ground planes to study the effect of screening currents in the ground plane on shielding the circuit from external magnetic fields.

Although a shift register would serve as a useful tool for assisting with verification of large scale superconducting circuits, it is only part of a more complex system that would be required for full built-in self-tests of superconducting logic. In the meantime, room temperature test equipment is needed to help push nTron and superconducting nanowire circuits from the few-nTron-per-circuit to dozen- or even hundred-nTron-per-circuit regime.

Part II

Design and characterization of an analog frontend for nanowire characterization and testing

Chapter 6

Design of a room-temperature low-noise analog frontend for interfacing with superconducting nanowire electronics

Measuring, characterizing, and debugging superconducting nanowire circuits requires specialized test equipment to be done efficiently. Conventional techniques using AWGs and DSOs that just barely work for testing circuits with a small number of nanowires will quickly become prohibitive from a cost and time perspective when scaling up to circuits with hundreds of nanowires. Increasing the complexity and scale of nTron circuits would be quite straightforward if measuring experimental circuits were just like simulation and we could probe device currents and hotspot resistances to understand why our circuit was producing a particular output. However, we can only estimate these quantities indirectly by looking at how the voltages at various nodes in the circuit evolve over time. Watching for the characteristic voltage spikes that we saw in simulation (Figures 3-10b and 3-11b) and experiment (Figures 4-7 and 4-8) tells us when a nanowire switches from superconductive to resistive and back again. If we have high confidence that the circuit was in a particular state be-

fore the switching event, then we can know with pretty good certainty what the state is after the switching event. Subsequent switching events can be used to re-evaluate and update our assumptions about the state of the circuit (and also inform us of any errors that may have occurred in the circuit).

For a circuit that is designed to be fast, these voltage pulses will be on the order of a few nanoseconds, so a high bandwidth (*i.e.* >1 GHz) digitizer is needed. Furthermore, as discussed at the end of section 4.1 and demonstrated in Figures 4-7 and 4-8, the inductance of a resistively-shunted nanowire can cause the formation of additional voltage spikes when a fast-edge clock signal is applied to the nanowire. In the case of the shift register, the additional inductance was an erroneous design decision, but it is plausible that there will be circuits that must be designed this way. These spikes have a characteristic shape and symmetry (a negative pulse always accompanies a positive pulse) that is different from that of a nanowire switching (which is always a single negative or positive pulse, depending on the original direction of current flow through the wire). The digitizer must have sufficient voltage resolution to produce waveforms that enable a human or computer to distinguish between a nanowire switching event and the voltage spikes caused by the charging/discharging of the nanowire kinetic inductance.

Watching these pulses go by on a DSO is straightforward enough for a small circuit, where it's feasible to hold a mental model of the circuit in our heads and the number of node voltages to measure is within the number of channels on the DSO. However, as circuit complexity scales, a custom data acquisition system becomes attractive to deal with increased channel count and to allow for automation of circuit state tracking and error calculation. In the world of quantum computing, the Xilinx RFSoC platform is attractive, and has been demonstrated as an effective tool for measurement and control of superconducting quantum systems [42, 43, 44]. Although intended for gate-based quantum computing, which uses substantially different signaling schemes as compared to classical electronics based on superconducting nanowires, the overall design goals of low noise, high bandwidth, and powerful real-time processing make the work of [42] a good starting point for the development of custom test equip-

ment for superconducting nanowire electronics. The RFSoC platform is most notable for its high bandwidth dataconverters: radio-frequency analog-to-digital converters (RFADCs) and radio-frequency digital-to-analog converters (RFDACs) which have analog bandwidths above their Nyquist bandwidth, allowing them to perform direct downconversion of bandlimited radio-frequency signals without an analog mixer [45]. This is not useful for testing nTron circuits, since all signals will be baseband, but it will be useful for testing microwave superconducting circuits with single-frequency tones. The programmability of the FPGA is key to making a fast and flexible system for automating the testing of superconducting nanowire circuits.

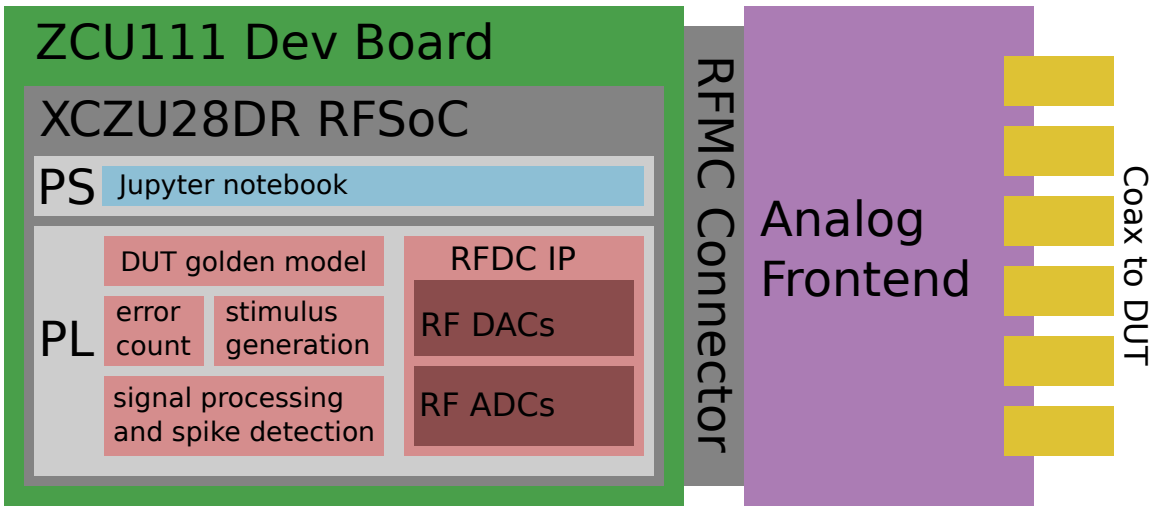

Figure 6-1: Proposed configuration for RFSoC to test nTron logic. The programmable logic includes signal generation to create the stimulus for the nTron logic, as well as signal processing to detect switching events in the device under test (DUT) and compare them with a golden model of the expected device behavior. A Jupyter notebook hosted on the multicore ARM CPU orchestrates the test and provides an easy-to-use interface for users working in the lab.

However, the RFSoC platform comes with two key challenges that must be addressed. For one, the firmware needs to be written to configure the programmable logic in such a way that makes it a useful tool for any scientist to characterize circuits made of superconducting nanowires. Firmware is not the main focus of this thesis, but some discussion of next steps for firmware development can be found in chapter 8. This thesis primarily discusses the second concern: the dataconverters require an

analog frontend to interface with the superconducting nanowires. The analog frontend's primary purpose is to perform amplification, but it also acts as a buffer to protect the expensive RFSoC from accidental damage due to improper termination of digital-to-analog converter (DAC) outputs or electrostatic discharge.

Due to the flexibility of the RFSoC as a test instrument, in addition to testing nTron circuits, we also aim to use the RFSoC for testing other superconducting electronics which require high-bandwidth, low-noise, automated testing infrastructure. For example, the RFSoC has been used to perform readout of frequency multiplexed arrays [46]. For this reason, the goal of the analog frontend design is to be applicable for a variety of applications and not just for testing digital electronics made with superconducting nanowires. Keeping the considerations for these other use cases in mind, we focus our attention primarily on the use case of nTron circuit testing and characterization, while making design decisions that avoid handicapping performance in these other use cases. A proposed diagram for a setup that uses the RFSoC to test nTron circuits is shown in Figure 6-1. The following chapters will focus on the design and characterization of an analog frontend designed for this purpose.

The process for designing and building the analog frontend can be broken down into three steps: architectural design and planning, component selection and schematic capture, and PCB layout. While they are listed as separate steps, the process does not run once through in this exact order, there are many loops as different designs are re-evaluated, and it is not uncommon for the first two steps to have significant overlap. The goal of this chapter is to not only explain the specific design decisions that went into making the analog frontend, but also highlight the general approach for doing this sort of design work.

6.1 ADC buffer design

Interfacing superconducting nanowire circuits with an RFADC requires an analog circuit that can perform two tasks: signal amplification (with variable gain) and single-ended to differential conversion. The stipulation for variable gain is required

to ensure applicability for testing a variety of nanowire electronics which may have drastically different output signal levels. Single-ended to differential conversion can be performed quite well with a balun. However, high voltage gains are not achievable with a balun in this specific application, since baluns amplify voltage by transforming from a low impedance to high impedance. In this case, the impedance ratio is fixed at 2, giving a fixed voltage gain of 4 if using a balun. Furthermore, the ability to adjust the gain is not available with a balun. For these reasons, we focus our attention to active electronics, basing the design heavily off of the work in [42]. This work uses the same amplifier chain, with an additional preamplification stage using an inexpensive LNA with lower noise figure and higher gain than the input amplifier used in [42].

6.1.1 Choke selection for LNA bias tees

The bias tee of an LNA is used to power the LNA with a voltage source without shorting its output to small-signal ground. Commercial LNAs sometimes have integrated bias tees, particularly in models packaged for benchtop use. However, surface-mount chip solutions often will leave bias tee implementation up to the circuit designer for cost and performance reasons. For one, on-chip inductors are typically lower performance than discrete off-chip inductors (not to mention significantly more costly due to the use of expensive chip area for large metal coils). Second, this allows the circuit designer to optimize the performance of the LNA for their application while saving cost. The simplest implementation of a bias tee uses a single inductor and capacitor, with the inductor separating the DC power supply from the LNA output, and the capacitor removing the DC offset from the LNA output signal [47]. An equivalent circuit model of an LNA output is shown in Figure 6-2.

One common misconception when specifying an RF choke for a high frequency application is that the self-resonant frequency of the choke must be well below the desired operating frequency of the circuit. Generally, larger value inductors require more turns and a larger package, increasing the effect of parasitic capacitance and lowering the self-resonant frequency of the inductor. Below the self-resonant frequency, the impedance of the inductor is mostly inductive (positive-imaginary), and

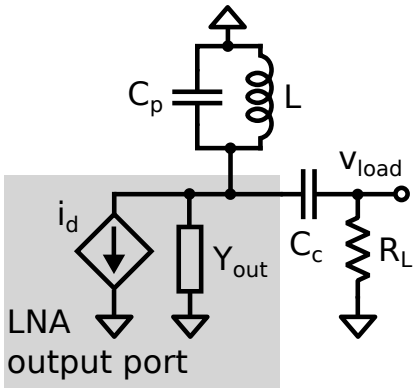

Figure 6-2: Small signal model of a low noise amplifier and bias tee. Parasitics and second order effects in the amplifier are not considered. Only the parasitic capacitance of the choke inductor is modeled; the DC resistance and magnetic losses are ignored.

above the self-resonant frequency, it's capacitive (negative imaginary). On resonance, the impedance is purely real and is quite large (often in the range of 1 kΩ-10 kΩ for µH inductors). For a tuning inductor, the sign of the impedance matters, as it is often being used to cancel some opposite reactive component to perform impedance matching or to build a resonant tank. In these cases, it is still possible to use inductors near their self-resonant frequency, but care must be taken to properly model the parasitics so that their effects are properly included in any final model of the circuit. For the application of a choke inductors, however, only the magnitude of the inductor impedance matters, and the bigger the better, so using inductors well above their self-resonant frequency is perfectly acceptable.

Figure 6-3 shows the frequency response of a variety of inductors that could be used for LNA chokes. The power delivered to the load (normalized to the power delivered if the reactive components are ignored) is calculated from equation 6.2, based on the equivalent small signal model from Figure 6-2:

$$H(s) = -\frac{sC_cR_L}{1+sC_cR_L}\left(\frac{1}{Y_{out}+sC_p+\frac{1}{sL}+\frac{sC_c}{sC_cR_L+1}}\right) \tag{6.1}$$

$$P_{L,n}(s) = \left|\frac{H(s)}{R_L}\right|^2 \frac{1}{1+R_LY_{out}}^2 \tag{6.2}$$

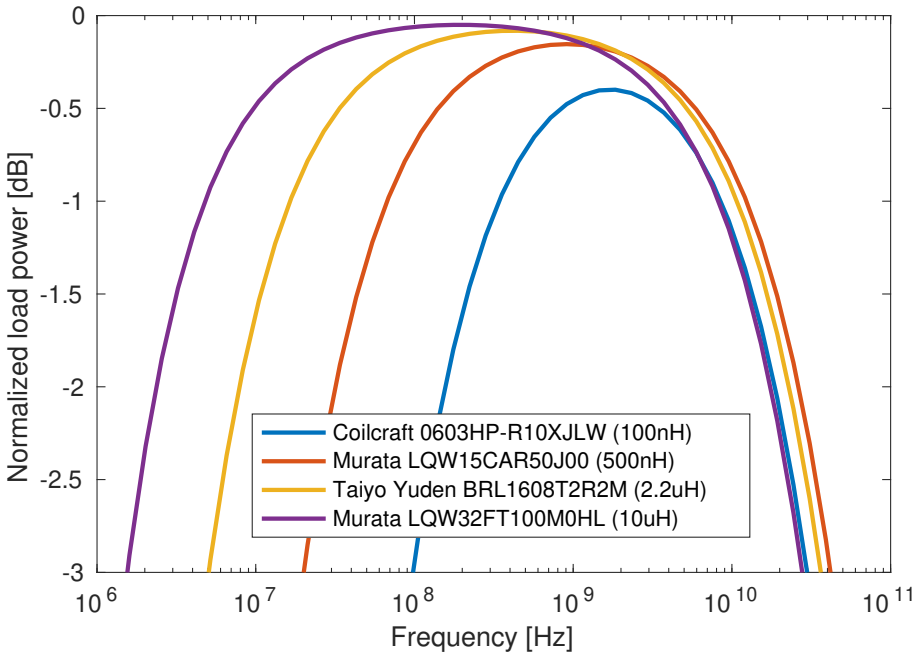

Figure 6-3: Power delivered to 50 Ω load for a matched LNA as a function of frequency for various choke selections

The output impedance of the amplifier $1/Y_{out}$ is assumed to be 50 Ω, since the amplifier is usually designed with this in mind. For the amplifier used in the ADC buffer, the magnitude of the input and output impedances vary between roughly 25 Ω and 150 Ω over the operating range of the device. This does not have a large effect on the shape of the frequency response in Figure 6-3; it mostly just impacts the actual power delivered to the load due to impedance mismatch.

It is interesting to note that despite the 100 nH choke having the highest self-resonant frequency (beating out some inductors by more than a decade in frequency), it has one of the lowest high frequency 3 dB cutoffs, due to its significantly lower inductance than the other inductors.

6.1.2 Analog fronted ADC buffer

Figure 6-5 shows the ADC buffer circuit. Amplification is first performed with a LNA (Texas Instruments TRF37D73), which is biased with a 10 µH choke inductor (Murata LQW32FT100M0HL). The AC-coupling capacitors C_c are realized with reverse-aspect-ratio capacitors (Murata LLL185R71E103MA01L), which have better

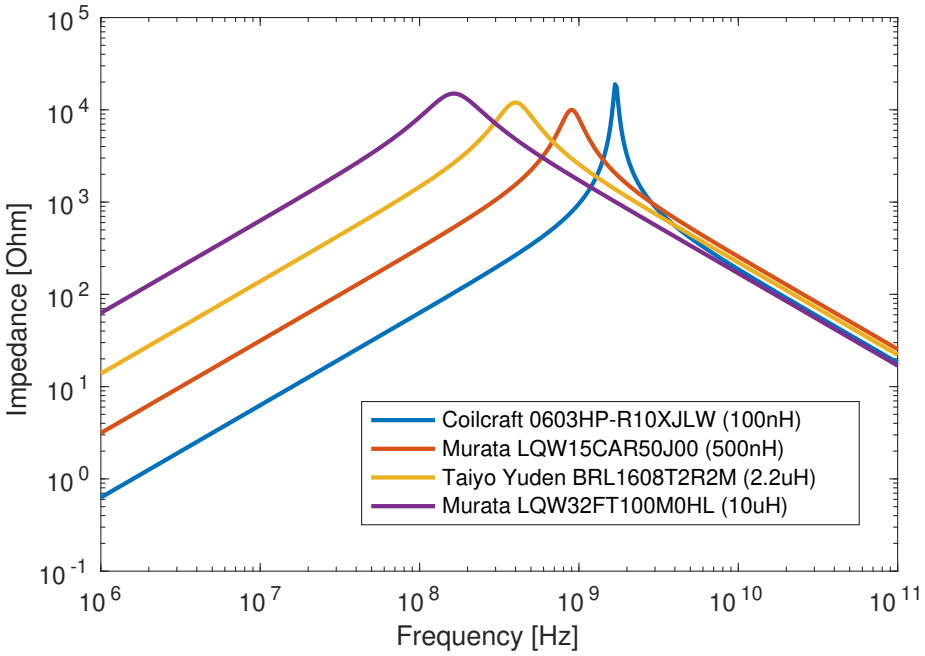

Figure 6-4: Inductor impedance as a function of frequency for various inductors. Larger inductors, despite having a much lower self-resonant frequency have similar high frequency impedance to smaller inductors. Their larger impedance at low frequencies therefore makes them better for broadband choke applications since they have high impedance over a wider bandwidth.

bandwidth than equivalent-area conventional capacitors due to the decreased effective series inductance from the wide, short package. Single-ended to differential conversion is performed by a fully-differential amplifier (FDA) (Texas Instruments LMH5401). For a discussion on selection of R_{G1}, R_{G2}, R_T, and R_{FB} for single-ended to differential conversion, see [48]. The output of the FDA is fed through a variable-gain amplifier (VGA) (Texas Instruments LMH6401) to adjust the signal level and differential low-pass filter (Minicircuits DLFCV-1750+) to prevent aliasing of high frequency signals during digitization by the RFADC. The gain of the VGA can be set in 1 dB increments with serial peripheral interface (SPI).

