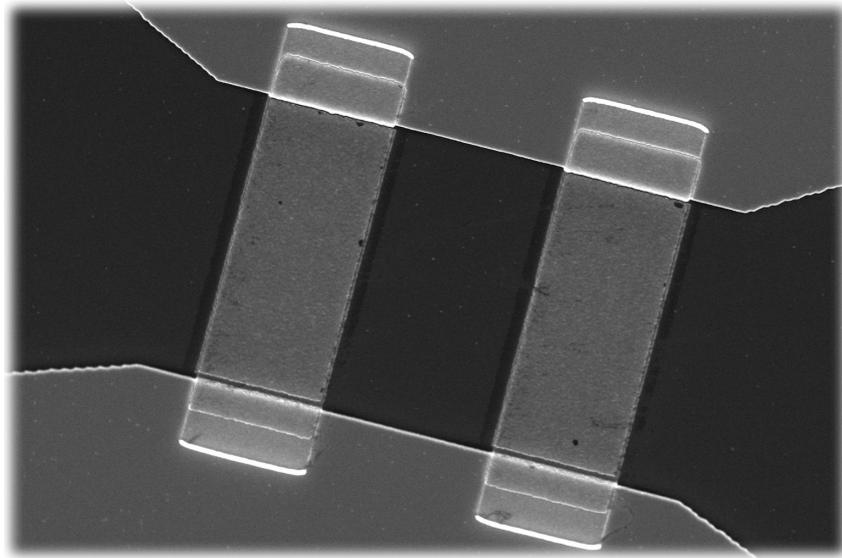


Fabrication and characterization of Josephson parametric amplifiers

by

M.H.J. de Jong
November 30, 2018



to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday October 17, 2018 at 09:00 AM.

Student number: 4221850
Project duration: January 16, 2018 – October 17, 2018
Thesis committee: Prof. Dr. G. A. Steele, TU Delft, supervisor
Prof. Dr. L. Kuipers, TU Delft
Dr. A. Endo, TU Delft
I. C. Rodrigues, TU Delft, daily supervisor
Dr. D. Bothner, TU Delft, daily supervisor

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

In this thesis, we have designed and fabricated a Josephson Parametric Amplifier (JPA) using a new double-angle evaporation method without a Dolan bridge. We have found and resolved several issues in the fabrication procedure, but it requires further tuning before being fully functional. We have also simulated the behaviour of a general parametric amplifier with an additional Duffing non-linear term, and found that this term appears to limit the oscillation amplitude. We have attempted to characterize Josephson junctions fabricated with the new double-angle evaporation procedure, but without much success. Using a different fabrication method, a JPA was made and successfully characterized. The maximum measured gain is 16 dB, with a bandwidth of 1 MHz. The noise temperature is comparable to the cryostat temperature of 250 mK, but it was not characterized accurately.

Contents

1	Introduction	1
2	Theory	3
2.1	Circuits	3
2.1.1	LC Oscillators	3
2.1.2	Microwave transmission lines	5
2.1.3	RLC oscillator capacitively side-coupled to a transmission line	7
2.2	Josephson junctions	9
2.2.1	Josephson relations	10
2.2.2	Josephson inductance	13
2.2.3	Flux quantization	13
2.2.4	SQUIDs	15
2.3	Parametric amplification	16
2.3.1	Parametric oscillation	17
2.3.2	From a parametric oscillator to a parametric amplifier	19
2.3.3	Modes of parametric amplification	21
2.3.4	JPA parameters	22
2.4	Noise and noise temperature	24
2.4.1	Thermal noise	24
2.4.2	Measuring noise temperatures	25
2.4.3	Noise temperature of an amplifier chain	25
3	Design, fabrication and simulation	27
3.1	Facilities	27
3.2	Design and fabrication	28
3.2.1	Junction and SQUID test design	29
3.2.2	Parametric amplifier design	30
3.2.3	Transmission cavity design	33
3.2.4	Josephson junctions	34
3.3	Fabrication issues	37
3.4	Measurement set-up and protocol	43
3.4.1	Room temperature Junction measurements	43
3.4.2	Fridge	44
3.4.3	Measurements	44
3.4.4	Printed circuit boards	45
3.5	Simulation of a general parametric amplifier	45
4	Results	51
4.1	Characterization of Josephson junctions	51
4.1.1	Room-temperature resistance	51
4.1.2	Critical current at 250 mK	53
4.1.3	SQUIDs in a magnetic field	55
4.2	Parametric Amplification	57

4.2.1	SQUID cavities	57
4.2.2	Biasing the SQUIDS	57
4.2.3	SQUID cavity power limits	58
4.2.4	Transmission amplitude versus pump strength	60
4.2.5	Gain versus phase delay	61
4.2.6	Noise Temperature	64
4.2.7	4-wave mixing	66
5	Conclusion	71
5.1	Josephson junctions	71
5.1.1	Fabrication	71
5.1.2	Characterization	72
5.2	JPA	73
5.2.1	Simulation	73
5.2.2	Characterization	73
A	JPA recipe	79

List of symbols

The list of symbols used in this thesis, their meaning and (if applicable) the units in which they are commonly expressed are used. Symbols denoted with a non-zero number subscript (e.g. ω_1, ω_2 , etc.) are to be taken as token symbols that do not have a set definition.

Symbol	Constant name	Value
e	Elementary charge	$1.60 \cdot 10^{-19} \text{ C}$
h	Planck's constant	$6.63 \cdot 10^{-34} \text{ m}^2\text{kg/s}$
\hbar	Reduced Planck's constant	$1.05 \cdot 10^{-34} \text{ m}^2\text{kg/s}$
k_B	Boltzmann constant	$1.38 \cdot 10^{-23} \text{ m}^2\text{kg/s}^2\text{K}$
Φ_0	Magnetic flux quantum	$2.07 \cdot 10^{-15} \text{ Wb}$
μ_0	Vacuum permeability	$4\pi \cdot 10^{-7} \text{ H/m}$

Table 1: Table of constants used in this thesis.

Symbol	Meaning	Unit
α	Parametric pump modulation amplitude (α_p normalized)	-
β	Duffing non-linearity parameter	-
Δ	Superconducting energy gap	J (eV)
	Detuning	Hz
ϵ_r	Relative dielectric permittivity	-
η	Quantum noise factor	-
θ	Angle (θ_0 offset, θ_{eq} equilibrium)	rad.
κ	Angular loss rate (κ_i internal, κ_e external)	Hz
λ_L	London penetration depth	m
μ	Energy eigenvalue of a superconductor	J (eV)
ρ	Density of charge-carrying particles in a superconductor	m^{-3}
	Resistivity	Ω/m^2
ϕ	Phase ($\Delta\phi$ Phase difference)	rad.
Φ	Magnetic flux (Φ_{ext} externally applied)	Wb
ψ	Wave function	-
ω	Angular frequency (ω_0 resonance, ω_d pump, ω_s signal, ω_i idler)	Hz
a	Growth rate	s^{-1}
a, b, d, y	Fit parameters (simple model)	-
$a_c, a_\omega, a_\phi, a_0$	Fit parameters (involved model)	-
A	Area	m^2
A, B, C, D	Transmission matrix formalism	-
A, A^*	Complex field (A_0 amplitude)	-
$A(t), B(t)$	Amplitude	-
\vec{A}	Magnetic vector field	Vs/m
B	Bandwidth	Hz
B	Magnetic field (B_{ext}^\otimes externally applied)	T
C	Capacitance (C_C coupling, C_1 per unit length, C_J Josephson, C_N Norton, C_{tot} total)	F
f_e	External drive	-
F	Force (F_0 amplitude)	N