Bypass capacitors in 6-5 are shown with just a single capacitor, but are implemented with multiple discrete capacitors with different values and package sizes, as well as distributed power planes which use the PCB dielectric to form large parallel plate capacitors with very high self-resonant frequencies for high quality filtering of power supply noise. This is a common strategy in PCB design, and used for all

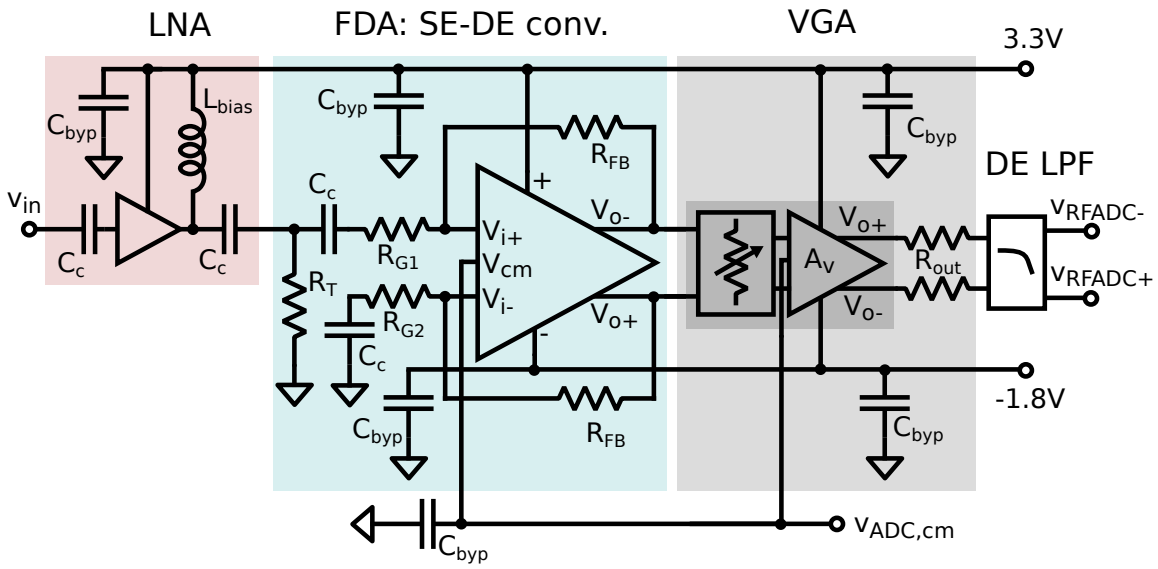

Figure 6-5: ADC buffer circuit. The single-ended input signal is amplified with an LNA and converted to a differential signal with an FDA. The amplitude of the signal is adjusted with a VGA, and filtered with a differential lowpass filter.

bypass/decoupling capacitances on the power planes in this board.

The split power supply for the FDA and VGA is chosen to maximize the output swing around the RFADC common-mode voltage of 1.2 V, improving linearity of the ADC buffer.

6.2 DAC buffer design

The DAC buffer serves two main purposes. One, it shifts the common-mode of the RFDAC output signal to be centered about 0 V. Second, it provides a known impedance termination to the RFDACs, preventing accidental damage to the RFSoC due to improper termination. The output impedance of the DAC buffer is designed to tolerate arbitrary load impedances without any risk of damage.

Figure 6-6 shows a schematic of the DAC buffer circuit. The DAC buffer consists of two main parts: a very high bandwidth FDA which is used to amplify the signal from the RFDAC and shift its common mode, and a lower bandwidth, very-low-offset opamp to correct the output common-mode voltage (which translates to the offset voltage of the signal v_{out}) to 0 V. Only a single bypass capacitor is shown for each rail

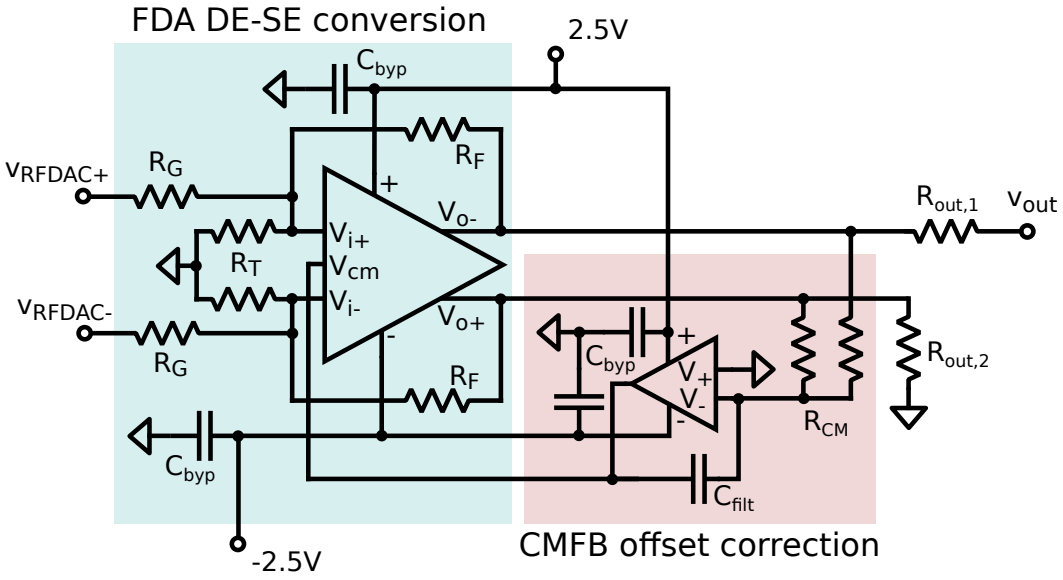

Figure 6-6: DAC buffer circuit. The differential RFDAC output is amplified and common-mode shifted by the FDA. A low-offset opamp is used to correct the common-mode offset of the FDA so that the signal v_{out} is centered around 0 V.

for each amplifier, however multiple capacitor values with different package sizes are used to optimize the performance of the bypass circuit. The resistors R_T at the input of the FDA perform common-mode shifting to bring the RFDAC output common-mode within the specified range of the FDA to reduce nonlinearity. Note that due to the negative feedback configuration of the FDA, the input nodes v_{i+} and v_{i-} are held at small-signal ground for the differential mode, so the differential mode is unaffected by the choice of R_T.

An opamp can perform both of these functions, however the bandwidth requirement forces us to use a FDA due to the lack of commercially available opamps with bandwidths above 2 GHz (some decompensated opamps have gain-bandwidth products as high as 8 GHz, but only have stable bandwidths of 1 GHz). Due to the differential output of FDAs, the output common-mode must be set through a common-mode feedback (CMFB) circuit [49]. FDAs include an integrated CMFB circuit to ensure maximum output swing, but for cost and power reasons, the offset of the CMFB error amplifier is usually 10 mV or more. The main use case of FDA is for driving differential electronics, such as a differential-input ADC, in which case, such a small common mode offset is perfectly acceptable. However, since we are only using one of

the outputs of the FDA, the offset matters a lot more here. Fortunately, an external, low-offset opamp can be used to correct the output offset to well below 1 mV as shown in Figure 6-6. The implementation of the external CMFB circuit follows the guidelines presented in [48]. The opamp is configured as an inverting low-pass filter with C_{filt} and R_{CM} to minimize common-mode noise. The cutoff frequency is 6 kHz.

In order to impedance match the DAC buffer output with 50 Ω and ensure an output common-mode voltage of 0 V, the FDA must be terminated with a 200 Ω load. As designed, the DAC buffer is impedance matched to a single-ended 50 Ω load and supplies a 0 V common-mode offset to that load. However, this limits the gain and maximum output power of the buffer circuit; the output matching network results in a −12 dB penalty.

6.3 PCB waveguide design

Routing PCB traces for the ADC and DAC buffer requires care to avoid performance degradation from distributed effects. For electrical signals with bandwidths above 100 MHz, the wavelength of the signal becomes comparable to the size of the PCB used to integrate the electronics that generate, measure, and condition these signals. Above 1 GHz, the wavelength in a typical PCB waveguide structure can be less than 10 cm, so any traces longer than roughly 1 cm need to be treated as distributed elements instead of a lumped wire.

For a multilayer circuit board, several options are available for waveguides. The simplest options for single-ended signals are stripline and microstrip waveguides, each of which has advantages and disadvantages depending on the bandwidth of the signals. Figure 6-7 shows two single-ended microstrip waveguides as well as single-ended and differential stripline waveguides.

One advantage of stripline waveguides is increased isolation from external electromagnetic fields which may introduce noise. This is less of a concern for tightly-coupled differential signals, since the environmental radiated noise is expected to mostly impact the common-mode signal, so the differential mode will remain relatively clean.

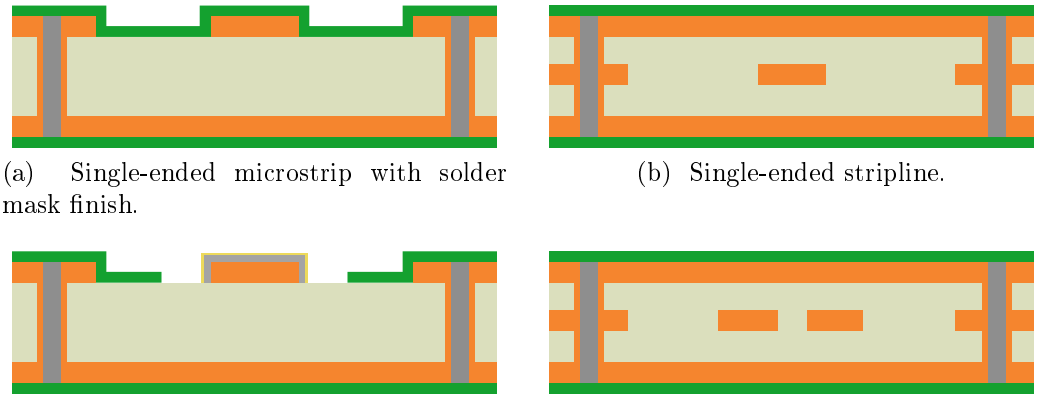

(a) Single-ended microstrip with solder mask finish.

(b) Single-ended stripline.

(c) Single-ended microstrip with ENIG finish.

(d) Differential stripline.

Figure 6-7: PCB cross-sections of microstrip and stripline waveguides with via-stitched guard traces. The mode propagates in the dielectric between copper layers. The solder mask polymer is shown in green, copper foil is shown in orange, and the prepreg/core laminates are light gray. The vias are clad in copper, and their drilled out centers are shown in gray. In order to avoid perturbing the mode impedance, the guard traces are placed far from the waveguide.

Because a microstrip is on the outer surface of the circuit board, it must either be covered in solder mask or be capped in a protective surface finish material, such as electroless nickel immersion gold (ENIG), to protect the bare copper from reacting with the environment and degrading over time. At high frequencies, solder mask is often dispreffered over finished copper due to increased dielectric losses and poor uniformity of the solder mask (which leads to non-uniform impedance over the length of a waveguide). Due to the skin effect, at high frequencies the majority of current flows at or near the surface of a conductor, increasing conduction losses in ENIG-plated waveguides due to the poor conductivity of nickel as compared with copper.

For these reasons, it is sometimes preferred to use a stripline waveguide for very high frequency signals, since the internal metal conductors do not need plating, so the only losses are conductor losses due to surface roughness of the copper foil on internal layers and the dielectric losses from the core and prepreg laminates that separate the copper layers.

However, below 10 GHz, the difference in loss due to surface finish is negligible for cm-long traces, unless extremely low insertion loss below 1 dB is desired [50].

Furthermore, in order to connect to the components on the surface of a PCB, an stripline waveguide requires interlayer transitions. It is possible to create vias with a variety of characteristic impedances (including 50 Ω), however most inexpensive multilayer PCB fabrication processes perform just a single drill step after laminating the layers together, thus the via would extend through all layers, creating a small capacitive stub. If the reactance of the stub impacts RF performance, then a higher cost high density interconnect (HDI) process which performs intermediate drill step(s) on the inner layers before lamination of the outermost layer(s) can be used, allowing stubless transitions between internal and outer layers. For these reasons, a microstrip waveguide may be a more attractive choice for signal bandwidths below a few GHz.

Coupled differential pair routing is desirable for differential signals because crosstalk reduction and susceptibility to radiated emissions can be improved by routing the traces adjacent to each other. However, due to the increased coupling between very close traces, the differential mode solution has a lower impedance than the combined impedance of each trace on its own. This can present a challenge for tightly coupled traces (*.e.g* in a broadside coupled stripline configuration), since the trace width has to be reduced to compensate for the coupling-induced reduction in odd mode impedance. Narrower trace widths are more sensitive to fabrication variation and can increase manufacturing costs of the PCB. In general, this is not an issue for edge-coupled differential pairs which have weaker coupling.

Guard traces with via stitching are necessary to ensure good continuity of the ground plane around the waveguide, and also provides additional isolation from nearby waveguides or other signals propagating in the dielectric between the reference planes. The general rule of thumb is to space the guard traces at least 2.5× as far from the waveguide trace as the dielectric thickness between the waveguide trace and ground plane. See chapter 3 of [51] for more information on isolation and guard traces for stripline waveguides (this discussion is also relevant to microstrip waveguides).

When specifying the widths of the impedance-controlled traces in a PCB, impedance calculators that use field-solvers are recommended over analytical approximations to

more accurately account for second order effects such as dispersion and surface roughness.

Proper dielectric selection is also important for high frequency signals, as waveguides that use standard FR-4 laminates used in low-speed inexpensive PCBs suffer from very high dielectric losses. Furthermore, the resin to fiberglass ratio and fiberglass weave are not well-controlled for standard FR-4, which can cause significant variations in the characteristic impedance of waveguides constructed with these dielectrics.

6.4 Power supply design

Power supply architectural design is dictated by the component selection for the signal chain. The first considerations that must be taken into account are the expected supply voltages and load currents, and then signal levels and PSRR should be taken into account to take appropriate steps to mitigate effects of power supply noise. Single-supply amplifiers simplify power design, but cannot generate bipolar outputs. Since we desire bipolar outputs, the power supply must generate both positive and negative voltage rails. Therefore, a switching regulator is necessary to convert the positive supply voltage from the RFSoC to a negative voltage. Furthermore, the power requirements of the amplifiers in the signal chain are sufficiently high so as to require a switching regulator to step down the RFSoC voltage.

Since we must use switching regulators for the power supply, proper component selection and layout are critical to limit power supply noise at the input of the ADC.

6.4.1 Linear regulators

A linear regulator is often implemented as a metal-oxide-semiconductor field effect transistor (MOSFET) biased in triode or bipolar junction transistor (BJT) biased in saturation with negative feedback, as shown in Figure 6-8. A MOSFET in the triode region or BJT in saturation operates as a variable resistor (set by the gate voltage/base current), so tuning the gate voltage/base current in closed-loop feedback

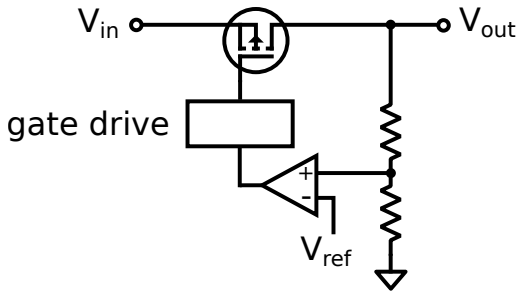

Figure 6-8: Linear regulator implemented with p-channel enhancement-mode MOSFET. A lower gate voltage increases the conductance of the channel, so the error amplifier increases the MOSFET resistance when the output voltage is too high, and decreases it when the voltage is too low.

allows the regulator to control the voltage supplied to the load. For a linear regulator, the voltage drop between the input voltage and regulated load voltage is dissipated in the form of heat. Therefore, it is usually best to operate the regulator with the lowest possible voltage drop between the input and output. The minimum dropout voltage is limited by the on resistance $R_{ds,on}$ of the MOSFET, and is therefore a function of the load current; higher load currents result in higher minimum dropout voltages. Typically, dropout voltage of inexpensive regulators is more than 100 mV. Ultra low-dropout linear regulators can provide very low dropout voltage (less than 100 mV) even under high load currents, which is desirable for high power applications, where the load current may be several amperes. Reducing the dissipated power in the linear regulator package down to a few watts (or even less than 1 W) is desirable from both a power efficiency and thermal perspective. For high voltage stepdowns, a linear regulator is dispreffered unless the expected load current is quite small.

6.4.2 Switching regulators

Linear regulators are very inefficient for large voltage stepdown, and at high power, the excess heat generated can present a challenge for reliability and lifetime of the power supply. For the analog frontend, the positive regulation chain needs to generate 3.3 V (at 220 mA per channel) and 2.5 V (at 70 mA per channel) from a 12 V supply. If we were to generate these voltages with linear regulators, they would dissipate over 2.5 W per channel (over 20 W in total for an eight-channel board). On top of

the huge waste of power, this approach is undesirable from a thermal perspective, as the heat must go somewhere, and most silicon electronics are only rated to junction temperatures of 125 °C. This would require a module with a thermal resistance less than 5 °C/W to prevent the regulator from overheating, entailing significant heat-sinking and potentially forced air flow. Switching regulators allow voltage conversion with minimal power dissipation, even for high step down voltages. It is not uncommon to achieve voltage conversion with power efficiency exceeding 90 % at high current loads, meaning the regulator would dissipate less than 1 W. Furthermore, switching regulators are necessary for generating the negative supply required for the analog frontend.

The main drawback of switching regulators is that they generate switching noise (both radiated and conducted) which can degrade the SNR or sensitive analog frontends. However, careful design of the switching regulator in terms of component selection and placement can suppress the switching noise so it has no impact on the noise performance of the frontend.

Operating principle of switching regulators

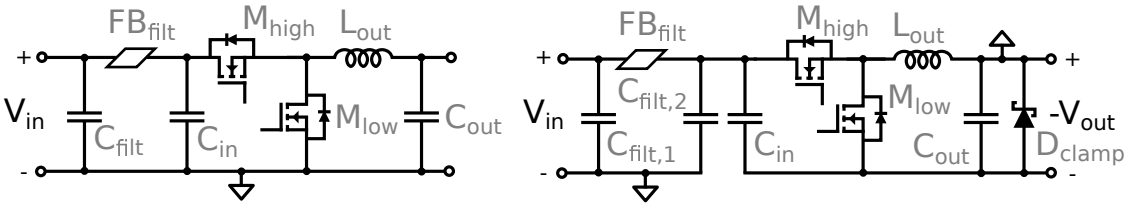

Figure 6-9: Synchronous buck and inverting buck-boost regulator topologies. Feedback control circuitry not shown.

A buck regulator and inverting buck boost regulator are shown in Figure 6-9. The transistors M_{high} and M_{low} pull the switching node between $V_{\text{in},+}$ and $V_{\text{out},-}$ (which in the case of the buck regulator is V_{in} and ground). This produces a square wave voltage on the switching node. A second-order lowpass filter (L_{out} and C_{out}) with a cutoff frequency well below the switching frequency or fundamental passes just the

DC component of the chopped waveform. The duty cycle of the square wave can be altered with pulse width modulation (PWM) to change the average value of the waveform, changing the output voltage of the converter. In reality, the second-order filter has a finite rolloff (40 dB per decade), so it is not able to infinitely suppress the fundamental and harmonics of the square wave, resulting in some output ripple. Furthermore, the physical filter will have non-idealities, such as self-inductance of the capacitor C_{out}, which creates high frequency noise due to the sharp transitions of the switching node voltage. Decreasing the frequency of the filter double-pole will improve attenuation of the switching frequency fundamental and harmonics, however, the maximum ripple attenuation that is practical for a second-order filter is between 40 dB and 80 dB. This is because the pole frequency of the filter cannot be too low, because the double pole of the filter introduces a 180° phase shift which must be compensated by the switching regulator controller. Most controllers are designed to boost phase margin one or two decades below the switching frequency of the regulator. More attenuation can be achieved by cascading multiple stages of filters, but care must be taken to ensure proper phase response of the feedback network to prevent instability. An input filter is also shown in 6-9. This input filter helps isolate upstream electronics from switching noise [52].

It is also important to note is that, due to the different topology of the inverting buck boost regulator, its transfer function has a right-half-plane zero. The dreaded right-half-plane zero increases loop gain above the zero frequency at 20 dB per decade like a left-half-plane zero, but simultaneously decreases the phase margin, making it difficult to stabilize the system. The expression for the right-half-plane zero frequency of the inverting buck boost regulator as described in [52] is given by equation 6.3:

$$\omega_z = \frac{R}{L_{\text{out}}} \left(\frac{(1-D)^2}{D} \right) \tag{6.3}$$

where R is the load impedance (*i.e.* desired output voltage divided by load current draw), L_{out} is the output inductor, and D is the duty cycle of the PWM signal. For high inductor values and high load currents, the zero frequency is pulled down,

increasing the risk of loop instability if it drops near or below the unity-gain frequency of the rest of the loop. As long as the magnitude of the output voltage is less than the input voltage, the duty cycle of an inverting regulator is low enough that the right-half-plane zero frequency is dominated by the load impedance to inductor ratio [53].

Feedback control and phase margin

As mentioned previously, the feedback network used to set the PWM duty cycle must have proper phase margin to ensure stability of the output voltage. If the combined response of the feedback system an open-loop phase shift greater than or equal to 180°, then the system will oscillate. This is a problem in linear regulators as well, but is exacerbated by the buck regulator output filter, which introduces a 180° phase shift (or more in the case of the inverting buck-boost regulator) that must be compensated for.

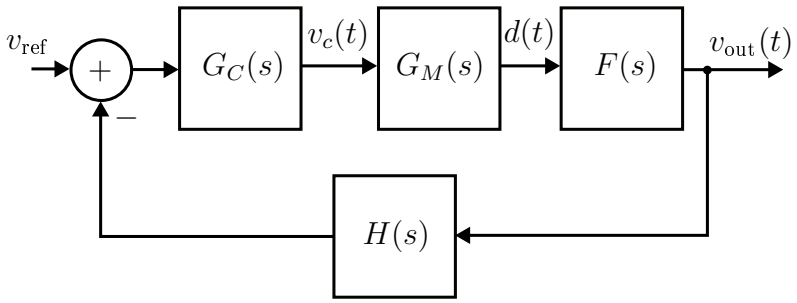

Figure 6-10: Feedback loop of switching regulator. The block $G_C(s)$ models the compensator, $G_M(s)$ models the pulse-width modulator, and the rightmost block $F(s)$ models the response of the switches and second-order LC low-pass filter. $H(s)$ models the feedback network (often implemented as a resistive divider, sometimes with a compensation capacitor).

Figure 6-10 shows the feedback loop of a typical switching regulator. The PWM generation is modeled by the transfer function $G_M(s)$, and the high/low-side switches and output filter are modeled by $F(s)$. A compensation circuit modeled by the transfer function $G_C(s)$ is typically constructed around the error amplifier to boost the phase margin of the loop [52]. Depending on the switching converter load conditions (and the filter transfer function), one or more left-half-plane zeros are used in

the compensation circuit. Sometimes, additional compensation is performed with the feedback network modeled by $H(s)$. Typically, the simplest implementation of a feedback network is just a resistive divider. However, a capacitor is sometimes included in parallel with the upper resistor in the divider to introduce a left-half-plane zero, creating a phase lead to help compensate the phase lag from the output filter $F(s)$. Most commercial switching regulators include internal compensation networks as part of the error amplifier circuit which are designed for a variety of load conditions.

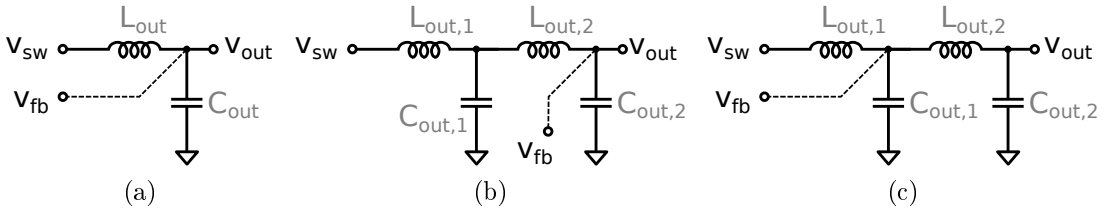

Figure 6-11: Various output filter and feedback configurations for a buck converter.

Depending on the output filter configuration, the phase margin can vary widely as shown in Figure 6-12. A variety of output filter and feedback configurations (shown in Figure 6-11) are used to close the loop of the switching converter. The first configuration (Figure 6-11a) is the standard configuration for a buck converter with a single filter stage. The second configuration illustrates an attempt at decreasing output ripple with improper feedback that will likely lead to an unstable system. The final configuration shows the same configuration as the second, but with the center of the filter tapped for the feedback voltage. This achieves the same level of ripple attenuation as the second configuration without the stability problems.

The Bode plot of $v_{\text{fb}}/v_{\text{sw}}$ in Figure 6-12a shows the transfer function from the switching node to the feedback node used to set the duty cycle of the switching converter. In the second configuration, where an additional filter stage is added with the hope of further attenuation of output ripple, the phase margin drops precipitously around 100 kHz due to the second-double pole relatively close to the dominant double-pole. This large of a phase drop so close to the dominant double-pole at 10 kHz is undesirable, as there is a good chance the open loop gain of the system in Figure 6-10 does not roll off fast enough to prevent instability, since the unity-gain

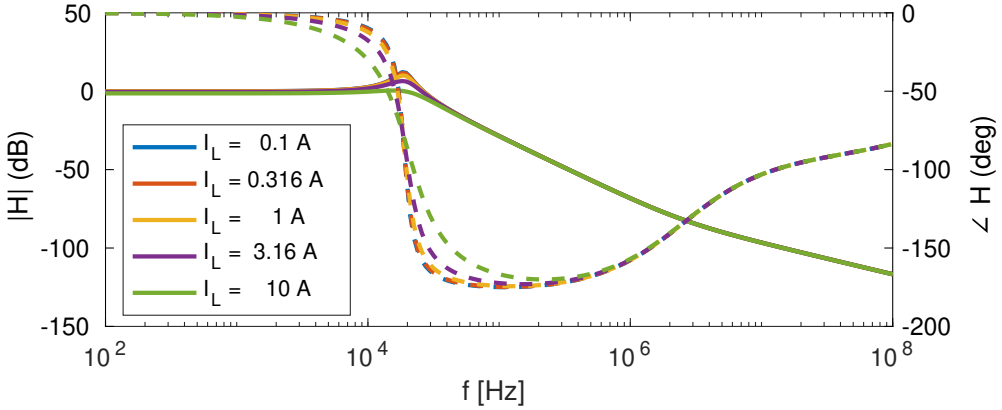

(a) Frequency response from input to output of single-stage low-pass filter (6-11a)

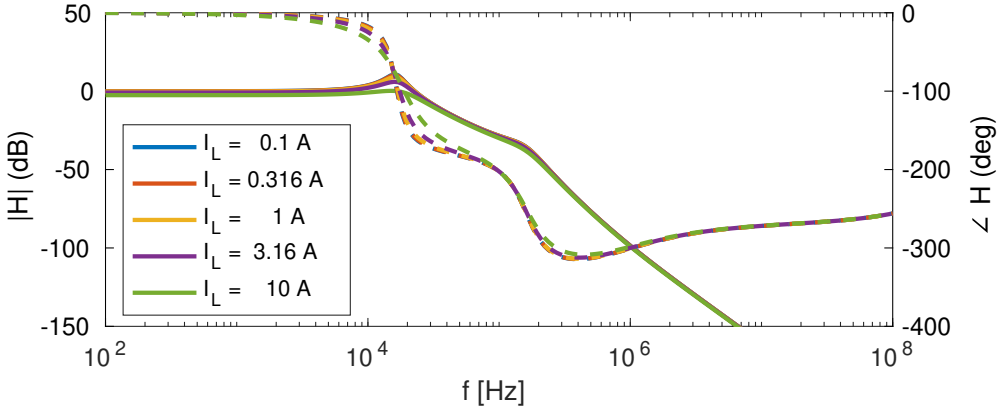

(b) Frequency response from input to output of double-stage low-pass filter (6-11b)

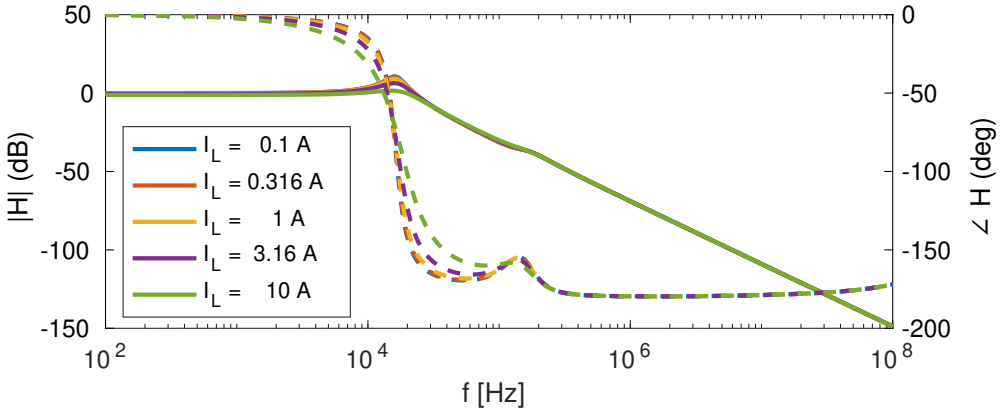

(c) Frequency response from input to center node of double-stage low-pass filter (6-12c)

Figure 6-12: Phase and magnitude of feedback input to error amplifier for various buck filter configurations. Components are modeled with realistic parasitics based on the self-resonant frequencies and equivalent series resistances reported on the component datasheets.

frequency does not change drastically, despite the extra 180° phase shift. In a system where the engineer knows the compensation $H(s)$ used in the feedback loop, it would be straightforward to ensure adequate phase margin and stability of this filter and feedback configuration. However, most designers opt for an integrated solution for the feedback control (either in the form of a switching controller or switching controller with integrated switches) to save cost and engineering hours. In this case, the compensation used in the feedback loop is often unknown, and only a range of recommended inductor and capacitor values for a single output stage are provided by the manufacturer. The final configuration, shown in Figure 6-11c, uses the same two-stage filter as before, but instead uses the center node of the filter as the feedback voltage for the closed loop duty cycle control. This allows for straightforward integration with commercial switching regulators without having to worry about loop stability, provided the appropriate values for the first stage are chosen, and that the quality factor of the second resonance is sufficiently low enough that the gain peaking does not push out the open-loop unity-gain frequency too much.

Switching regulator layout

Layout of the switching regulators is critical for noise performance, since the loops formed between the input and output of the regulators act as antennas, and radiate considerable electromagnetic interference when the switches in the regulator are toggled. Furthermore, the switching node can act as an antenna if it is too long, resulting in further increases in radiating switching noise. This radiated switching noise can couple to sensitive electronics and degrade their noise performance. Fortunately, although critical, the layout of switching regulators is quite simple as long as two guidelines are followed: minimize the area of loops that undergo large current transients, and minimize the area of the switching node to decrease antenna effects. Figure 6-13 shows an example of a recommended layout for a switching regulator from [54]. This particular regulator has separate supply pins for the analog feedback control circuitry (`AVIN`, `AGND`), presumably to improve the output noise of the regulator by increasing isolation from the large switching transients present on the input and

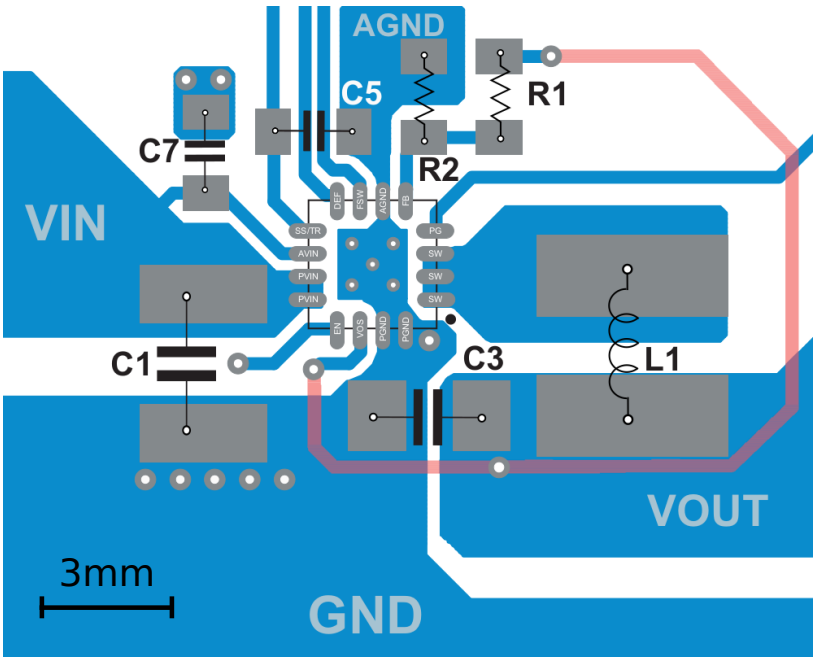

Figure 6-13: Recommended layout for a switching regulator. The input loop (`PGND` to `PVIN` through C_1) and output loop (`SW` to `PGND` through L_1 and C_3) are kept as small as possible. The area of the switching node is kept as small as possible.

output power rails. Note that the input and output capacitors are placed as close as possible to the integrated circuit to minimize loop area. The switching node is kept as small as is reasonable as well.

6.4.3 Analog frontend power supply

The two switching regulators (Texas Instruments TPS62130) generate −3.0 V and 3.9 V as shown in Figure 6-14. These two rails are then further regulated with linear regulators (Texas Instruments TPS7A94 and Analog LT3015) to generate −2.5 V, −1.8 V, 2.5 V and 3.3 V rails. In order to prevent the positive rail from pulling the negative output of the inverting regulator positive before the inverting regulator is able to initialize, the two switching converters are sequenced so that the buck regulator turns on after the inverting regulator [53]. This is achieved with the level-shifting circuit that uses the power-good (`PG`) output of the inverting regulator to enable the buck regulator.