Symbol	Meaning	Unit
G	Conductance (G_l per unit length)	S
\hat{H}	Hamiltonian	-
I	Current (I_0 amplitude, I_c critical, I_{screen} screening)	A
\vec{j}	Current density	A/m ²
K	Coupling parameter	-
l	Length	m
L	Inductance (L_l per unit length, L_J Josephson, L_{SQUID} of a SQUID, L_{geom} geometric)	H
m	Mass	g
n	number (n_p of charge carriers)	-
P	Power	W
q	charge	C
Q	Quality factor (Q_e external, Q_i internal)	-
$r(t)$	Amplitude	-
R	Resistance (R_l per unit length, R_N Norton, R_{tot} total, R_J Josephson)	Ω
S	Scattering parameter (S_{11} reflection, S_{21} transmission, S_{in} input)	-
t	time	s
T	Temperature (T_C critical, T_{input} input, $T_{\text{amplifier}}$ amplifier, T_{chain} chain)	K
V	Voltage	V
\vec{v}	velocity	m/s
x	position	m
X	Dummy variable	-
Z	Impedance (Z_0 characteristic, Z_{in} input, Z_{ext} external)	Ω

Table 2: List of symbols used in this thesis.

Introduction

This master thesis details work done on the subject of Josephson Parametric Amplifiers (JPAs) performed in the group of Prof. Dr. Gary Steele, in the section Quantum Nanoscience of the department of Applied Sciences of the TU Delft, between January and October 2018. First, a brief introduction to the concept of JPAs will be given, followed by a short motivation. An overview of the relevant theory is included in Chapter 2, the design, fabrication and measurement processes will be explained in Chapter 3, and we will also perform simulations of a general parametric amplifier here. Characteristic properties of the fabricated junctions and JPAs will be treated in Chapter 4, and finally conclusions will be drawn and a brief outlook on further research will be given in Chapter 5.

Amplifiers in general are common devices, and have been fundamental to the functioning of radios, mobile phones, music systems, etc. since their inception. A parametric amplifier in general requires the parameters of an oscillator (e.g. the length of a swing, the capacitance of an LC circuit, etc.) to be varied with a specific frequency. The effect of parametric amplification was first noticed by Faraday in 1831 [1].

Electric parametric amplifiers were developed between 1910 and 1920 (Fig. 1.1a) for use in telephone connections [2], and in the '50s and '60s the frequency range of parametric amplifiers was extended to the microwave regime [3, 4]. Nowadays, parametric amplifiers can reach optical frequencies [5] (Fig. 1.1b) and several types are commercially available.

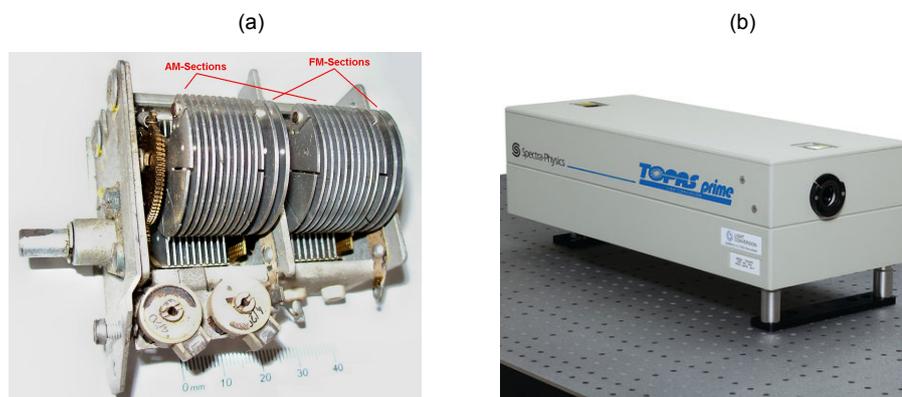


Figure 1.1: a) Parametric amplifier for radio signals. The plates form a variable capacitance [6] used for tuning the device frequency, the parametric variation comes from another element, potentially a variable-capacitance diode (varactor). b) Commercially available optical parametric amplifier [7].

The devices studied in this thesis are based on a Josephson junction, which is a structure made of two superconducting parts separated by a thin, non-superconducting barrier [8]. Here, two Josephson junctions in parallel form a Superconducting QUantum Interference Device, a SQUID [9], which functions as a tunable, non-linear inductor that is the key element of the Josephson Parametric Amplifier (JPA). It requires cryogenic temperatures to work, and it is operated in the microwave regime.

The value of the Josephson parametric amplifier lies in the fact that it can be a phase-sensitive amplifier, and potentially does not add any significant noise to the signal it is amplifying [10, 11]. This makes JPAs vital for experiments where weak signals are measured, and their applications for research are thus very broad. JPAs have been used in the search for dark matter [12, 13], for fast readout of qubits [14, 15] and for sensitive magnetometers [16]. JPAs can also be used to generate squeezed states of photons [17, 18]. Superconducting circuits at microwave frequencies also form an important part of quantum computation efforts [19], and since JPAs fit in the framework of superconducting quantum circuits, they could serve an important role in the communication between quantum computers.

The main motivation for this thesis lies in the fact that Josephson parametric amplifiers are not commercially available¹. Using JPAs to amplify the desired signals could decrease measurement times in various experiments. Also, conventional JPAs have limited bandwidth, and although tuning them is an option in some configurations, it would be more convenient if they were designed to operate at frequencies that fit with the experiments. The goal of this thesis is the design, fabrication and characterization of a Josephson parametric amplifier in the 6-8 GHz frequency range.

To realise this goal, a novel way of fabricating Josephson junctions is tested. With this method, it should be possible to create larger junctions, with higher critical currents, which can be used to create JPAs that work at higher signal strengths. Then, a design is made for a parametric amplifier. After fabrication of the design, the JPA will be tested and characterized in a dilution fridge. However, results of the fabrication show that the new method needs more fine-tuning before being ready for application. Using a different, established, method of fabricating junctions, a JPA was made, and subsequently characterized.

¹Optical parametric amplifiers are commercially available, but operate in a different frequency regime, and radio frequency amplifiers are as well, but they do not operate well at superconducting temperatures.

2

Theory

This chapter contains the theory necessary to understand the results of the experiments performed. As the focus of this thesis is experimental rather than theoretical, the theory will be kept brief and as general as possible. Sec. 2.1 describes circuits and their properties, Sec. 2.2 describes Josephson junctions. Sec. 2.3 is focussed on parametric amplification, and specifically on Josephson Parametric Amplifiers. Finally, Sec. 2.4 will treat noise and noise temperature of amplifiers.