The input and output of the switching regulator are filtered with pi-LC filters

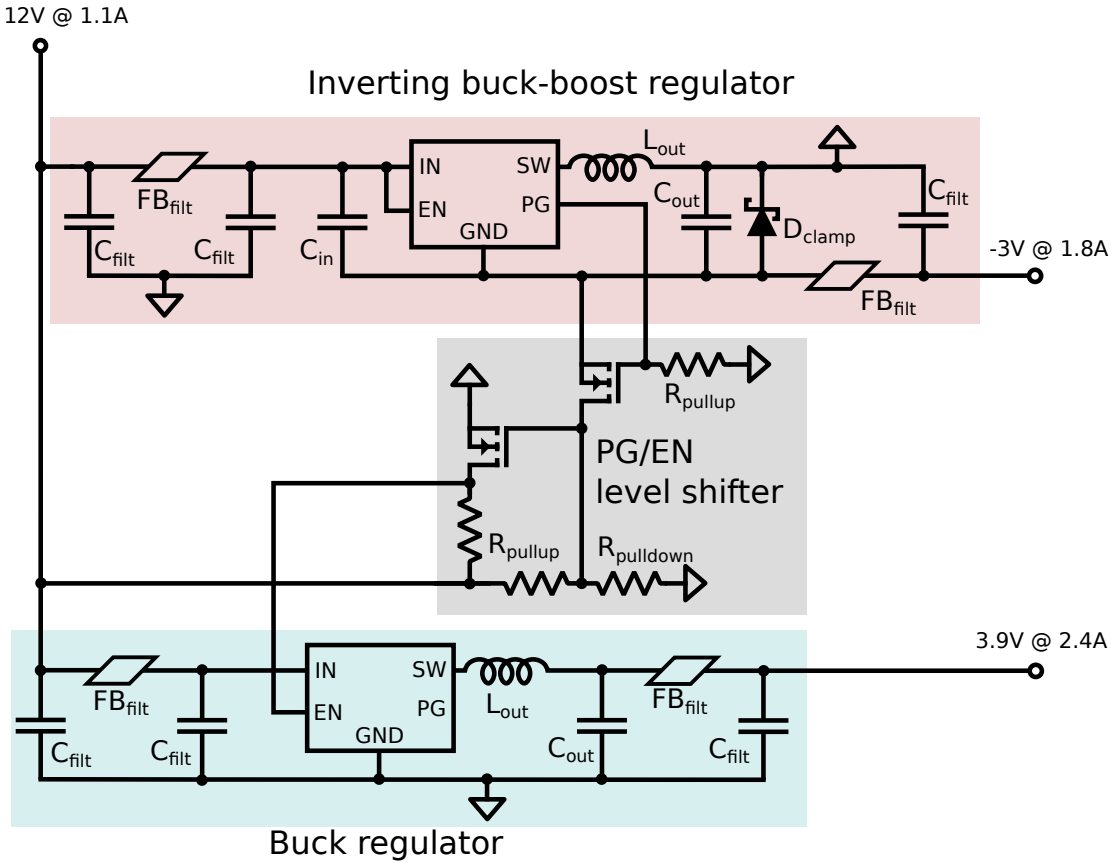

Figure 6-14: Generation of negative and positive rails with switching regulators. A level shifting circuit is used to convert the power-good (PG) output of the inverting regulator into an enable signal for the buck regulator. This is done to help prevent charging of the negative rail to a positive voltage, which could cause problems in the feedback loop of the inverting regulator. The clamp diode also helps mitigate precharging of the negative rail due to inrush current through the input capacitor C_{in} on startup.

implemented with ferrite beads. The low quality factor of the filter resonance that results from using a ferrite bead serves two purposes. For one, it reduces gain peaking which could amplify noise near the resonance of the filter. Second, the overdamped resonance has little impact on the phase margin of the control loop, whereas a high-Q resonant filter on either the input or the output would reduce the phase margin, potentially making the regulator unstable [52]. Just like for the ADC and DAC buffers, the bypass capacitors shown in Figure 6-14 are implemented with multiple capacitors of different value and package size, again to improve filtering at high frequencies by increasing the effective self-resonant frequency. The same is true for the linear

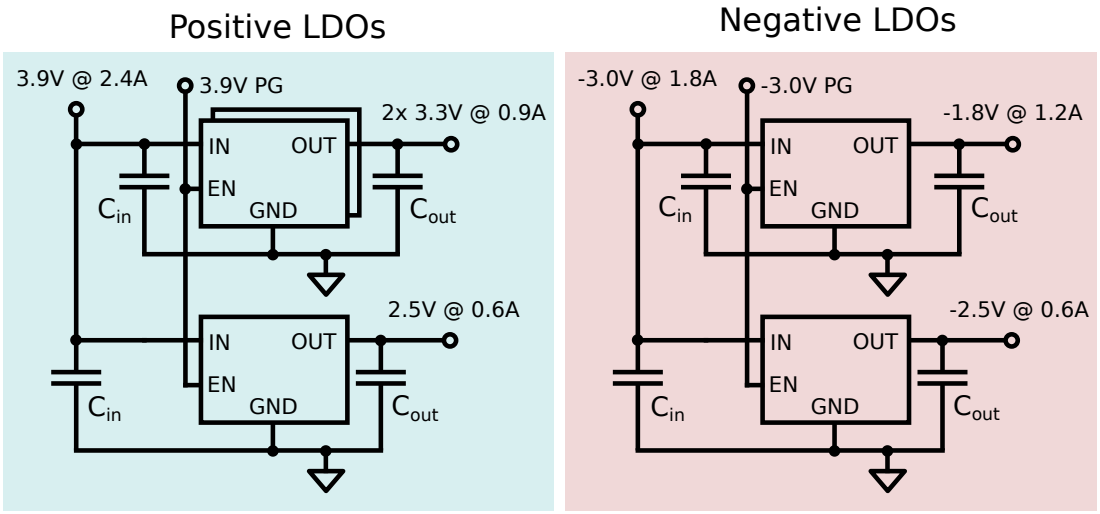

Figure 6-15: Active filtering of low frequency noise and switching fundamental with low noise low dropout linear regulators.

regulator output capacitors in Figure 6-15.

After passing through the two-stage output filter of the switching regulator and the linear regulator stage, the amplitude of the fundamental from the switching regulator is expected to be close to or even below the noise floor of the linear regulator (rms noise voltage of roughly 1 µV). Indeed, as we will see in the next chapter, there is no noticeable signal in the spectrum measured by the RFADC at the switching frequency of the regulator.

Chapter 7

Characterization of the analog frontend

Testing superconducting nanowires requires equipment that is sensitive to small voltages, and introduces a minimal amount of noise. For example, a nanowire with a critical current of 2 µA that is shunted with 50 Ω would produce a voltage spike with an amplitude of roughly 100 µV. Being able to recognize the pulse of a switching nanowire and distinguish it from other signals (such as the L/R differentiator discussed at the end of section 4.1) means that the amplification process should also introduce a negligible amount of distortion. Therefore, the receive signal chain of the test equipment should have high gain, low noise, and high linearity. Furthermore, for the input and output impedance of the test equipment should be as closely matched to 50 Ω as possible to minimize signal reflections. The reflected signals at best make it difficult to analyze the output of the superconducting circuit, and at worst actually impact the performance of the circuit.

For these reasons, it is important to fully characterize the gain, linearity, noise performance, and input/output impedance of the analog frontend to verify that it will be able to interface with the superconducting circuits we wish to test.

7.1 Low-frequency gain, linearity, and noise characterization with RFSoC

Using the RFSoC to characterize all of these quantities would be desirable since it provides an opportunity to develop firmware infrastructure and know-how that will be useful when developing the FPGA firmware for characterizing superconducting nanowire electronics. However, scattering parameter measurements (to analyze input and output impedance) are not possible without constructing specialized microwave structures and analog circuits, so those measurements are best left to a vector network analyzer (VNA). Furthermore, the bandwidth of the RFADCs in the RFSoC is limited by the sampling rate f_s of the RFADC to 2 GHz. Note that the RFADCs can operate in the second Nyquist zone from $f_s/2$ to f_s, so it would be possible to measure the gain of the DAC buffer with the RFSoC, but this would be tedious, and a VNA would still be needed to extract phase information and the other scattering parameters. Therefore, the RFSoC is only used for measuring the low frequency (<100 MHz) gain, linearity, and noise of the analog frontend.

7.1.1 High purity tone generation with direct digital synthesis (DDS)

The RFSoC can be used to test the analog frontend by sending pure sinusoidal signals into the frontend and measuring the amount of distortion and noise the frontend introduces. If the stimulus generated by the RFSoC has significant distortion or noise, then it is more difficult to disentangle the contribution from the analog frontend to the overall noise and distortion of the measured response. As mentioned in section 2.5, DDS is a technique that allows for the generation of sinusoidal tones with very high SFDR and SNR.

DDS implementations can be made quite compact with phase quantization, however there is a tradeoff between memory utilization and spectral purity caused by quantization noise. In an ideal world, there would be no phase quantization, however

a 16 bit wide lookup table for 24 bit phase factors takes up 256 Mibit, which is an enormous amount of memory on an FPGA. External SDRAM memory is an option for lookup tables due to its high storage density, however the bandwidth of SDRAM is typically limited to a few Gbit/s, so generating several multi-GS/s signals using lookup tables stored in SDRAM is out of the question. In order to save space, the lower bits of the DDS phase accumulator are often truncated. However, this truncation process introduces phase noise, and due to the periodic nature of the truncated residuals, the phase noise manifests as a series of spurs in frequency domain, as can be seen in Figure 7-1b.

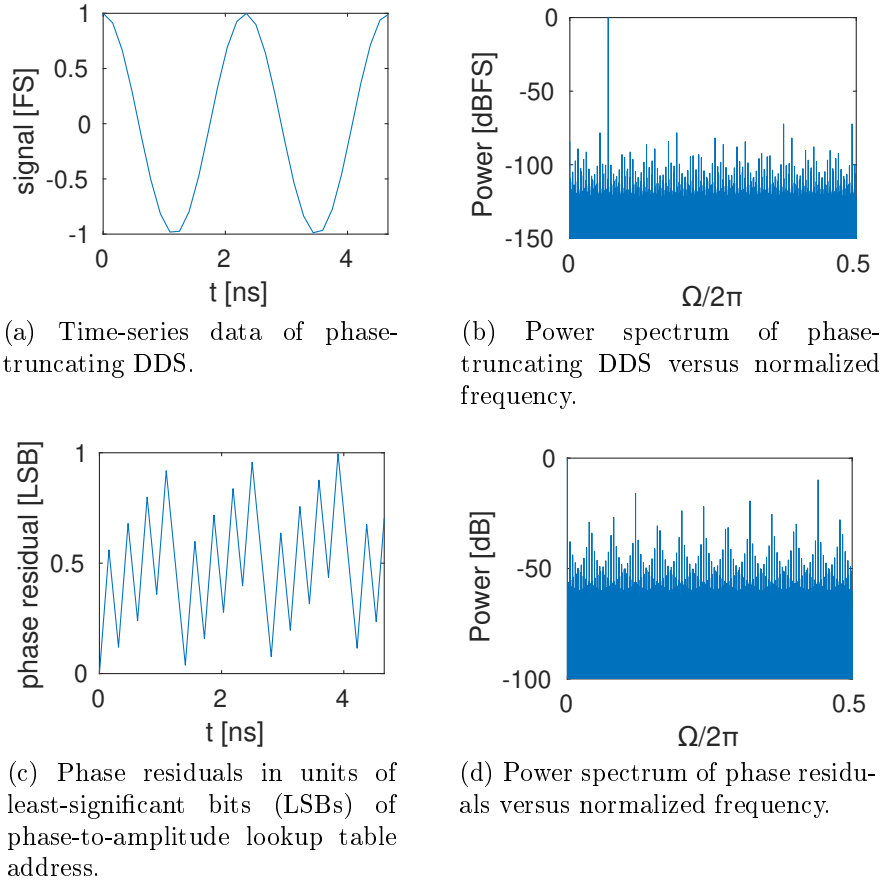

(a) Time-series data of phase-truncating DDS.

(b) Power spectrum of phase-truncating DDS versus normalized frequency.

(c) Phase residuals in units of least-significant bits (LSBs) of phase-to-amplitude lookup table address.

(d) Power spectrum of phase residuals versus normalized frequency.

Figure 7-1: Time-series data and power spectrum of sinusoidal waveform generated by truncating phase residuals. The effect of phase quantization on spectral purity is quite visible. The residuals and their power spectrum are also plotted. DDS was configured with $M = 24, N = 12, B = 16$.

The SFDR of a sinusoidal signal is a measure of the height of the largest spur

relative to the fundamental, typically measured in dBc (dB to carrier). If the SFDR of the DDS generator is too low, then we won't be able to measure the nonlinearity introduced by the analog frontend because the digital signal will already have spurs in the frequency domain before it's converted to an analog signal and passed through the analog frontend. Ideally, the output sinusoid that is used as stimulus to the analog frontend would have the largest possible SFDR, but phase and amplitude quantization in the digital domain and nonlinearities and noise in the RFDAC limit the realizable SFDR. In the digital domain, the dominant factor that limits SFDR is phase quantization. As we can see in 7-2b, introduction of a dither signal before phase quantization can flatten the spurs in the frequency domain, improving SFDR.

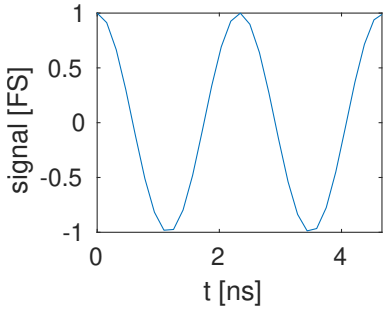

(a) Time-series data of phase-dithering DDS.

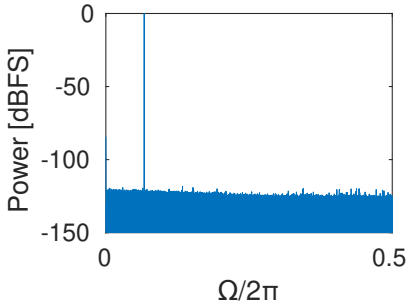

(b) Power spectrum of phase-dithering DDS versus normalized frequency.

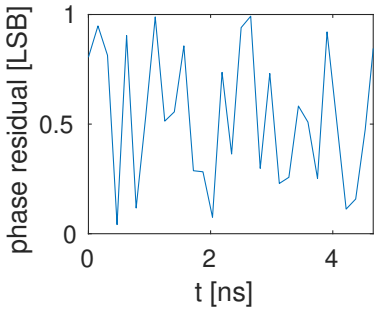

(c) Phase residuals after dithering.

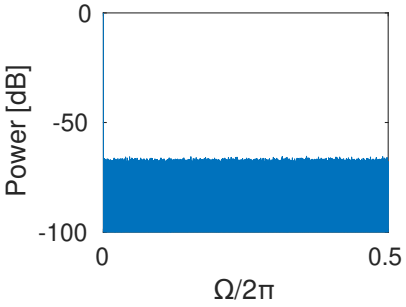

(d) Power spectrum of phase residuals versus normalized frequency.

Figure 7-2: Time-series data and power spectrum of residuals for a phase-dithering DDS. Phase-dithering improves SFDR by roughly 35 dB at the cost of a 3 dB reduction in SNR. The spectrum of the phase-dithering DDS residuals is much flatter than that of the phase-truncating DDS. DDS configured with $M = 24, N = 12, B = 16$.

Due to the truncating process which generates the residuals, they have a sawtooth-

like time-series behavior, which when translated to the frequency domain, results in a lot of spurs, as can be seen in Figure 7-1d The effect of phase quantization on the generated signal $x[n]$ can therefore be understood by considering Equation 7.1:

$$\begin{aligned} x[n] &= \cos(\hat{\theta}[n]) \\ &= \cos(\theta[n] - \widetilde{\theta}[n]) \\ &= \cos(\theta[n])\cos(\widetilde{\theta}[n]) + \sin(\theta[n])\sin(\widetilde{\theta}[n]) \\ &\approx \cos(\theta[n]) + \sin(\theta[n])\widetilde{\theta}[n] \\ &\approx \cos(\theta[n])\sqrt{1+\widetilde{\theta}[n]^2} \end{aligned} \tag{7.1}$$

where $\theta[n]$ are the unquantized, full-precision phases, $\hat{\theta}[n]$ are the quantized phases, and $\widetilde{\theta}[n]$ are the residuals as a result of quantization. For small quantization errors ($\widetilde{\theta}[n] \ll \theta[n]$), the approximation is accurate, so the phase modulation from the quantization of $\theta[n] \to \hat{\theta}[n]$ looks like amplitude modulation by the factor $\sqrt{1+\widetilde{\theta}[n]^2}$. Amplitude modulation results in frequency mixing, so if we take the pure sinusoidal waveform generated by the sequence of phases $\theta[n]$ and mix it with the residuals $\widetilde{\theta}[n]$, then it is unsurprising that the process of phase quantization produces the spurs we see in the spectrum in Figure 7-1b.

The simplest technique for flattening these spurs is to dither the phase before quantization is performed. Phase dithering is the addition of a random signal generated by a white process with variance equal to the least significant bit of the truncated phase [21]. The dither signal essentially drowns out the residuals, resulting in a flat residual spectrum (7-2d), which when mixed with the pure sinusoidal tone, results in a relatively clean spectrum (7-2b). In this case, the SFDR increased substantially from 72 dBc to 107 dBc.

It is important to emphasize that although the SFDR is improved by dithering, the SNR actually decreases slightly (the amount depending of course on the amplitude of the dither signal). This is because the dithering process doesn't actually remove the phase quantization noise, it just spreads it out. In doing so, it actually introduces a small amount of noise. Other more advanced techniques, such as Taylor expansion

correction can improve both SFDR and SNR simultaneously [55].

7.1.2 Measurement setup

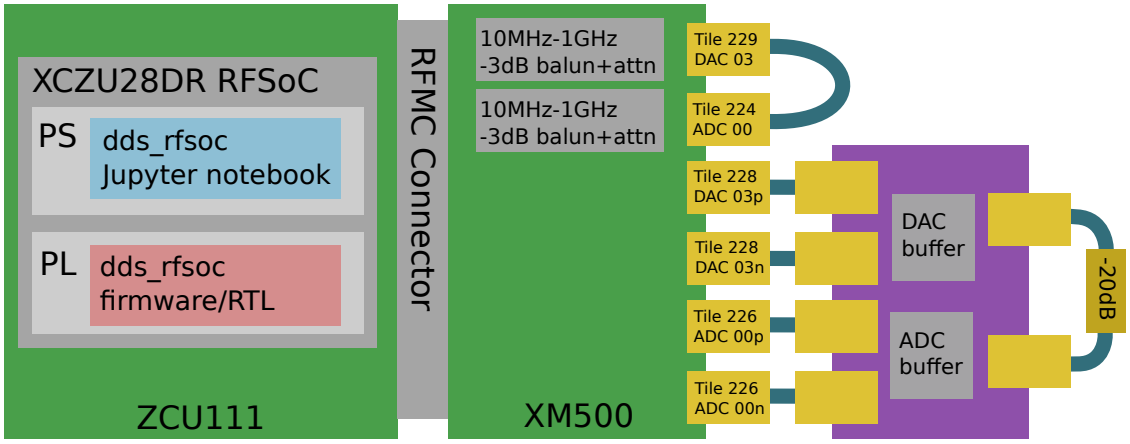

Figure 7-3: Diagram of setup for characterization of analog frontend with the RFSoC. An RFADC/RFDAC pair is connected in loopback with two baluns on the XM500 breakout board. The analog frontend is also configured in loopback; the output of the DAC buffer is fed through a 20 dB attenuator into the ADC buffer input. The value of the attenuator was varied for the linearity and noise measurements to access a greater dynamic range of input powers referred to the ADC buffer input.

Custom firmware [56] running on the RFSoC was used to measure the linearity, noise, and low frequency gain of the analog frontend. The PYNQ python productivity toolkit for Xilinx Zynq SoCs was used to simplify the development of software that runs on the ARM CPU and interfaces with the FPGA firmware. The FPGA firmware was used to generate a sinusoidal tone with DDS and measure the response of both the analog frontend and a pair of low-frequency (10 MHz-1 GHz) baluns provided on the XM500 breakout board. Since the baluns are passive devices, they should introduce negligible nonlinearity and noise, so by comparing the response of the balun circuit with the analog frontend, we can quantify the linearity and noise degradation introduced by the analog frontend. The analog frontend was configured in loopback with the DAC buffer output passed through one or more attenuators into the ADC buffer input. Because the DAC buffer outputs about 7 dBm at full scale, the attenuator was necessary to prevent it from saturating the input of the LNA, which has an input 1 dB compression point (IP1dB) of -2 dBm. The total attenuation was varied

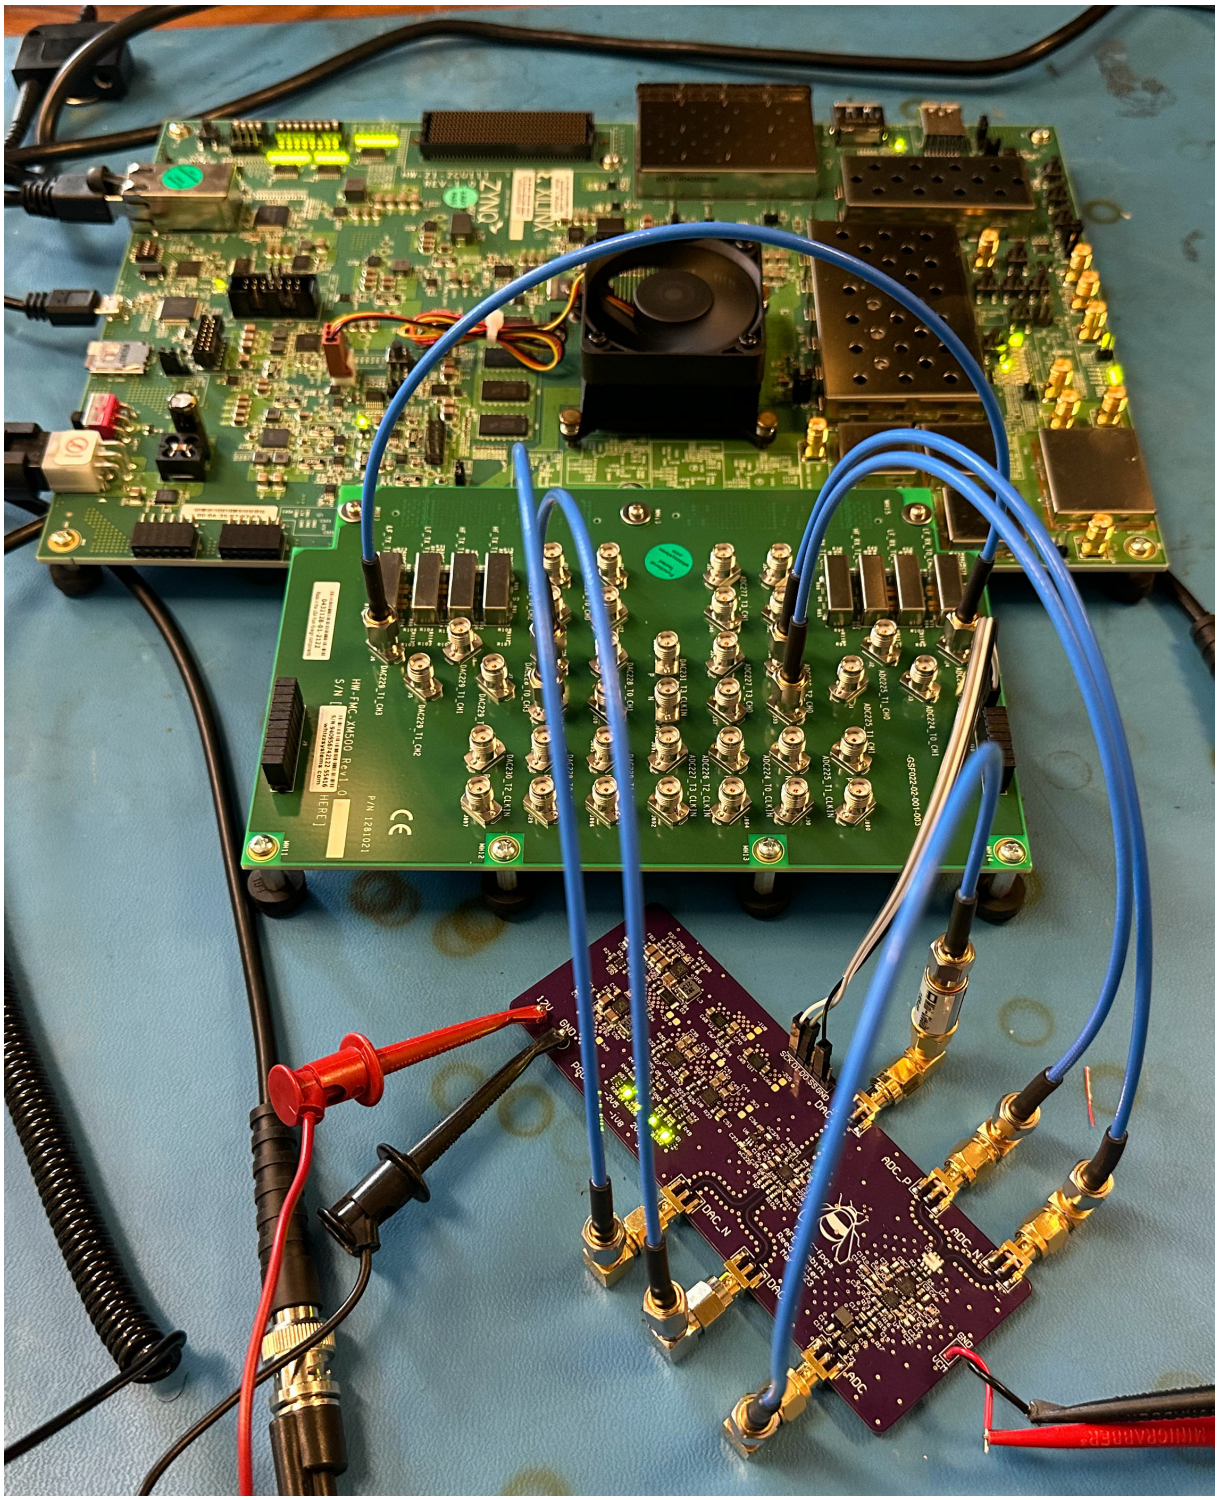

Figure 7-4: Setup of analog frontend characterization with the RFSoC. The ZCU111 and XM500 are the two green boards is in the back of the image, and the purple board is the prototype analog frontend. A 12 V power supply and 1.2 V low-noise voltage source are connected to the analog frontend to power it and provide a common-mode bias for the DC-coupled RFADC inputs.

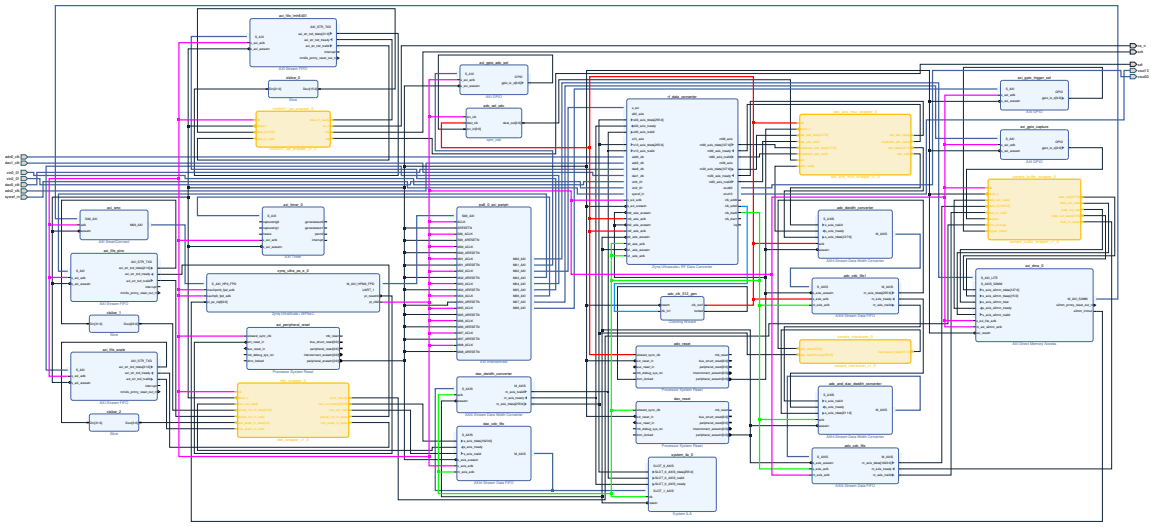

Figure 7-5: Block diagram of programmable logic. Separate clock domains are color coded, and custom modules are highlighted yellow. The rest of the modules are Xilinx IP blocks. The clocks for the RFDC IP data interfaces are generated from the RFDC PLL output clocks. Both the RFDAC and RFADC are run at a sample rate of 4.096 GS/s. The RFDAC interface is clocked at 256 MHz (16 samples per clock), and the RFADC interface at 512 MHz (8 samples per clock). The RFDC PLL input clocks are generated by the LMK04208 PLL configured to output 122.8 MHz and two LMX2594 PLLs configured to output 409.6 MHz.

between 20 dB and 57 dB for the linearity and noise experiments. The low frequency gain measurement was performed with 20 dB of attenuation.

Figure 7-4 shows the setup with the RFSoC connected to the XM500 breakout board. The XM500 breakout board contains both single-ended inputs/outputs that are connected to the dataconverters through baluns as well as DC-coupled differential inputs/outputs. One of the single-ended, balun-connected RFDAC outputs is sent to one of the single-ended, balun-connected RFADC inputs for the balun loopback configuration. The analog frontend is connected to the DC-coupled differential inputs/outputs. The power for the analog frontend is supplied with a bench-top power supply with 12 V at 240 mA output current. The ADC buffer is biased at 1.2 V according to the specification for the RFADC common-mode voltage. An external supply is used to bias the ADC buffer common-mode since the high density mezzanine connector includes the RFADC common-mode bias voltage (which will be used in later revisions of the analog frontend which use the mezzanine connector), but

it is not broken out on the XM500 breakout board. Generation of the clock input for the radio-frequency dataconverter (RFDC) phase-locked loop (PLL) is performed by two PLL and clock conditioning circuits from Texas Instruments on the RFSoC development board. A voltage-controlled crystal oscillator feeds a LMK04208 PLL which is configured to generate a 122.8 MHz clock which is then converted into two 409.6 MHz clocks (one each for the RFADC and RFDAC) by two LMX2594 PLLs. An internal PLLs that is part of the RFDC generates the necessary sampling clocks for analog-to-digital and digital-to-analog conversion.

The block diagram of the FPGA firmware is shown in Figure 7-5. The samples are generated with DDS in a slow clock domain (150 MHz) that is synchronous with the memory-mapped interface that the quad-core ARM CPU uses to communicate with the FPGA fabric. Standard Xilinx IP for clock domain crossing are used to step between the slow clock domain and faster clock domains used for interfacing with the RFDCs. A 1 MS buffer captures the digitized output of the RFADC, and can be configured to trigger manually (when sent a trigger from software on the ARM CPU) or automatically whenever the frequency of the DDS module is updated.

7.1.3 Low-frequency gain measurements

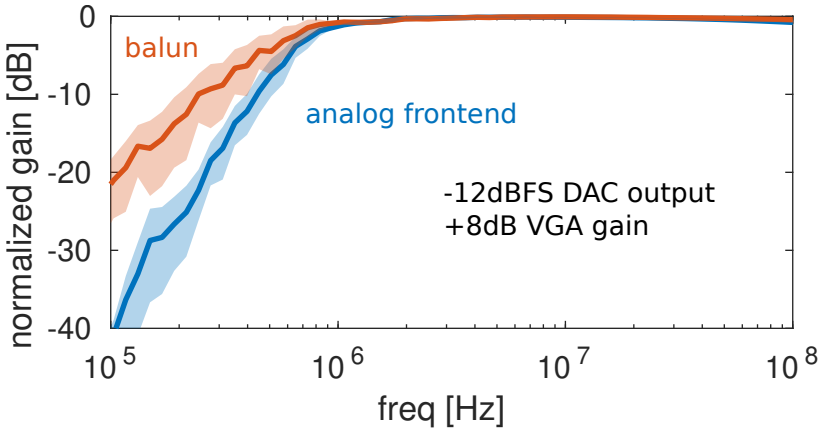

Figure 7-6: Low frequency gain characteristics of analog frontend and balun loopback configuration. The balun is only rated to work down to 10 MHz, but clearly it extends over a decade in frequency below that. The rolloff of the analog frontend response is twice as fast as that of the balun, due the existence of two inverted poles (instead of just one).

Characterization of low-frequency gain of the analog frontend is straightforward with DDS by measuring the amplitude of the digitized waveform over a range of single-tone stimuli at various frequencies. The DDS synthesizer was implemented with a phase factor width of 24 bits running at a sample rate of 4.096 GS/s, enabling the RFSoC to generate sinusoidal tones with frequencies as low as 244 Hz. This is more than low enough for characterizing the ADC buffer, which is expected to have a high-pass cutoff around 1 MHz.

Figure 7-6 shows the measured low-frequency gain of the analog frontend in comparison to the balun loopback configuration. As we can see from the plot, the response of both configurations begins to roll off slightly below 1 MHz, with a slope of roughly 20 dB per decade for the balun and 40 dB per decade for the analog frontend. This steeper slope from the analog frontend is indicative of either a double pole or two separate poles that are close in frequency.

The LNA is expected to have a high-pass behavior for two reasons, both of which give rise to inverted poles in the frequency domain response. The dominant inverted pole is due to the L/R cutoff from the RF choke used to bias the LNA, which is expected to have a -3 dB cutoff around 1 MHz. AC-coupling capacitors used to isolate the biasing of the LNA input from the DC component of the input signal give rise to the second inverted pole, which has a -3 dB cutoff near 400 kHz. These pole frequencies agree well with the observed trend shown in Figure 7-6.

7.1.4 Linearity and noise performance measurements

To quantify the linearity and noise performance of the ADC input buffer, SFDR and signal to noise and distortion ratio (SINAD) were measured for a variety of RFDAC output powers across six frequencies spanning three decades in frequency. Note that SINAD is an equivalent to SNR in the low-distortion limit for single-tone inputs (SNR will always exceed SINAD). The VGA gain was tuned to find the optimal SINAD and SFDR for each input power and frequency. The linearity and noise performance of the DAC buffer were not characterized due to time constraints.

To measure SFDR and SNR, a 1 MS-long transient signal is captured by the FPGA

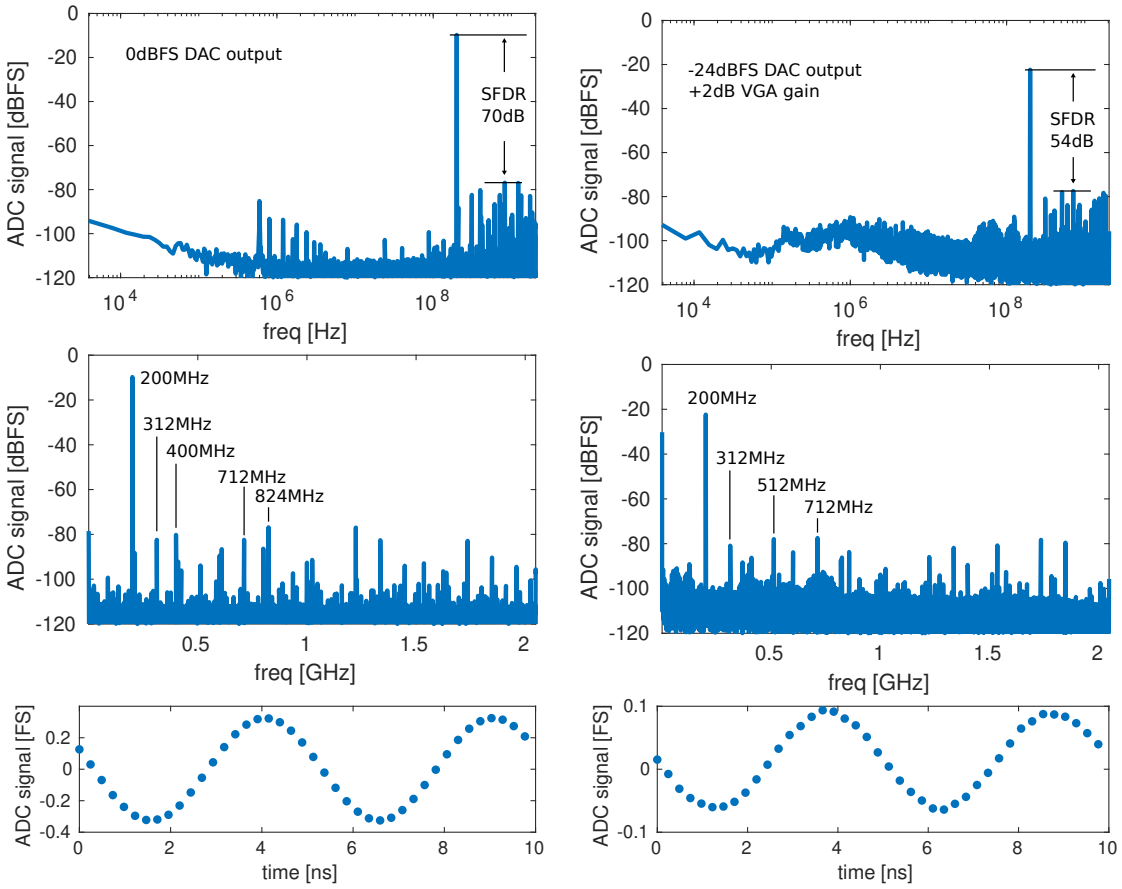

(a) Response of balun loopback for 200 MHz tone at 0 dBFS RFDAC output power.

(b) Response of AFE in loopback for 200 MHz tone at −24 dBFS RFDAC output power.