2.1 Circuits

This section gives an overview of the relevant concepts and equations from circuit theory. Background and more detailed information can be found in textbooks such as [20]. Microwave-frequency circuits are of particular importance, and the interested reader is advised to look at microwave engineering textbooks such as [21].

2.1.1 LC Oscillators

One of the most essential electrical circuits is an LC oscillator. They are present in nearly all electronic equipment, e.g. in amateur radios such as Fig. 2.1a. In this work, we work with LC oscillators on the planar surface of a chip, such as shown in Fig. 2.1b. Despite its many forms, the basic properties of an LC oscillator remain the same and it is generally depicted schematically by the circuit diagram of Fig. 2.2a (top). It contains a capacitor with capacitance C and an inductor with inductance L . The current through the oscillator is sinusoidal in time,

$$I = I_0 \cos(\omega_0 t + \phi), \quad (2.1)$$

with an amplitude I_0 , frequency ω_0 and phase offset ϕ . In the ideal case, it is an oscillator that dissipates no energy: Once it is oscillating, it will continue oscillating forever at its resonance frequency. This is also called the natural frequency of a LC oscillator, and it is given by [20]

$$\omega_0 = \frac{1}{\sqrt{LC}}. \quad (2.2)$$

In the more realistic case, there will be some resistance R in the circuit. The non-ideal (dissipative) parts of the circuit components are often grouped in a single resistor. For the LC oscillator, this leads to a RLC oscillator, shown schematically in Fig. 2.2a (bottom). An RLC oscillator will have the same natural frequency as the LC oscillator, but without a force driving it, the oscillation amplitude will decay over time, as shown in Fig. 2.2b.

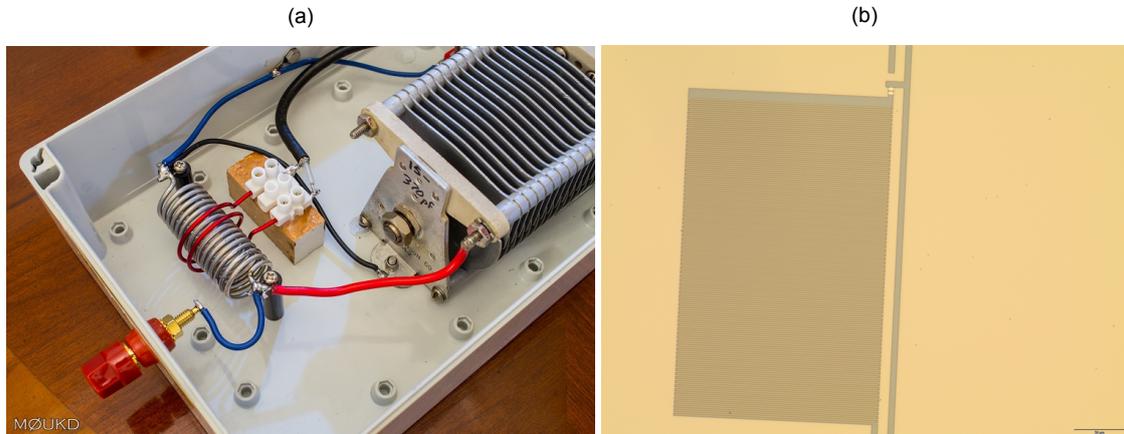


Figure 2.1: a) LC oscillator used for amateur radio. Clearly visible are the capacitor plates on the right, and the coil on the left. The inductance of the coil is made tunable via a second coil (in red) around the main coil. This tunability allows a radio to 'tune in' to different stations [22]. b) Microscope image of an LC oscillator fabricated on a chip. The washboard-like structure is an interdigitated capacitor.

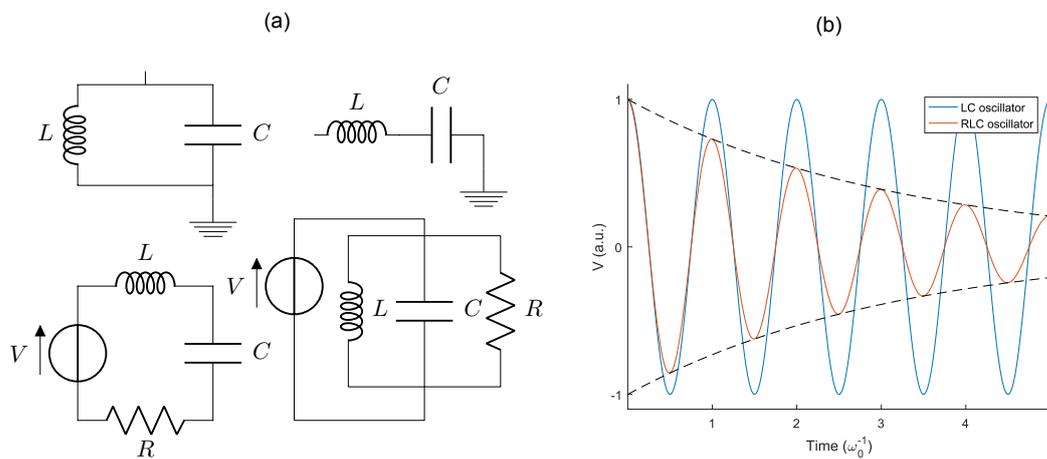


Figure 2.2: a) Circuit diagrams of a parallel (top left) and series (top right) LC oscillator, and a parallel (bottom right) and series (bottom left) RLC oscillator. b) Amplitude of voltage between the capacitor plates over time. For similar initial conditions, both an LC and RLC oscillator show oscillating behaviour (both voltage and current), but the resistor in the RLC set-up damps the oscillations. The dashed lines trace the decay of the oscillation amplitude.

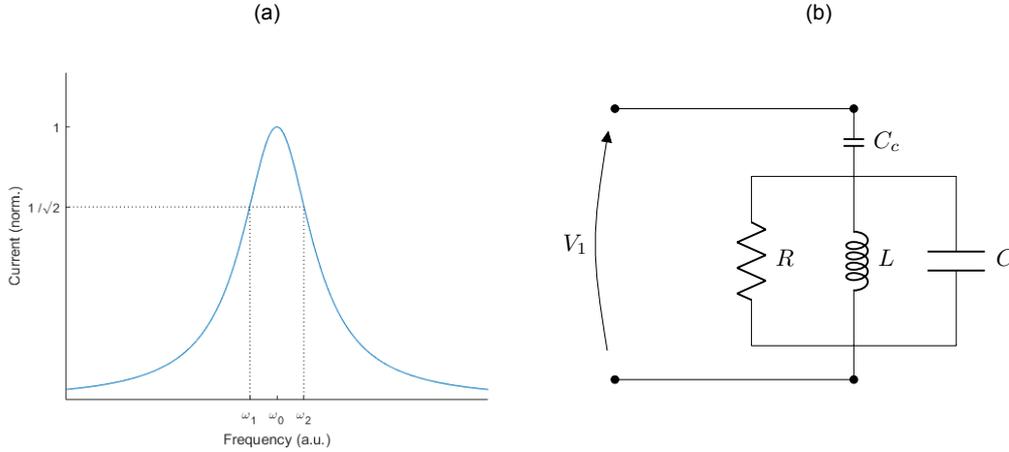


Figure 2.3: a) Oscillation amplitude for an RLC oscillator, with the half-power frequencies ω_1 and ω_2 denoting the bandwidth, and with it the Q-factor. b) Circuit diagram of a parallel RLC circuit capacitively coupled to a line with a potential difference V_1 between the two ports. Instead of directly connecting the circuit the line, it is coupled by a coupling capacitor, C_c