Figure 7-7: Power spectrum and transient response of analog frontend in comparison to balun loopback for a stimulus tone of 200 MHz. Note that the distortion and noise are small enough relative to the signal tone that they are virtually invisible in the transient plot. The analog frontend increased the noise power of the digitized signal by 6 dB and reduced the SFDR by 16 dB as compared to the balun loopback (however, the input power to the ADC buffer is 37 dB lower than the input to the RFADC balun). The spurs in the spectrum of the analog frontend are dominated by frequency mixing products of the 200 MHz fundamental with harmonics of the sampling clock, while the spurs in the spectrum of the balun loopback configuration are both harmonics of the 200 MHz fundamental and products of upconversion with the sampling clock subharmonic.

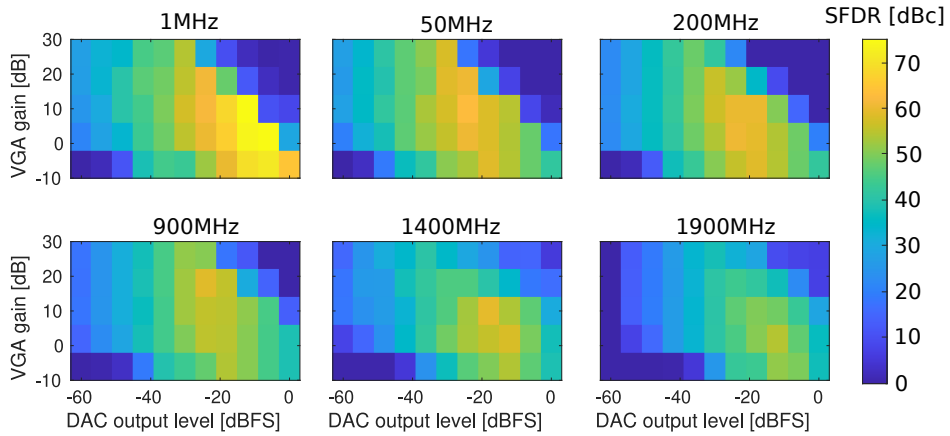

Figure 7-8: SFDR of analog frontend versus RFDAC power and VGA gain for various frequencies. Figure 7-7 shows The loopback attenuation was −20 dB. These plots show that, for a given signal level at the input of the ADC buffer, the VGA gain can be set so that the RFADC receives high linearity tone. For low output power and low amplification or high output power and high amplification, the maximum achievable SFDR is reduced.

fabric, and then postprocessed with `scipy` and `numpy` in the Jupyter notebook running on the ARM CPU. The SFDR was estimated by taking a Fourier transform of the digitized waveform and comparing the relative heights of the fundamental and next largest spur (harmonic or anharmonic, excluding DC). SINAD was estimated similarly to the method used by Matlab's `snr()`. First, a periodogram is calculated using a Kaiser window ($\beta = 38$) to estimate the power spectral density of the digitized waveform. Then, the fundamental is found and subtracted from the spectrum. Finally, the total noise and distortion power is calculated by integrating the remaining spectrum. This can be used to calculate the SINAD by dividing the power contained in the fundamental by the integrated noise and distortion power.

Figure 7-7 shows an example transient signal and spectrum captured by the RFSoC. Spurs from the sampling clocks are visible in 7-7b, and mixing products are visible in both 7-7a and 7-7b. The log frequency plot in 7-7b shows that there is negligible switching noise from the power supply (which would be at integer multiples of 2 MHz). The analog frontend has a 16 dB lower SFDR and 6 dB lower SINAD than the balun in this particular setup. However, it is important to note that the input power to the ADC buffer is 37 dB lower than the input power to the balun-connected

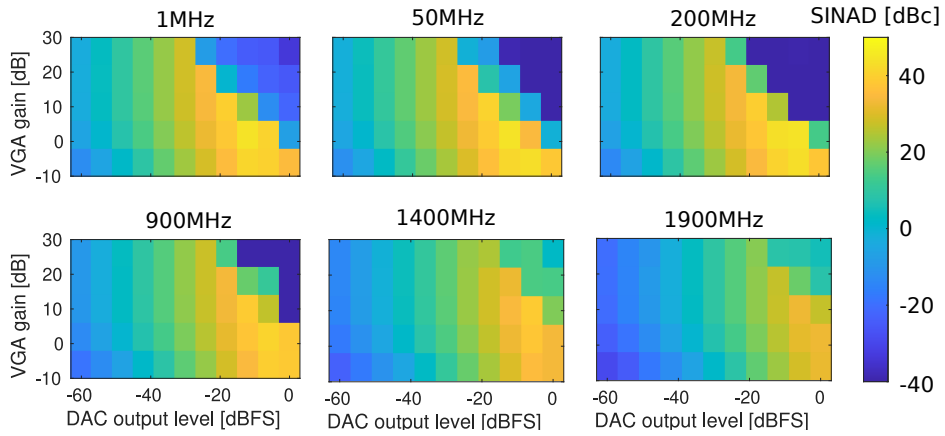

Figure 7-9: SINAD of analog frontend versus RFDAC power and VGA gain for various frequencies. The loopback attenuation was −20 dB. At high input power and high VGA gain, the SINAD is degraded substantially due to nonlinearities in the ADC buffer. The SINAD is also diminished for low input power. The key takeaway from this figure is that the SINAD (and SNR) can be kept relatively high even for small input powers at high frequencies, provided the VGA gain is set appropriately.

RFADC, so this isn't quite a fair comparison.

Linearity and noise performance are limited for very large and very small signals as shown in Figures 7-8 and 7-9. In these plots, both SFDR and SINAD are worse for low VGA gain settings when the input power to the ADC buffer is small, and for high VGA gain settings when the input power is large. These two cases are illustrated in Figure 7-10. The reduction in SFDR and SINAD for high signal power and high VGA gain (7-10a) is due to nonlinearities in the ADC buffer and the RFADC itself (since the ADC buffer can overdrive the input range of the RFADC). The reduction in SFDR and SINAD for low signal power and low gain (7-10b) is due to two factors. The first is that the signal is simply weaker, so it is closer to the noise floor and any anharmonic spurs (*e.g.* from sampling clock feedthrough, such as the 512 MHz tone in Figure 7-7b) increase noise and distortion power. Second, the RFDAC output power is programmatically decreased by right shifting and quantizing the digital samples before digital-to-analog conversion, so a RFDAC output level of −60 dBFS has quantization noise that limits the SFDR to about 30 dB and SINAD to about 24 dB.

If we take the VGA gain setting for each input power which maximizes SFDR and SINAD, we can estimate the minimum signal power that can be detected by the

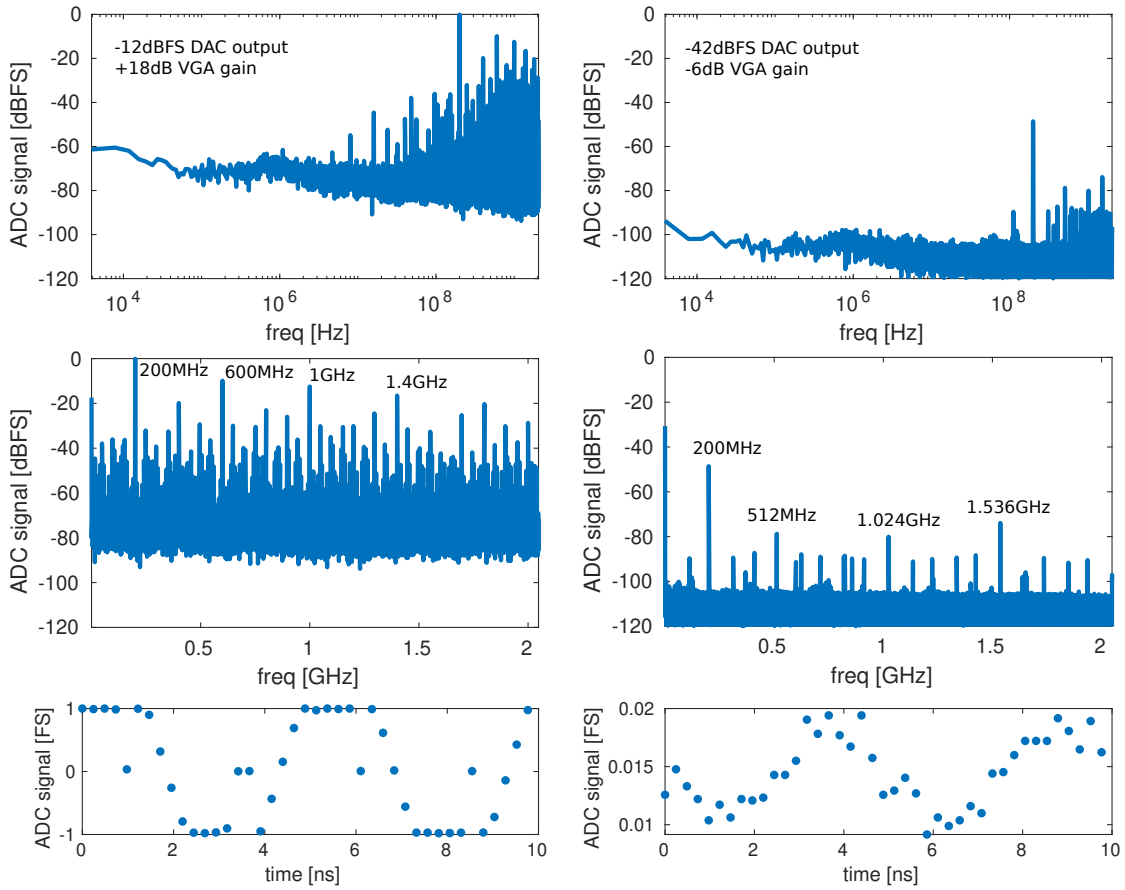

(a) Response of AFE in loopback with high amplification and strong signal resulting in clipping. This combination of gain setting and input power corresponds to the upper right corner of the plots in Figures 7-8 and 7-9. The RFDAC output power is −12 dBFS, and the VGA is configured for 18 dB of gain. Based on these gain settings, the input level to the RFADC given an ideal linear AFE would be roughly 11 dBFS.

(b) Response of AFE in loopback for low amplification, resulting in poor SNR and SFDR due to relatively large anharmonic spurs from the RFADC and RFDAC sample clocks (spurs at 512 MHz, 1.024 GHz and 1.536 GHz). This combination of gain setting and input power corresponds to the lower left corner of the plots in Figures 7-8 and 7-9. Note that the noise floor is a about 10 dB lower than in Figure 7-7b, indicating that the noise performance is (unsurprisingly) dictated by the LNA.

Figure 7-10: Response of analog frontend with improper gain selection settings for the input power. The loopback attenuation was −20 dB for both measurements. In (a), the VGA gain is set too high for the input power into the ADC buffer, resulting in lots of harmonic spurs. In (b), the VGA gain is set too low for the input power into the ADC buffer, resulting in lots of anharmonic spurs from sampling clock feedthrough.

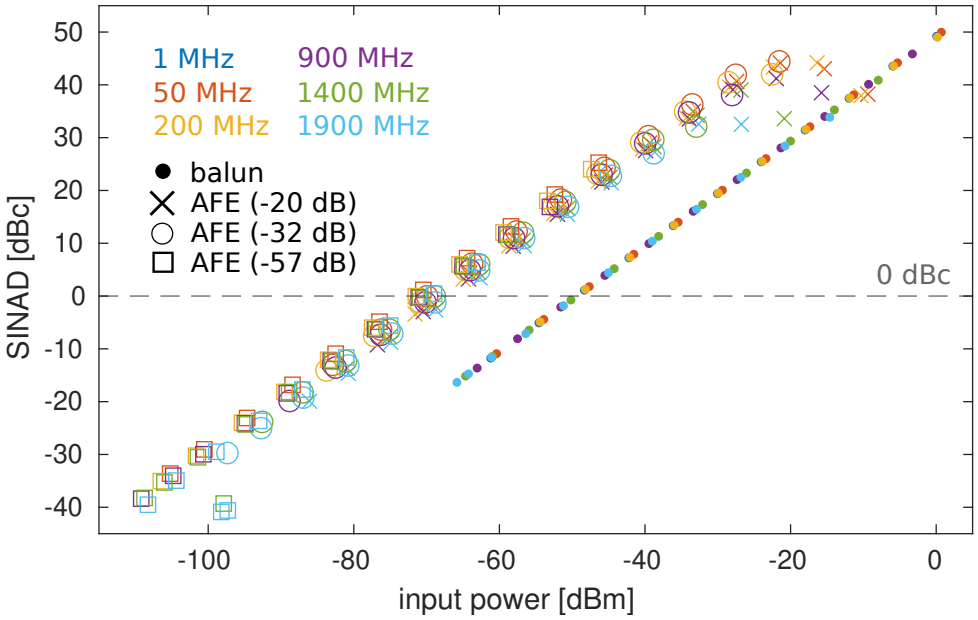

Figure 7-11: Plot of SINAD versus input power, comparing the analog frontend to the balun loopback. The loopback attenuation was varied between −20 dB, −32 dB, and −57 dB. The input power is referred to the input of the ADC buffer and calculated by dividing the measured signal power by the gain of the ADC buffer (as measured in section 7.2). From this plot, we can see that the analog frontend improves SINAD by about 20 dB as compared to the bare RFADC input. Detection of signals with 0 dB SINAD is possible for input powers as small as −70 dBm. The analog frontend performs about 10 dB worse than the theoretical limit as set by the receiver noise power. At high input powers, the LNA compresses, as can be seen for the measurements with 20 dB attenuation between the DAC buffer and ADC buffer. At very low input powers, the power spectral density of the signal is less than that of the noise, so the input power calculation is inaccurate.

ADC. This is an important metric for understanding if the analog frontend will be able to characterize low power superconducting nanowire electronics, where signal voltages may have amplitudes less than 100 µV. Figures 7-11 and 7-12 show the measured SINAD and SFDR for various input powers (referred to the input of the ADC buffer). Figure 7-11 suggests that the analog frontend can detect input signals as low as −70 dBm with 0 dBc SINAD (*i.e.* the signal will be on the same level as the noise). This input power would correspond to an rms signal voltage of 71 µV. However, if we turn our attention to Figure 7-12, we can see that an input power of −70 dBm has an SFDR of roughly 30 dB. Even though the signal power is the same as the total noise power, the signal is much more narrowband than the noise, so its power

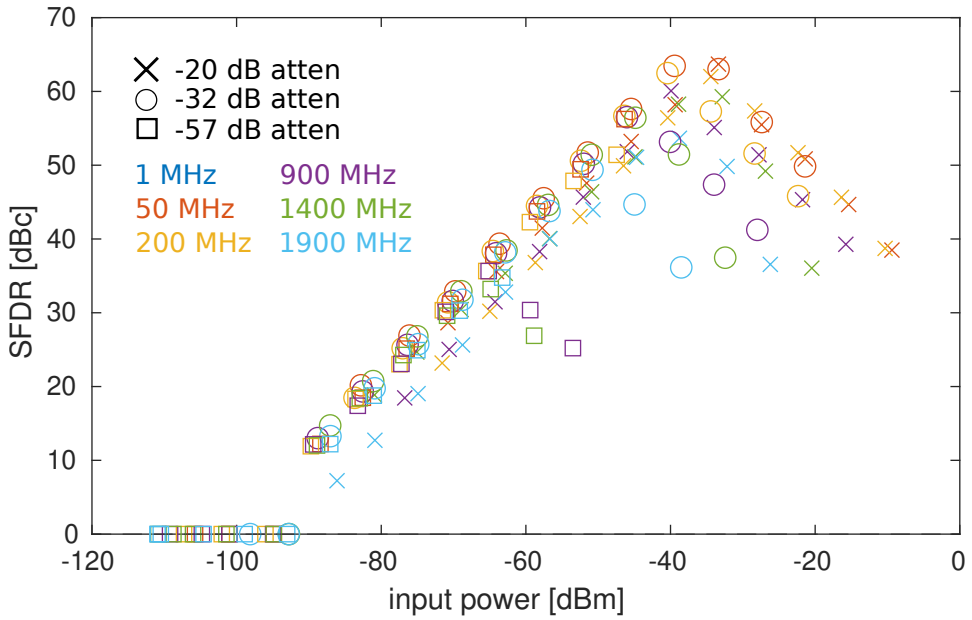

Figure 7-12: Plot of SFDR versus input power, comparing the analog frontend to the balun loopback. The loopback attenuation was varied between −20 dB, −32 dB, and −57 dB. If the frequency content of the input signal is well known, signals as small as −88 dBm can be detected. At high input powers, the LNA compresses and generates harmonic spurs which degrade the SFDR. At input power levels below −90 dBm, the signal is completely hidden by noise. Some outliers (particularly) for the −57 dB attenuation test can be seen; this is an artifact of the input power calculation, which assumes a linear response of the ADC buffer.

spectral density is actually well above the noise floor. This means that it is possible to detect signals with much lower powers than −70 dBm if they are bandlimited and their properties are well known. The transient waveform of a switching nanowire with a known bias current, load impedance, and L/R time constant can be very accurately modeled, which means that matched filtering can extend the sensitivity of detection [20]. The SFDR is above 10 dB for signal powers down to about −90 dBm, which suggests that it may be possible to detect switching events from nanowires with rms signal voltages of just 7 µV. Furthermore, if the signal power spectral density is below that of anharmonic spurs from sample clock feedthrough, but above the noise floor, the signal may still be recoverable with matched filtering.

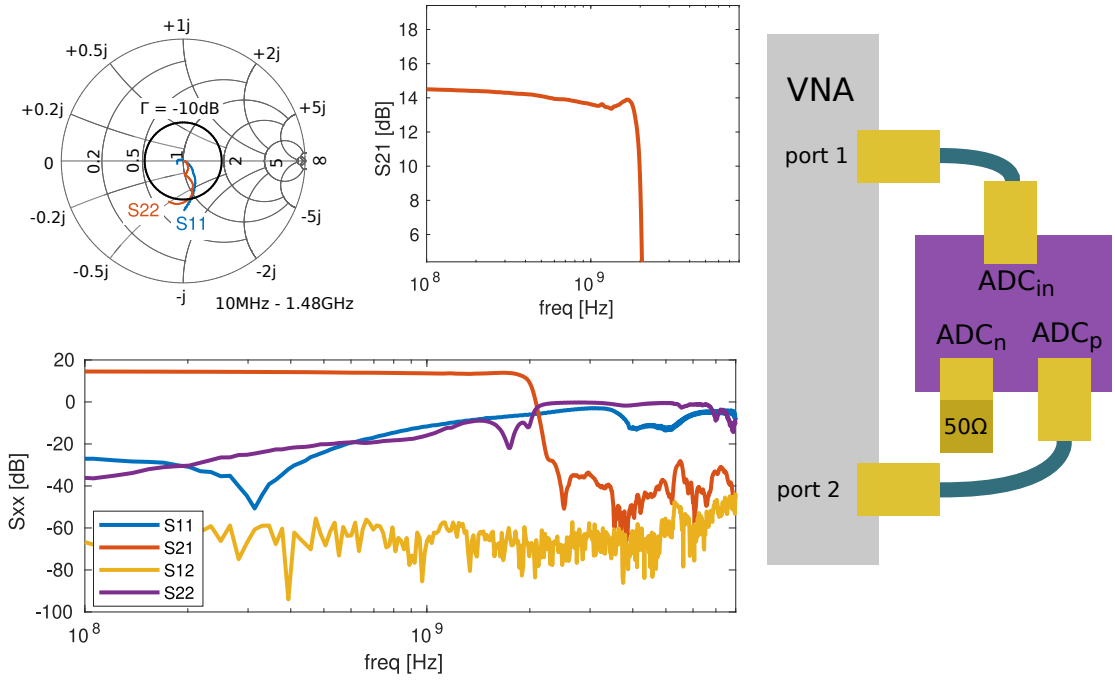

Figure 7-13: Small-signal scattering parameters of ADC buffer, measured with a two-port VNA using an input power of −30 dBm. One output of the buffer was terminated with 50 Ω and the other was connected to port 2 of the VNA. With the VGA at its lowest gain setting, the single-ended-to-differential gain is 20.5 dB.

7.2 Frequency-domain characterization

A VNA was used to measure the scattering parameters of the finished analog board. The input and output reflection coefficients S_{11} and S_{22}, reverse isolation S_{12}, and gain S_{21} were characterized. Ideally, an amplifier has low S_{11}, S_{22} and S_{12} combined with a large S_{21}. In general when measuring differential amplifiers with a VNA, either the VNA must have more than two ports, or baluns must be used. However, in this case the input impedance and gain actually do not change drastically when driven single-ended, provided the second input to the amplifier is terminated with a 50 Ω load.

Figures 7-13 and 7-14 show the test setup and measured scattering parameters of the ADC and DAC buffers (respectively). Figure 7-15 shows the setup and measured common-mode response of the DAC buffer. High common-mode rejection ratio (CMRR) is important for differential-to-single-ended amplifiers in general, but the RFSoC has quite good balance characteristics (*i.e.* there is no common-mode signal

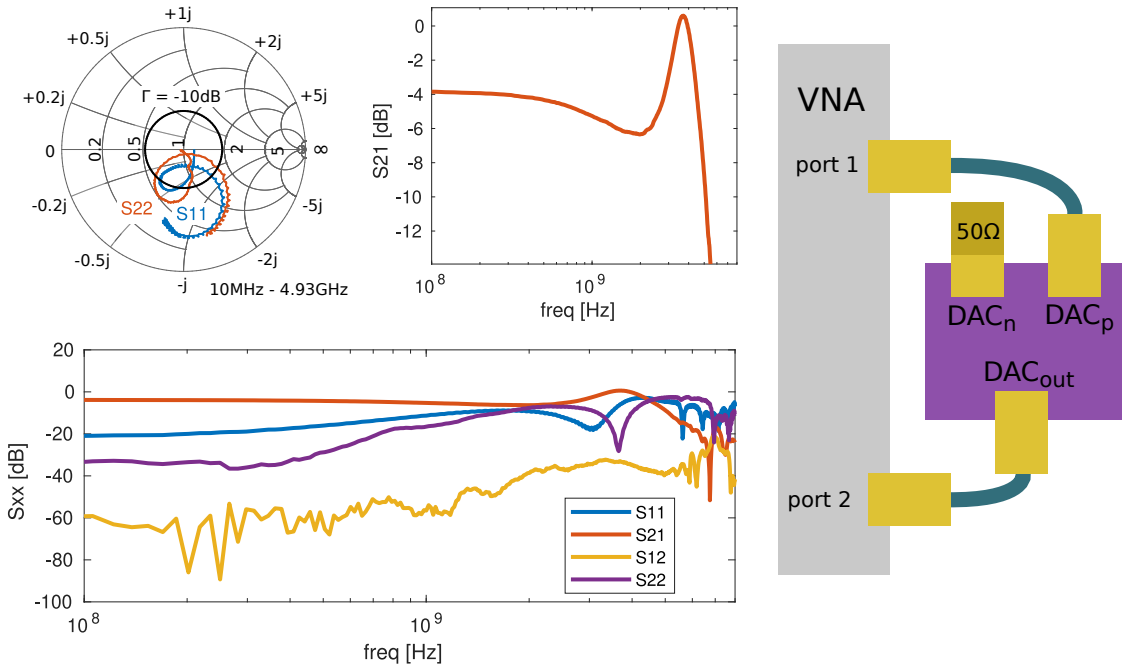

Figure 7-14: Small-signal scattering parameters of DAC buffer, driven single-ended. Measurement was performed with a two-port VNA using an input power of −30 dBm. One input of the buffer was terminated with 50 Ω and the other was connected to port 1 of the VNA. Due to the high CMRR as shown in Figure 7-15, the differential-to-single-ended gain is the same as the measured single-ended-to-single-ended gain. The gain of the buffer is −3.8 dB. There is about 1.5 dB rolloff before the gain peaking around 4 GHz, due to the internal compensation network of the FDA used in the buffer.

from the RFDAC output), so it is less critical in this specific case.

The single-ended gain of the ADC buffer was measured to be 14.5 dB, corresponding to a single-ended-to-differential gain of 20.5 dB. The expected gain of the buffer based on the component specifications is 22 dB, which is in good agreement with the measured gain. This measurement was performed with the VGA at its lowest gain setting, indicating that the maximum gain of the ADC buffer is 32.5 dB.

The single-ended gain of the DAC buffer was measured to be −3.8 dB (which is virtually the same as the differential-to-single-ended gain for this amplifier configuration).

As seen in the Smith charts in Figure 7-13 and 7-14, the input and output impedance are excessively capacitive. This is predominantly due to the intrinsic impedance mismatch of the VGA and FDA at high frequencies. Overall the reflection

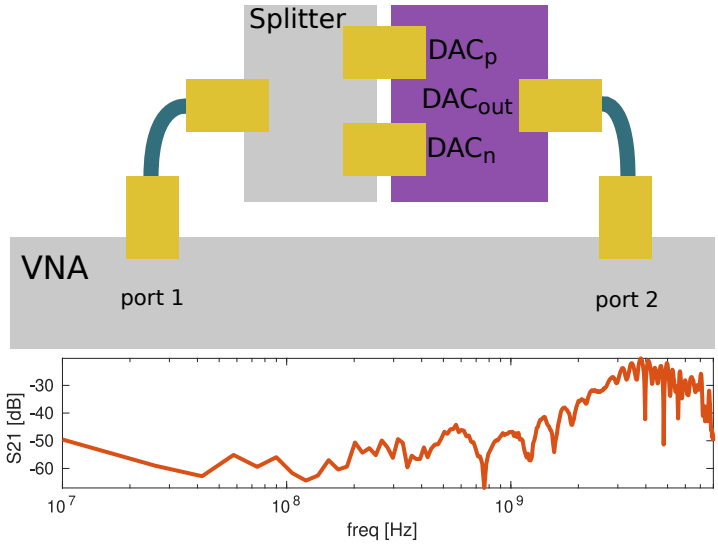

Figure 7-15: S_{21} of DAC buffer when driven with a $-30\,\text{dBm}$ common-mode input. The CMRR is better than 30 dB below 2.5 GHz.

coefficients for the ADC buffer are reasonably low, remaining below $-10\,\text{dB}$ for the most part. Note that the frequency scale only goes up to 1.48 GHz and that the output reflection coefficient S_{22} is very large in the stopband of the anti-alias filter. This is unsurprising due to the filter's construction as a cascade of LC filters. Furthermore, the high reflection coefficient in the filter stopband is perfectly acceptable, since there will not be any signal with frequency content in the filter stopband to reflect off of the impedance mismatch.

Unlike the ADC buffer, the high frequency reflection coefficients for the DAC buffer exceed $-10\,\text{dB}$ by quite a bit above 3.5 GHz, reaching a maximum of $-2.8\,\text{dB}$. This is likely not a problem for the RFDAC (which requires proper termination with a $100\,\Omega$ differential load), but if the buffer's S_{11} causes noticeable reflections, the board can be redesigned with an attenuation network before the amplifier.

The ripply behavior visible in the Smith chart is an artifact of imperfect calibration of the VNA.

Chapter 8

Future of testing of superconducting electronics with the RFSoC

Part II of this thesis presented the design and characterization of an analog frontend for testing superconducting nanowire electronics. Characterization of the analog frontend suggests it may be capable of detecting and amplifying signals with amplitudes as low as 70 µV or lower if a matched filter is used. This sensitivity will enable it to be used with an RFSoC to perform automated testing and characterization of large scale superconducting nanowire circuits. However, the analog frontend is only part of the proposed data acquisition setup.

The FPGA firmware required to enable testing of nTrons with the RFSoC has not been discussed much, despite being a potentially more substantial undertaking than the analog frontend design. However, due to similarities with [42], there is a very clear path forward for developing firmware that will enable fast, cost-effective test equipment for superconducting nanowires in the near term future. The three key parts of the firmware are stimulus generation, signal acquisition and processing, and device under test (DUT) state tracking.

The details of stimulus generation for testing nTron circuits are quite different from that of [42], however they are still straightforward. Ramp and pulse waveforms commonly used in testing of nTron-based circuits [1, 3, 18, 57] can readily be generated with very simple circuits that are easy to implement on an FPGA. Linear-feedback

shift registers are compact, simple-to-implement random number sources that can be used for generating pseudorandom stimulus for digital electronics [58].

Depending on the signal strength, signal acquisition may push the limits of what the RFSoC is capable of if matched filtering is necessary on all channels simultaneously. Matched filtering can be done on FPGAs [59], but can require hundreds of hardware multipliers for performing the filter convolution at full data-rate. The downsampling filters in the RFADC will be necessary if sufficiently many channels have low enough signal levels to require matched filtering.

Tracking the state of the DUT is relatively simple, but designing the interface a researcher would use to interact with the state tracker is a somewhat open-ended problem. In general though, the solution will require a state machine that describes the expected transitions of the internal DUT state (*e.g.* presence/absence or direction of circulating supercurrents) based on the provided stimuli. Each time a switching event is observed in the superconducting DUT, it would be compared with the expected switching event based on the state machine transition. The state machine must be reconfigurable through an easy to program interface that is accessable to the researcher using the data acquisition setup. It is crucial that the firmware is designed in such a way that researchers do not have to be experts in FPGA development to use it. This aspect of the design should lean heavily on the work in [42].

Appendix A

LTSpice symmetric nTron model code

```
************************************
* nTron                            *
************************************
.subckt ntron_symm g d s
+ params: width_g=12n
+ width_s=200n width_d=200n width_c=108n
+ sq_g=22 sq_d=200 sq_s=22 sq_c=10 thickness=4n
+ sheetRes=400 Tc=10.5 Tsub=4 Jc=46G C=1 A1=0.4
*Jc at T=Tsub in units of amps/m2
*sheetRes has units of ohms/sq
*Tc is critical temp and has units of K
*Tsub is substrate temp and has units of K
*C is constrition factor

*Critical currents
.PARAM Isw_g={Jc*width_g*thickness*C}
.PARAM Isw_d={Jc*width_d*thickness*C}
.PARAM Isw_s={Jc*width_s*thickness*C}
.PARAM Isw_c={Jc*width_c*thickness*C}
```

```
*Thermal Parameters
.PARAM kappa = 2.44e-8*Tc/(sheetRes*thickness);
* thermal conductivity W/m K
.PARAM heatCapacity = 4400
* heat capacity J/m^3 K
.PARAM hc = 50k
* thermal conductivity of surface W/m^2 K

*Electrical/Superconducting Parameters
.PARAM inductivity = {1.38p*sheetRes/Tc}
* H/square

.PARAM Lind_g = {inductivity*sq_g}
.PARAM Lind_d = {inductivity*sq_d}
.PARAM Lind_s = {inductivity*sq_s}
.PARAM Lind_c = {inductivity*sq_c}

.PARAM  minSquares = {1/sheetRes}
* # squares for minimum resistance
* Normal resistances
.PARAM Rnorm_g = {sheetRes*sq_g}
.PARAM Rnorm_d = {sheetRes*sq_d}
.PARAM Rnorm_s = {sheetRes*sq_s}
.PARAM Rnorm_c = {sheetRes*sq_c}

* units of ohms/square
.PARAM psi={sheetRes*(Jc*thickness)**2/(hc*(Tc-Tsub))}
* psi is the Stekly parameter.
.PARAM vo={1*sqrt(hc*kappa/thickness)/heatCapacity}
* vo is characteristic velocity
```

```
.PARAM Ihs_g={sqrt(2/psi)*Isw_g}
.PARAM Ihs_s={sqrt(2/psi)*Isw_s}
.PARAM Ihs_c={sqrt(2/psi)*Isw_c}
.PARAM Ihs_d={sqrt(2/psi)*Isw_d}

.PARAM Vthresh_g={minSquares*sheetRes*Ihs_g}
.PARAM Vthresh_s={minSquares*sheetRes*Ihs_s}
.PARAM Vthresh_d={minSquares*sheetRes*Ihs_d}
.PARAM Vthresh_c={minSquares*sheetRes*Ihs_c}

*Unitless Parameters
.PARAM delta={0.01}
* a small offset value for avoiding
* singularity in hotspot velocity

*Gate switching parameters
.PARAM beta = 12.82e-6

*********************************************************************************
** MAIN CIRCUIT GATE **
*********************************************************************************
* inductor
Lg g N_g_int Flux=((1-(v(N_g_res)/{Rnorm_g})+0.001)*{Lind_g}
+/(2*cos((2/3)*asin((0.6*abs(x))/{Isw_g}))-1))*x Rser = 1e-100

* hotspot resistor
B_Rg N_g_int center V=if(v(N_g_res)>0,v(N_g_res)*i(B_Rg),0)
* v(N_g_res) is resistance of hotspot
```

```
** S/C SENSE SUBCIRCUIT **
**Superconducting to Resistive Transition
*dependent source used to store state
B_g_state N_g_state 0 V=if((abs(i(Lg))>{Isw_g})
+|(abs(v(N_g_int)-v(center))>{Vthresh_g}),1,0)
R_g_state N_g_state 0 1
* v(N_g_state) is 0 if wire s/c, 1 if not

** HOTSPOT GROWTH INTEGRATOR SUBCIRCUIT **
B_g_hotspot N_g_res 0 V=sdt(if(v(N_g_state) & v(N_g_res)<{Rnorm_g},
+(2*sheetRes*vo/width_g)*
+(psi*(i(Lg)/{Isw_g})**2-2)
+/((sqrt((({psi}*(i(Lg)/{Isw_g})**2-1)
++abs({psi}*(i(Lg)/{Isw_g})**2-1))/2)+{delta})),0),
+0, V(N_g_state)<0.5)

*******************************************************************************
** MODULATION OF CHANNEL CRITICAL CURRENT **
*******************************************************************************
R_Isw_suppress N_Isw_channel 0 1e-6
B_Isw_suppress 0 N_Isw_channel I=if(abs(i(Lg))>{Isw_g}
+,{A1}*exp(-(abs(i(Lg))-{Isw_g})/{beta}),1)

*******************************************************************************
** CHANNEL HOTSPOT **
*******************************************************************************
* source/drain inductor
Ls s N_s_int Flux=((1.001-(v(N_s_res)/{Rnorm_s}))*{Lind_s}
+/(2*cos((2/3)*asin((0.6*abs(x))/{Isw_s}))-1))*x Rser = 1e-100
Ld d N_d_int Flux=((1.001-(v(N_d_res)/{Rnorm_d}))*{Lind_d}
```

```
+/(2*cos((2/3)*asin((0.6*abs(x))/{Isw_d}))-1))*x Rser = 1e-100
Lcs N_s N_cs Flux=(0.5*(1.001-(v(N_c_res)/{Rnorm_c}))*{Lind_c}
+/(2*cos((2/3)*asin((0.6*abs(x))/{Isw_c}))-1))*x Rser = 1e-100
Lcd N_d N_cd Flux=(0.5*(1.001-(v(N_c_res)/{Rnorm_c}))*{Lind_c}
+/(2*cos((2/3)*asin((0.6*abs(x))/{Isw_c}))-1))*x Rser = 1e-100

* channel hotspot resistors
B_Rcs N_cs center V=if(v(N_c_res)>0, v(N_c_res)/2*i(B_Rcs), 0)
B_Rcd N_cd center V=if(v(N_c_res)>0, v(N_c_res)/2*i(B_Rcd), 0)
B_Rs N_s N_s_int V=if(v(N_s_res)>0, v(N_s_res)*i(B_Rs), 0)
B_Rd N_d N_d_int V=if(v(N_d_res)>0, v(N_d_res)*i(B_Rd), 0)
* v(N_c_res) is resistance of channel hotspot

** S/C SENSE SUBCIRCUIT **
**Superconducting to Resistive Transition
*dependent source used to store state
B_c_state N_c_state 0 V=if((abs(i(Lcs))>{Isw_c}*abs(i(R_Isw_suppress)))
+|(abs(i(Lcd))>{Isw_c}*abs(i(R_Isw_suppress)))
+|(abs(v(N_cs)-v(N_cd))>{Vthresh_c}),1,0)
B_s_state N_s_state 0 V=if((abs(i(Ls))>{Isw_c}*abs(i(R_Isw_suppress)))
+|(abs(v(N_s)-v(N_s_int))>{Vthresh_s}),1,0)
B_d_state N_d_state 0 V=if((abs(i(Ld))>{Isw_c}*abs(i(R_Isw_suppress)))
+|(abs(v(N_d)-v(N_d_int))>{Vthresh_d}),1,0)
R_c_state N_c_state 0 1
R_s_state N_s_state 0 1
R_d_state N_d_state 0 1
* v(N_s_state) is 0 if wire s/c, 1 if not

** HOTSPOT GROWTH INTEGRATOR SUBCIRCUIT **
B_c_hotspot N_c_res 0 V=sdt(if((v(N_c_state)|v(N_s_state)|V(N_d_state))
```

```
+&v(N_c_res)<{Rnorm_c}, +(2*sheetRes*vo/width_c)*
+(psi*(max(abs(i(Lcs)),abs(i(Lcd)))/{Isw_c})**2-2)
+/((sqrt((({psi}*(max(abs(i(Lcs)),abs(i(Lcd)))/{Isw_c})**2-1)
++abs({psi}*(max(abs(i(Lcs)),abs(i(Lcd)))/{Isw_c})**2-1))/2)+{delta})),0),
+0, V(N_c_state)<0.5)

B_s_hotspot N_s_res 0 V=sdt(if(v(N_s_state)&v(N_c_res)>={Rnorm_c/2}
+&v(N_s_res)<{Rnorm_s},+(2*sheetRes*vo/width_s)*+(psi*(i(Ls)/{Isw_s})**2-2)
+/((sqrt((({psi}*(i(Ls)/{Isw_s})**2-1)
++abs({psi}*(i(Ls)/{Isw_s})**2-1))/2)+{delta})),0),0,V(N_s_state)<0.5)
B_d_hotspot N_d_res 0 V=sdt(if(v(N_d_state)&v(N_c_res)>={Rnorm_c/2}
+&v(N_d_res)<{Rnorm_d},+(2*sheetRes*vo/width_d)*+(psi*(i(Ld)/{Isw_d})**2-2)
+/((sqrt((({psi}*(i(Ld)/{Isw_d})**2-1)
++abs({psi}*(i(Ld)/{Isw_d})**2-1))/2)+{delta})),0),0,V(N_d_state)<0.5)

.ends ntron_symm
*$
```

Bibliography

[1] Alessandro Buzzi et al. "A nanocryotron memory and logic family". In: *Applied Physics Letters* 122.14 (Apr. 2023). 142601. ISSN: 0003-6951. DOI: 10.1063/5.0144686 (cit. on pp. 31, 47, 53, 94, 97, 145).