From LC to RLC: Q-factor and losses

The resistor in the RLC oscillator does not change the resonance frequency ω_0 , but it affects the system by introducing losses. A common measure of these losses is the quality factor, referred to as the Q-factor [23]. It is given by the range of driving frequencies where the system will oscillate above a certain amplitude (relative to the amplitude exactly at resonance), normalized to the resonance frequency. The conventional cut-off for the oscillation amplitude is the half-power point: the frequencies where the oscillations contain half the power of the oscillations at resonance, shown schematically as ω_1 and ω_2 in Fig. 2.3a. Note that half the power corresponds to $\frac{1}{\sqrt{2}}$ times the current.

The band of frequencies between these two end-points is called the bandwidth B . The quality factor Q is defined as

$$Q = \frac{\omega_0}{\omega_2 - \omega_1} = \frac{\omega_0}{B}. \quad (2.3)$$

We can also relate this to the intensity decay rate, κ , the conventionally used parameter in (microwave-)optics systems such as the ones studied in this thesis, via

$$Q = \frac{\omega_0}{\kappa}. \quad (2.4)$$

Instead of directly driving oscillators (as in Fig. 2.2a) with a voltage or current source in the circuit, we couple oscillators to a transmission line to the outside of the cryostat, as shown in Fig. 2.3b. In the next sections, we will treat these transmission lines and how to couple a parallel RLC circuit to one.

2.1.2 Microwave transmission lines

We will make use of microwave-frequency signals to drive and probe our samples. These electromagnetic waves with a frequency typically between 300 MHz and 300 GHz require consideration in design of the circuit elements [21], because the electrical length of a microwave signal can be comparable to or smaller than the length of the circuit elements. Thus, for example, the voltage might not be constant through a wire, even if it is a perfect conductor [21].

Another issue to take into consideration is the radiation from unshielded wires. Losses caused by electromagnetic field radiating out of the wires can be avoided by using two wires of opposite polarity next to each other. This causes the fields to cancel, and avoids most of the losses. Having these two wires in a certain geometry next to each other is done in microwave transmission lines.

Microwave transmission lines serve to transmit the signals from point A to point B, similar to the wires

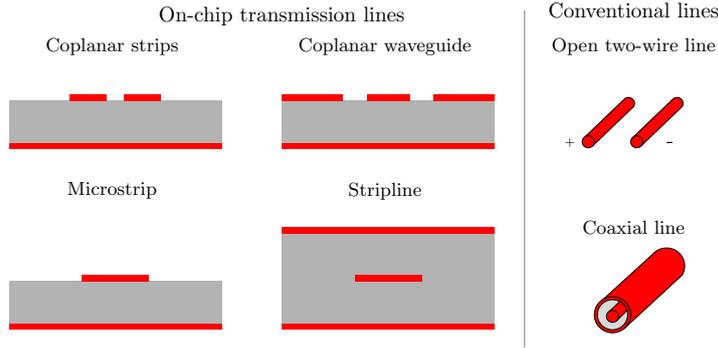


Figure 2.4: Schematic of several microwave transmission line types. Red represents a metallic conductor, gray represents a dielectric medium. The bottom red plane (and also the top plane for the stripline, and the side-planes for the coplanar waveguide) are the ground planes, the other red areas are the lines carrying the signal.

in a conventional circuit. In this thesis, they are present in two different geometries: In planar, on-chip geometries to connect our fabricated structures, and in open geometries to connect to the microwave sources and measurement apparatus. A schematic overview of some common geometries is given in Fig. 2.4. The coaxial (shortened to coax) line and the coplanar waveguide will be the ones used in this thesis.

Due to the fact that the electrical signal wavelength is shorter than the transmission line length, we can not describe the transmission line as simple (ideal) wire. Instead, we must take an approximation per unit length of the transmission line. This can be done using four parameters in a lumped-element model as shown in Fig. 2.5: Resistance per unit length, R_1 (Ω/m), conductance per unit length, G_1 (S/m), inductance per unit length, L_1 (H/m) and capacitance per unit length, C_1 (F/m) [21]. These four parameters are used to describe the characteristic impedance of a transmission line,

$$Z_0 = \sqrt{\frac{R_1 + i\omega L_1}{G_1 + i\omega C_1}}, \quad (2.5)$$

which is 50Ω for a typical coaxial cable. The characteristic impedance is the potential difference divided by the current for a cable of infinite length, so its unit is Ω instead of Ωm . For high frequency (typically above 10 kHz [24]) and a good conductor, the characteristic impedance of a transmission line can be approximated as

$$Z_0 = \sqrt{\frac{L}{C}}. \quad (2.6)$$

This is the case for the commercial coax lines used for the measurements performed in this work. Besides commercially available coax lines to connect our sources to our sample and back to the measurement equipment, we must also transmit signals from the edge of our sample to systems under study (in the middle). For this, we use coplanar waveguides.

For us, the most important property of a coplanar waveguide is that its characteristic impedance must match that of the rest of the microwave transmission lines, to prevent reflections at the interfaces. To ensure this, it is necessary to design the waveguide so that it has the right capacitance and inductance per unit length. The capacitance per unit length of a coplanar waveguide can be calculated using conformal mapping, and involves elliptical integrals [25]. There are tools available to calculate this for certain geometries, e.g. AppCAD. The coplanar waveguides used in this thesis have been simulated (Sec. 3.2.2) to obtain the desired characteristic impedance of 50Ω .

We are interested in coupling an RLC-oscillator to a transmission line. There are three ways in which we can couple a circuit: Directly coupled, capacitively coupled, and inductively coupled. In the capacitively coupled case, the transmission line is coupled to our circuit through a capacitor. In the inductively coupled case, the transmission line is coupled through an inductor, and in the directly coupled case it is connected through a wire. In our designs, we will use all three methods.

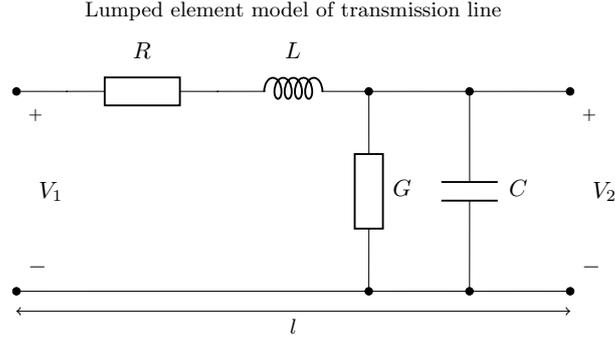


Figure 2.5: A lumped-element model of a typical transmission line of unit length l . Generally, this is designed to be 50Ω , to prevent unwanted reflections at interfaces between different transmission lines.

2.1.3 RLC oscillator capacitively side-coupled to a transmission line

For the rest of this thesis, a lot of the systems we will use will have the form of an RLC oscillator capacitively side-coupled to a transmission line, shown schematically in Fig. 2.6a. We are interested in finding out how the circuit will behave if we send in a certain signal (e.g. a changing voltage, or a changing current). Suppose we send in a signal through the leftmost port (transmission line) to the oscillator. Part of this signal could be reflected back, so we can measure it at the left port. Another part could be transmitted, so we measure it at the right port. Generally, this is described in terms of the fraction of the power of a signal that is reflected or transmitted. For a two-port system (i.e. a system connected by two transmission lines), reflected power fraction is denoted by S_{11} , and the transmitted power fraction is denoted by S_{21} , which are called the scattering parameters [21].

To obtain these scattering parameters for our system (the capacitively side-coupled RLC oscillator), we first translate the transmission lines into circuit elements. The new circuit is shown in Fig. 2.6b. At the top, we have translated the transmission lines into resistors with impedance Z_0 . The dashed lines indicate the border between the RLC oscillator (internal) and the transmission line (external) part of the circuit. We treat the system as two oscillators, each with a separate loss rate and quality factor. The loss rate of RLC circuit itself is called the internal loss rate, κ_i , and losses in the coupling to the outside environment are called the external loss rate, κ_e . The total loss rate is given by

$$\kappa = \kappa_i + \kappa_e. \quad (2.7)$$

Similarly, we can split the total Q-factor of the circuit into an external, Q_e , and internal part, Q_i .

To simplify the analysis of the circuit drawn in Fig. 2.6b, we can replace the external part of the circuit by its Norton equivalent [20], with resistor R_N and capacitor C_N . The circuit diagram for this is shown at the bottom of Fig. 2.6b.

First, we want to know R_N and C_N in terms of the transmission line impedance Z_0 and coupling capacitor C_C . We combine the impedance of these last two into the external impedance, Z_{ext} ,

$$Z_{\text{ext}} = \left(\frac{1}{Z_0} + \frac{1}{Z_0} \right)^{-1} + \frac{1}{i\omega C_C} = \frac{Z_0}{2} + \frac{1}{i\omega C_C}. \quad (2.8)$$

We can also determine the external impedance of the Norton-equivalent part of the circuit,

$$\frac{1}{Z_{\text{ext}}} = \frac{1}{R_N} + i\omega C_N. \quad (2.9)$$

If we combine these two equations, we can obtain expressions for R_N and C_N . We can simplify these using the approximations $\omega \approx \omega_0$ and $\omega C_C \approx Z_0$, and get

$$\begin{aligned} R_N &= \frac{4 + \omega^2 C_C^2 Z_0^2}{2\omega^2 C_C^2 Z_0} \approx \frac{2}{\omega_0 C_C^2 Z_0} \\ C_N &= \frac{4C_C}{4 + \omega^2 C_C^2 Z_0^2} \approx C_C. \end{aligned} \quad (2.10)$$

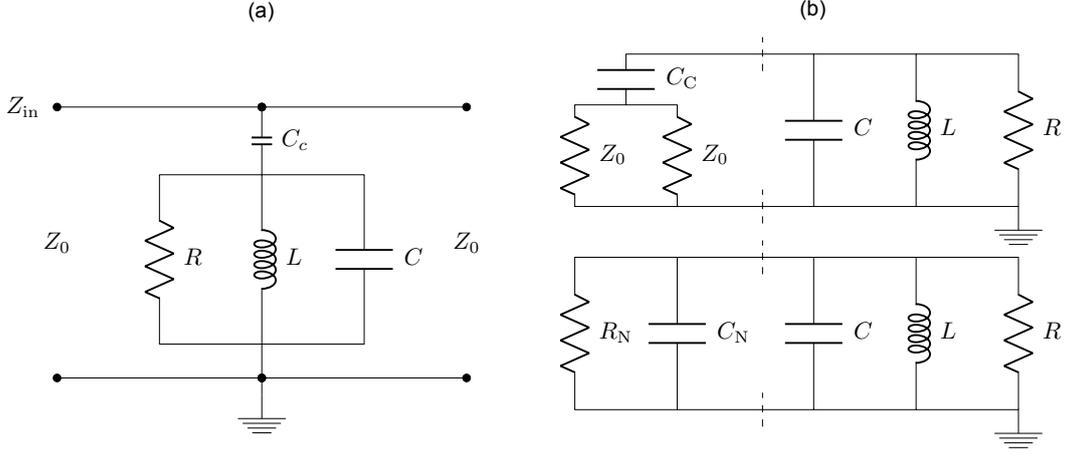


Figure 2.6: a) Circuit diagram of an RLC oscillator capacitively side-coupled to two transmission lines with impedance Z_0 . b) Equivalent circuits for the RLC oscillator capacitively side-coupled to a transmission line. At the top, we have replaced the transmission line with a resistor of impedance Z_0 , where we have two of them in parallel due to the two transmission lines. At the bottom, we have taken the Norton equivalent circuit, a resistor with resistance R_N and a capacitor with capacitance C_N . The dashed lines indicate the separation between the RLC oscillator (right) and the transmission line (left) part of the circuit.

With these, we can calculate the total resistance and capacitance of our circuit, and also find its resonance frequency,

$$\begin{aligned}
 R_{\text{tot}} &= \frac{RR_N}{R + R_N} \\
 C_{\text{tot}} &= C + C_C \\
 \omega_0 &= \frac{1}{\sqrt{L(C + C_C)}}.
 \end{aligned} \tag{2.11}$$

With the total resistance, capacitance and resonance frequency of the circuit, we can determine the total quality factor, Q , and separate it into an internal and external quality factor,

$$\begin{aligned}
 Q &= \omega_0 R_{\text{tot}} C_{\text{tot}} = (Q_i^{-1} + Q_e^{-1})^{-1} \\
 Q_i &= \omega_0 R(C + C_C) \\
 Q_e &= \frac{2(C + C_C)}{\omega_0 C_C^2 Z_0}.
 \end{aligned} \tag{2.12}$$

We can similarly also find expressions for the internal and external loss rate,

$$\begin{aligned}
 \kappa_i &= \frac{1}{R(C + C_C)} \\
 \kappa_e &= \frac{\omega_0^2 C_C^2 Z_0}{2(C + C_C)}.
 \end{aligned} \tag{2.13}$$

Before we can proceed to the scattering parameters, we only need to find the input impedance Z_{in} . This is the impedance between the input transmission line (as denoted in Fig. 2.6a) and the ground.