[2] Matteo Castellani et al. "A Nanocryotron Ripple Counter Integrated with a Superconducting Nanowire Single-Photon Detector for Megapixel Arrays". In: arXiv:2304.11700 (Apr. 23, 2023). DOI: 10.48550/arXiv.2304.11700. arXiv: 2304.11700[physics] (cit. on pp. 31, 47, 53).

[3] Kai Zheng et al. "A Superconducting Binary Encoder with Multigate Nanowire Cryotrons". In: *Nano Letters* 20.5 (May 13, 2020). Publisher: American Chemical Society, pp. 3553–3559. ISSN: 1530-6984. DOI: 10.1021/acs.nanolett.0c00498 (cit. on pp. 31, 145).

[4] A. N. McCaughan et al. "The thermally coupled imager: A scalable readout architecture for superconducting nanowire single photon detectors". In: *Applied Physics Letters* 121.10 (Sept. 2022). 102602. ISSN: 0003-6951. DOI: 10.1063/5.0102154 (cit. on pp. 31, 61).

[5] Risheng Cheng et al. "A 100-pixel photon-number-resolving detector unveiling photon statistics". In: *Nature Photonics* 17.1 (Jan. 2023), pp. 112–119. ISSN: 1749-4893. DOI: 10.1038/s41566-022-01119-3 (cit. on p. 31).

[6] Karl K Berggren et al. "A superconducting nanowire can be modeled by using SPICE". In: *Superconductor Science and Technology* 31.5 (Apr. 2018), p. 055010. DOI: 10.1088/1361-6668/aab149 (cit. on pp. 31, 37, 53–55).

[7] M. Castellani. "Design of Superconducting Nanowire-Based Neurons and Synapses for Power-Efficient Spiking Neural Networks". MA thesis. Politecnico di Torino, 2020. URL: https://webthesis.biblio.polito.it/id/eprint/15926 (cit. on pp. 31, 53–55).

[8] Tareq "Torque" El Dandachi. "Efficient simulation of Large-Scale Superconducting Nanowire Circuits". MA thesis. Massachusetts Institute of Technology, 2023. URL: https://hdl.handle.net/1721.1/150197 (cit. on pp. 31, 55, 59).

[9] Michael Tinkham. *Introduction to Superconductivity*. 2nd. Dover Publications, 2004. ISBN: 978-0-486-43503-9 (cit. on pp. 35, 36, 38).

[10] Anna Herr et al. "Scaling NbTiN-based ac-powered Josephson digital to 400M devices/cm^2". In: arXiv:2303.16792 (Apr. 29, 2023). DOI: 10.48550/arXiv.2303.16792. arXiv: 2303.16792[physics] (cit. on p. 36).

[11] Iman Esmaeil Zadeh et al. "Superconducting nanowire single-photon detectors: A perspective on evolution, state-of-the-art, future developments, and applications". In: *Applied Physics Letters* 118.19 (May 13, 2021), p. 190502. ISSN: 0003-6951. DOI: 10.1063/5.0045990 (cit. on p. 36).

[12] Adam N. McCaughan and Karl K. Berggren. "A Superconducting-Nanowire Three-Terminal Electrothermal Device". In: *Nano Letters* 14.10 (Oct. 2014), pp. 5748–5753. ISSN: 1530-6984. DOI: 10.1021/nl502629x (cit. on pp. 36, 37, 51).

[13] Reza Baghdadi et al. "Multilayered Heater Nanocryotron: A Superconducting-Nanowire-Based Thermal Switch". In: *Phys. Rev. Appl.* 14 (5 Nov. 2020), p. 054011. DOI: 10.1103/PhysRevApplied.14.054011 (cit. on p. 36).

[14] Adam N. McCaughan et al. "Using Geometry To Sense Current". In: *Nano Letters* 16.12 (Dec. 14, 2016). Publisher: American Chemical Society, pp. 7626–7631. ISSN: 1530-6984. DOI: 10.1021/acs.nanolett.6b03593 (cit. on p. 36).

[15] Aviv Glezer Moshe, Eli Farber, and Guy Deutscher. "Granular superconductors for high kinetic inductance and low loss quantum devices". In: *Applied Physics Letters* 117.6 (Aug. 11, 2020), p. 062601. ISSN: 0003-6951. DOI: 10.1063/5.0017749 (cit. on p. 36).

[16] 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). Publisher: American Physical Society, p. 034006. DOI: 10.1103/PhysRevApplied.11.034006 (cit. on p. 37).

[17] Adam N. McCaughan et al. "A kinetic-inductance-based superconducting memory element with shunting and sub-nanosecond write times". In: *Superconductor Science and Technology* 32.1 (Nov. 2018). Publisher: IOP Publishing, p. 015005. ISSN: 0953-2048. DOI: 10.1088/1361-6668/aae50d (cit. on p. 37).

[18] Kai Zheng et al. "Characterize the switching performance of a superconducting nanowire cryotron for reading superconducting nanowire single photon detectors". In: *Scientific Reports* 9.1 (Nov. 8, 2019). Number: 1 Publisher: Nature Publishing Group, p. 16345. ISSN: 2045-2322. DOI: 10.1038/s41598-019-52874-3 (cit. on pp. 39, 145).

[19] Thomas H. Lee. In: *The Design of CMOS Radio-Frequency Integrated Circuits*. 2nd ed. Cambridge University Press, 2004. Chap. 11. ISBN: 978-0-521-83539-8 (cit. on p. 41).

[20] Alan V. Oppenheim and George C. Verghese. In: *Signals Systems and Inference*. Global ed. Pearson Education, 2017. Chap. 13. ISBN: 978-1-292-15620-0 (cit. on pp. 41, 140).

[21] Chris H. Dick and Fred J. Harris. "Direct digital synthesis: some options for FPGA implementation". In: *Reconfigurable Technology: FPGAs for Computing and Applications*. Ed. by John Schewel et al. Vol. 3844. International Society for Optics and Photonics. SPIE, 1999, pp. 2–10. DOI: 10.1117/12.359534. URL: https://doi.org/10.1117/12.359534 (cit. on pp. 42, 129).

[22] Daniel Gomez-Garcia et al. “Linear chirp generator based on direct digital synthesis and frequency multiplication for airborne FMCW snow probing radar”. In: *2014 IEEE MTT-S International Microwave Symposium (IMS2014)*. 2014, pp. 1–4. DOI: 10.1109/MWSYM.2014.6848668 (cit. on p. 42).

[23] Jri Lee et al. “Design of 56 Gb/s NRZ and PAM4 SerDes Transceivers in CMOS Technologies”. In: *IEEE Journal of Solid-State Circuits* 50.9 (Sept. 2015). Conference Name: IEEE Journal of Solid-State Circuits, pp. 2061–2073. ISSN: 1558-173X. DOI: 10.1109/JSSC.2015.2433269 (cit. on p. 45).

[24] W. S. Boyle and G. E. Smith. “Charge Coupled Semiconductor Devices”. In: *Bell System Technical Journal* 49.4 (Apr. 1970), pp. 587–593. ISSN: 00058580. DOI: 10.1002/j.1538-7305.1970.tb01790.x (cit. on p. 45).

[25] S. Nagasawa et al. “New Nb multi-layer fabrication process for large-scale SFQ circuits”. In: *Physica C: Superconductivity*. Proceedings of the 21st International Symposium on Superconductivity (ISS 2008) 469.15 (Oct. 15, 2009), pp. 1578–1584. ISSN: 0921-4534. DOI: 10.1016/j.physc.2009.05.219 (cit. on p. 45).

[26] A.F. Kirichenko, O.A. Mukhanov, and A.I. Ryzhikh. “Advanced on-chip test technology for RSFQ circuits”. In: *IEEE Transactions on Applied Superconductivity* 7.2 (June 1997). Conference Name: IEEE Transactions on Applied Superconductivity, pp. 3438–3441. ISSN: 1558-2515. DOI: 10.1109/77.622124 (cit. on p. 45).

[27] Qing-Yuan Zhao et al. “A compact superconducting nanowire memory element operated by nanowire cryotrons”. In: *Superconductor Science and Technology* 31.3 (Feb. 2018), p. 035009. DOI: 10.1088/1361-6668/aaa820 (cit. on pp. 46, 97).

[28] Andrew J. Kerman et al. “Electrothermal feedback in superconducting nanowire single-photon detectors”. In: *Phys. Rev. B* 79 (10 Mar. 2009), p. 100509. DOI: 10.1103/PhysRevB.79.100509 (cit. on p. 53).

[29] Jens Timmer and Michel Koenig. “On generating power law noise.” In: *Astronomy and Astrophysics* 300 (1995), p. 707 (cit. on p. 59).

[30] 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 (cit. on p. 61).

[31] Emma E. Wollman et al. "Kilopixel array of superconducting nanowire single-photon detectors". In: *Optics Express* 27.24 (Nov. 25, 2019). Publisher: Optica Publishing Group, pp. 35279–35289. ISSN: 1094-4087. DOI: 10.1364/OE.27.035279 (cit. on p. 61).

[32] Shigeyuki Miyajima et al. "High-time-resolved 64-channel single-flux quantum-based address encoder integrated with a multi-pixel superconducting nanowire single-photon detector". In: *Optics Express* 26.22 (Oct. 29, 2018). Publisher: Optica Publishing Group, pp. 29045–29054. ISSN: 1094-4087. DOI: 10.1364/OE.26.029045 (cit. on p. 61).

[33] Emma E. Wollman et al. "Recent advances in superconducting nanowire single-photon detector technology for exoplanet transit spectroscopy in the mid-infrared". In: *Journal of Astronomical Telescopes, Instruments, and Systems* 7.1 (Jan. 2021). Publisher: SPIE, p. 011004. ISSN: 2329-4124, 2329-4221. DOI: 10.1117/1.JATIS.7.1.011004 (cit. on p. 63).

[34] Xu Tao et al. "A high speed and high efficiency superconducting photon number resolving detector". In: *Superconductor Science and Technology* 32.6 (Apr. 2019). Publisher: IOP Publishing, p. 064002. ISSN: 0953-2048. DOI: 10.1088/1361-6668/ab0799 (cit. on p. 65).

[35] Werner Herr and B Muratori. "Concept of luminosity". In: (2006). DOI: 10.5170/CERN-2006-002.361 (cit. on p. 65).

[36] Adam N. McCaughan et al. "PHIDL: Python-based layout and geometry creation for nanolithography". In: *Journal of Vacuum Science & Technology B* 39.6 (2021), p. 062601. DOI: 10.1116/6.0001203 (cit. on p. 67).

[37] Reed Foster. *qnn-layouts*. DOI: 10.5281/zenodo.7877875. URL: https://github.com/reed-foster/qnn-layouts (cit. on p. 67).

[38] Brenden A Butters. "Digital and Microwave Superconducting Electronics and Experimental Apparatus". PhD thesis. Massachusetts Institute of Technology, 2022. URL: https://hdl.handle.net/1721.1/143150 (cit. on pp. 72, 73).

[39] Reed Foster. *keysight-automation*. DOI: 10.5281/zenodo.7877877. URL: https://github.com/reed-foster/keysight-automation (cit. on p. 75).

[40] 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* 37.5 (Sept. 25, 2019), p. 051207. ISSN: 2166-2746. DOI: 10.1116/1.5119516 (cit. on pp. 92, 93).

[41] Biao Yang et al. "The State of the Art: Lightweight Cryocoolers Working in the Liquid-Helium Temperature Range". In: *Journal of Low Temperature Physics* 206.5 (Mar. 1, 2022), pp. 321–359. ISSN: 1573-7357. DOI: 10.1007/s10909-021-02664-x (cit. on p. 93).

[42] Leandro Stefanazzi et al. "The QICK (Quantum Instrumentation Control Kit): Readout and control for qubits and detectors". In: arXiv:2110.00557 (Mar. 10, 2022). DOI: 10.48550/arXiv.2110.00557. arXiv: 2110.00557[physics, physics:quant-ph] (cit. on pp. 102, 105, 145, 146).

[43] Kun Hee Park et al. "ICARUS-Q: Integrated control and readout unit for scalable quantum processors". In: *Review of Scientific Instruments* 93.10 (Oct. 6, 2022), p. 104704. ISSN: 0034-6748. DOI: 10.1063/5.0081232 (cit. on p. 102).

[44] Mats O. Tholén et al. "Measurement and control of a superconducting quantum processor with a fully integrated radio-frequency system on a chip". In: *Review of Scientific Instruments* 93.10 (Oct. 31, 2022), p. 104711. ISSN: 0034-6748. DOI: 10.1063/5.0101398 (cit. on p. 102).

[45] *High-Speed, Analog-to-Digital Converter Basics*. SLAA510. Texas Instruments. 2011 (cit. on p. 103).

[46] Adrian K. Sinclair et al. "Breaking the 10 mW/pixel Limit for Kinetic Inductance Detector Readout Electronics". In: arXiv:2305.00928 (May 1, 2023). DOI: 10.48550/arXiv.2305.00928. arXiv: 2305.00928[astro-ph] (cit. on p. 104).

[47] *Wideband Bias Tee Design Using 0402, SMD Components*. AN-2061. Rev. A. Analog Devices. 2020 (cit. on p. 105).

[48] *LMH5401 8-GHz, Low-Noise, Low-Power, Fully-Differential Amplifier*. SBOS710D. Revision D. Texas Instruments. 2018 (cit. on pp. 108, 111).

[49] Paul R. Gray et al. In: *Analysis and Design of Analog Integrated Circuits*. 5th ed. John Wiley & Sons, Inc., 2009. Chap. 12. ISBN: 978-0-470-24599-6 (cit. on p. 110).

[50] John Coonrod. "Circuit Materies and High-Frequency Losses of PCBs". In: *The PCB Magazine* (Feb. 2012), pp. 24–29 (cit. on p. 112).

[51] *UltraScale Architecture PCB Design*. UG583. v1.25. Xilinx. 2023 (cit. on p. 113).

[52] Robert W. Erickson and Dragan Maksimović. *Fundamentals of Power Electronics*. 3rd. Cham: Springer International Publishing, 2020. ISBN: 978-3-030-43879-1 978-3-030-43881-4. DOI: 10.1007/978-3-030-43881-4 (cit. on pp. 117, 118, 123).

[53] *Using the TPS6215x in an Inverting Buck-Boost Topology*. SLVA469D. Revision D. Texas Instruments. 2023 (cit. on pp. 118, 122).

[54] *TPS6213x 3-V to17-V, 3-A Step-Down Converter In 3-mm × 3-mm QFN Package*. SLVSAG7F. Revision F. Texas Instruments. 2021 (cit. on p. 121).

[55] *DDS Compiler v6.0 Product Guide*. PG141. Xilinx. 2021 (cit. on p. 130).

[56] Reed Foster. *rfsoc_dds*. DOI: doi.org/10.5281/zenodo.7948478. URL: https://github.com/reed-foster/rfsoc_dds (cit. on p. 130).

[57] Brenden A. Butters et al. "A scalable superconducting nanowire memory cell and preliminary array test". In: *Superconductor Science and Technology* 34.3 (Jan. 2021). Publisher: IOP Publishing, p. 035003. ISSN: 0953-2048. DOI: 10.1088/1361-6668/abd14e (cit. on p. 145).

[58] *PN Generators using the SRL macro*. XAPP211. v1.2. Xilinx. 2004 (cit. on p. 146).

[59] R. Andraka and A. Berkun. "FPGAs make a radar signal processor on a chip a reality". In: *Conference Record of the Thirty-Third Asilomar Conference on Signals, Systems, and Computers (Cat. No.CH37020)*. Conference Record of the Thirty-Third Asilomar Conference on Signals, Systems, and Computers (Cat. No.CH37020). Vol. 1. ISSN: 1058-6393. Oct. 1999, 559–563 vol.1. DOI: 10.1109/ACSSC.1999.832392 (cit. on p. 146).