The input impedance of a simple RLC circuit, as in the bottom left of Fig. 2.2a, can be found from the impedances of all of the components in parallel

$$\frac{1}{Z_{\text{in}}} = \frac{1}{R} + \frac{1}{i\omega L} + i\omega C. \tag{2.14}$$

Then, we can add a coupling capacitor, as in Fig. 2.3b, which we add separately from the other terms in the input impedance, because it is in impedance,

$$Z_{\text{in}} = \frac{R}{1 + iR(\omega C - \frac{1}{\omega L})} + \frac{1}{i\omega C_C}. \tag{2.15}$$

If we take the transmission line into account as a resistor, the calculations become more complicated. However, we can achieve an approximate result by performing a Taylor expansion around $\omega \approx \omega_0$ and performing a trick using complex frequency [21] to take into account the losses. We obtain for the RLC oscillator capacitively coupled to a transmission line

$$Z_{\text{in}} \approx \frac{L(C + C_C)^2}{C_C^2} (\kappa_i + 2i(\omega - \omega_0)) \quad (2.16)$$

Now, we can add the second transmission line in parallel and obtain the circuit of Fig. 2.6b (top). This second transmission line can be added as an additional impedance in parallel,

$$\frac{1}{Z_{\text{in}}} = \frac{C_C^2}{L(C + C_C)^2 (\kappa_i + 2i(\omega - \omega_0))} + \frac{1}{Z_0} \quad (2.17)$$

or, rewriting this,

$$Z_{\text{in}} = Z_0 \frac{\kappa_i + 2i(\omega - \omega_0)}{2\kappa_e + \kappa_i + 2i(\omega - \omega_0)} \quad (2.18)$$

Now, we are ready to find the scattering parameters, S_{11} and S_{21} , that we are interested in. We can do this using the ABCD-matrix for a two-port system [21],

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \frac{1}{Z_{\text{in}}} & 1 \end{pmatrix}. \quad (2.19)$$

The scattering parameters are

$$\begin{aligned} S_{11} &= \frac{A + \frac{B}{Z_0} - CZ_0 - D}{A + \frac{B}{Z_0} + CZ_0 + D} = \frac{-\frac{Z_0}{Z_{\text{in}}}}{2 + \frac{Z_0}{Z_{\text{in}}}} = \frac{-\kappa_e}{\kappa_i + \kappa_e + 2i(\omega - \omega_0)} \\ S_{21} &= \frac{2(AD - BC)}{A + \frac{B}{Z_0} + CZ_0 + D} = \frac{2}{2 + \frac{Z_0}{Z_{\text{in}}}} = \frac{\kappa_i + 2i(\omega - \omega_0)}{\kappa_i + \kappa_e + 2i(\omega - \omega_0)}. \end{aligned} \quad (2.20)$$

Now that we have the scattering parameters, we can gain insight in how we expect our system to react when we send in a certain signal. For a frequency range near the resonance frequency, we have plotted S_{11} and S_{21} for several combinations of κ_i and κ_e in Figs. 2.7a and 2.7b. These combinations describe how well the RLC circuit (in our case) is coupled to the environment. The combinations are:

- Undercoupled: Most of the loss happens inside the RLC oscillator, so $\kappa_e < \kappa_i$, and $Q_e > Q_i$.
- Critically coupled: The losses inside the RLC oscillator equal the losses to the outside of the oscillator, so $\kappa_e = \kappa_i$, and $Q_e = Q_i$.
- Overcoupled: Most of the loss happens to the outside of the RLC oscillator, so $\kappa_e > \kappa_i$, and $Q_e < Q_i$.

Concluding this section, we have described LC and RLC oscillators, as well as transmission lines. We have detailed how to couple these oscillators to a transmission line, and derived what we can expect to measure when we send in a certain signal through one of the transmission lines.

2.2 Josephson junctions

A Josephson junction consists of two superconducting parts separated by a thin non-superconducting (insulator or normal conductor) part such that the Cooper pairs in the superconductor can tunnel between the two parts, as shown schematically in Fig. 2.8a. The relation between current and voltage of such a device were predicted by B.D. Josephson in 1962 [8], which earned him the 1973 Nobel prize together with L. Esaki and I. Giaever, who had earlier reported observations of Cooper pairs tunneling between the superconducting parts. Confirmation of the Josephson effect was done in 1963 by Anderson and Rowell [26].

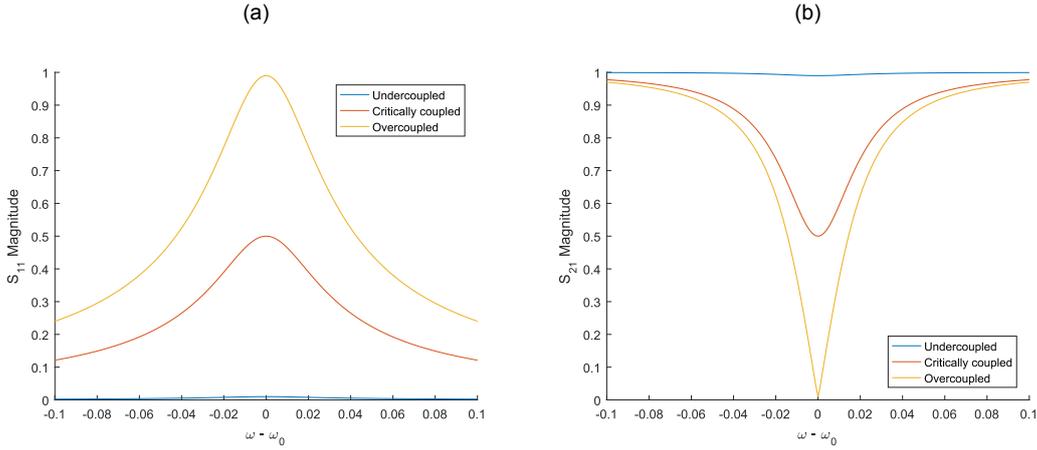


Figure 2.7: a) Reflected power coefficient S_{11} for various frequencies near the resonance frequency of the RLC oscillator. The reflected power shows a peak near the resonance frequency. Explanations of undercoupled, critically coupled and overcoupled are given in the main text. b) Transmitted power coefficient S_{21} for various frequencies near the resonance frequency of the RLC oscillator. The transmitted power shows a dip near the resonance frequency.

2.2.1 Josephson relations

Josephson described the current $I(t)$ through and voltage $V(t)$ across the junction as dependent on the phase difference $\Delta\phi$ of the wave functions of the two superconducting parts [8], as sketched in Fig. 2.8a,

$$\begin{aligned} I(t) &= I_c \sin(\Delta\phi) \\ V(t) &= \frac{\hbar}{2e} \frac{\partial \Delta\phi}{\partial t}. \end{aligned} \quad (2.21)$$

These are commonly referred to as the Josephson equations. When the phase difference is constant ($V = 0$), there can be a supercurrent through the junction, consisting of tunnelling Cooper pairs [8]. This supercurrent will not lead to an increase in voltage, and the junction will have zero resistance. However, above a certain critical current, I_c , which is an important junction parameter [27], the junction will behave Ohmic. This behaviour is seen in the slope of the I-V (current-voltage) curve in Fig. 2.8b for sufficiently large voltage, $|V| > \frac{2\Delta}{e}$, where Δ is the superconductor energy gap, is related to quasiparticle tunneling [28].

The anomalous current at $V = 0$ is called the DC (Direct Current) Josephson effect, which is one of three effects predicted by Josephson [8]. The other two are the AC- (Alternating Current) and inverse AC Josephson effects, which allow a Josephson junction to function as respectively a voltage-to-frequency and a frequency-to-voltage converter [27]. The AC Josephson effect happens when a constant non-zero voltage $V \neq 0$ is put across the junction. The phase difference $\Delta\phi$ will vary linearly over time, which results in an AC current with amplitude I_c and frequency $\frac{2e}{\hbar}V$. The inverse AC Josephson effect occurs when a Josephson junction is driven with a combination of an AC and a DC current, which generates a set of voltage steps in the I-V curve, the so-called Shapiro steps [29].

Josephson junctions come in different geometries and types, several of which are shown schematically in Fig. 2.9. All of these have in common that some type of barrier separates the superconductors. This can be a constriction (as in a) and e) in the figure), an insulating or barrier layer (as in c) and d)), or a boundary junction (as in b) or f)) [30]. In this thesis, all fabricated junctions consist of the superconductor-insulator-superconductor type.

Derivation of the Josephson relations

We can derive the Josephson relations from more fundamental concepts. Here, we follow the example of [28] and [27]. The starting point of the wave function of a superconductor. In general, this takes the

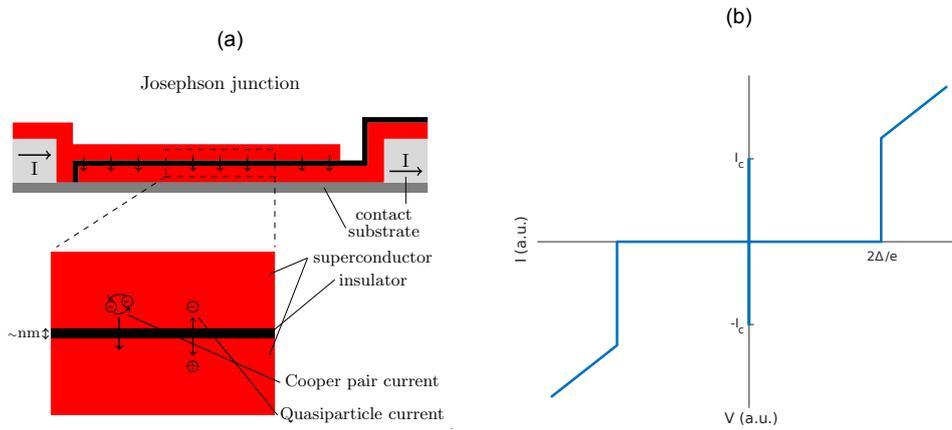


Figure 2.8: a) Schematic of two different types of Josephson junction. Red represents a superconducting material (e.g. Al), black an insulating layer (e.g. Al_2O_3), dark grey a substrate (e.g. Si) and light grey are the contacts (generally also superconducting). Zoom-in schematically shows two important factors contributing to the current: The supercurrent from Cooper pairs tunnelling through the barrier, and the quasi-particle tunnelling current. b) I-V curve of an ideal Josephson junction. The current peak at $V = 0$ consists of Cooper pairs tunnelling, while the slope for large positive and negative voltage originates from quasi-particles tunnelling across the barrier.

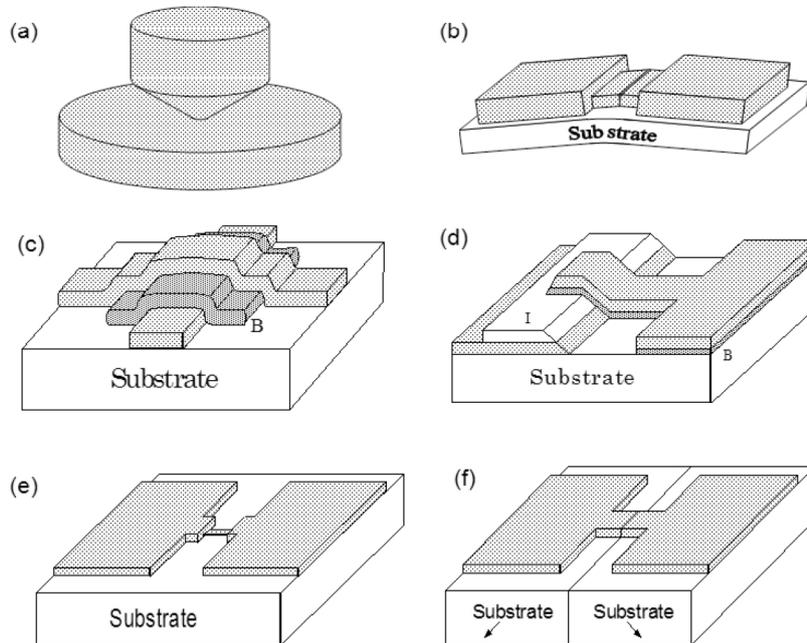


Figure 2.9: An overview of several types of Josephson junctions: a) Point contact junction, b) break junction, c) crossed electrodes and barrier (B) junction, d) edge junction with insulator (I) and barrier (B) layers, e) nanobridge junction, f) bicrystal grain boundary junction [30].

form

$$\psi = \sqrt{\rho}e^{i\phi} \quad (2.22)$$

where we treat ρ as the (constant) density of the charge-carrying particles in the superconductor, Cooper pairs, and we call ϕ the phase of the wave function. The Schrödinger equation,

$$i\hbar\frac{\partial\psi}{\partial t} = \hat{H}\psi \quad (2.23)$$

describes the time evolution of the wave function, where \hat{H} is the Hamiltonian of the system. In steady state, the Hamiltonian working on the wave function results in a constant factor that represents the energy of the state, the energy eigenvalue,

$$i\hbar\frac{\partial\psi}{\partial t} = \mu\psi. \quad (2.24)$$

Suppose we take two separate pieces of superconducting material, ψ_1 and ψ_2 with Cooper pair densities ρ_1 and ρ_2 , and phases ϕ_1 and ϕ_2 respectively. If we bring the two pieces sufficiently close together, a coupling will start to play a role and can be taken into account using a simple model if it is small enough. We get two coupled equations,

$$\begin{aligned} i\hbar\frac{\partial\psi_1}{\partial t} &= \mu_L\psi_1 + K\psi_2 \\ i\hbar\frac{\partial\psi_2}{\partial t} &= \mu_R\psi_2 + K\psi_1. \end{aligned} \quad (2.25)$$

The coupling is described by a coupling constant K . If the superconductors are identical, the energies will be equal, $\mu_1 = \mu_2$. If we then apply a potential difference V between the two superconductors and offset the zero of energy, we get $\mu_1 = eV$, $\mu_2 = -eV$. Here $2e$ is the charge of the Cooper pairs, twice the electron charge. We obtain

$$\begin{aligned} i\hbar\frac{\partial\psi_1}{\partial t} &= eV\psi_1 + K\psi_2 \\ i\hbar\frac{\partial\psi_2}{\partial t} &= -eV\psi_2 + K\psi_1. \end{aligned} \quad (2.26)$$

Now, we can substitute our assumption of the wave function form, $\psi_1 = \sqrt{\rho_1}e^{i\phi_1}$ and $\psi_2 = \sqrt{\rho_2}e^{i\phi_2}$. Then, we separate the two equations each in a real and imaginary part. To simplify, we write the phase difference between the superconductors as $\Delta\phi = \phi_2 - \phi_1$,

$$\begin{aligned} \frac{\partial\rho_1}{\partial t} &= \frac{1}{\hbar}K\sqrt{\rho_1\rho_2}\sin(\Delta\phi) \\ \frac{\partial\rho_2}{\partial t} &= -\frac{1}{\hbar}K\sqrt{\rho_1\rho_2}\sin(\Delta\phi) \\ \frac{\partial\phi_1}{\partial t} &= -\frac{K}{\hbar}\sqrt{\frac{\rho_2}{\rho_1}}\cos(\Delta\phi) + \frac{eV}{\hbar} \\ \frac{\partial\phi_2}{\partial t} &= -\frac{K}{\hbar}\sqrt{\frac{\rho_1}{\rho_2}}\cos(\Delta\phi) - \frac{eV}{\hbar} \end{aligned} \quad (2.27)$$

The first two of these equations show that the change in pair density between the two superconductors is opposite in sign, and equal in size, as expected from conservation of charge. To avoid charging of the individual superconductors, there must be a current flowing through the circuit connecting the two superconductors to the outside world (i.e. the source of the voltage difference). This shows that these two equations already have the correct form,

$$I = I_c \sin(\Delta\phi). \quad (2.28)$$

Where the I_c is a constant. An equation for this constant can be obtained from more detailed microscopic description of the situation, but what follows is called the Ambegaokar-Baratoff relation [31],

$$I_c = \frac{\pi\Delta(T)}{2eR_n} \tanh\left(\frac{\Delta(T)}{2k_B T}\right), \quad (2.29)$$

where $\Delta(T)$ is the energy gap (temperature dependent), R_n is the normal state resistance of the junction (due to quasiparticle tunneling) and k_B is the Boltzmann constant.

The other two equations of Eq. (2.27), can be subtracted from each other to get on the left hand side the derivative of the phase difference, $\frac{\partial\phi}{\partial t}$, and on the other side the voltage V with some constants. Shifting these to the left hand side returns the other Josephson equation in Eq. (2.21).

2.2.2 Josephson inductance

In this thesis, we utilize Josephson junctions as inductors in a LC oscillator. We treat the junction as a circuit element, as shown schematically in Fig. 2.10a. To do this accurately, we need to find the inductance of a single Josephson junction, L_J . We follow the derivation of [27].

The general form of a voltage across an inductance is

$$V = L \frac{\partial I}{\partial t}. \quad (2.30)$$

Using the chain rule, we can split the derivative with respect to time: $\frac{\partial I}{\partial t} = \frac{\partial I}{\partial \Delta\phi} \frac{\partial \Delta\phi}{\partial t}$ and we can use both Josephson relations, Eq. (2.21), to obtain expressions for both parts,

$$\begin{aligned} \frac{\partial I}{\partial \Delta\phi} &= I_c \cos(\Delta\phi) \\ \frac{\partial \Delta\phi}{\partial t} &= \frac{2eV}{\hbar} = \frac{2\pi V}{\Phi_0} \end{aligned} \quad (2.31)$$

Where Φ_0 is the magnetic flux quantum (Sec. 2.2.3). We can combine these to obtain an equation of the form of the general voltage-over-inductor relation,

$$V = \frac{\Phi_0}{2\pi I_c} \frac{1}{\cos(\Delta\phi)} \frac{\partial I}{\partial t}. \quad (2.32)$$

Where we retrieve the Josephson inductance

$$L_J(\phi) = \frac{\Phi_0}{2\pi I_c} \frac{1}{\cos(\Delta\phi)} = \frac{\Phi_0}{2\pi I_c} \frac{1}{\sqrt{1 - \sin^2(\Delta\phi)}} = \frac{\Phi_0}{2\pi I_c} \frac{1}{\sqrt{1 - \frac{I^2}{I_c^2}}}. \quad (2.33)$$

We see that the Josephson inductance is non-linear, it depends on the phase difference between the superconducting parts of the junction.

2.2.3 Flux quantization

We take a side-step from the Josephson junctions, to explain a phenomenon called flux quantization. In a superconductor, the magnetic flux is quantized, with one magnetic flux quantum being $\Phi_0 = \frac{h}{2e} \approx 2.07 \cdot 10^{-15}$ Wb. We can see effects of this if the superconductor surrounds a non-superconducting part, for example if we have a hole in a ring-shaped superconductor, or if we consider the center of a magnetic vortex. Flux quantization was first predicted phenomenologically by London in 1948 [32], and experimental evidence for flux quantization was given in 1961 for cylinders [33] and rings [34].

Flux quantization follows from the form of the superconducting wave function, Eq. (2.22) $\psi = \sqrt{\rho} e^{i\phi}$, where again ϕ is the phase of the wave function. Adding 2π to the phase ϕ will result in an identical wave function, due to the exponent. So when a superconductor has a hole (e.g. it is loop or ring shaped, or has a non-superconducting part in its interior), the wave function can contain a phase change of $n2\pi$ (integer n) over the loop without it changing the state the superconductor is in.

There are two main factors which affect the phase of the wave function: The movement of the charge carriers (Cooper pairs) and the magnetic field, expressed by the supercurrent density \vec{j} and the vector

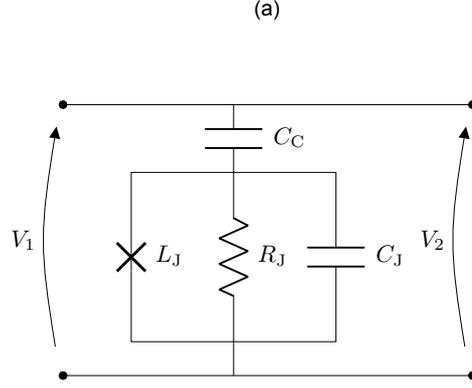


Figure 2.10: a) Model of a Josephson junction capacitively coupled to a transmission line. The junction, denoted by a cross, can be treated as a resistively and capacitively shunted inductor with inductance L_J . It is capacitively coupled to a transmission line by a capacitor C_C .

potential \vec{A} respectively (see also Appendix I of [9]). To see how the phase of the wave function changes around a hole on a loop or ring shaped superconductor, we study the gradient of the phase and integrate it over a path around that loop.

First, the gradient of the superconducting phase ($\nabla\phi$) is

$$\nabla\phi = \frac{1}{\hbar}(m\vec{v} + q\vec{A}) = \frac{1}{\hbar} \left(\frac{m}{qn_p} \vec{j} + q\vec{A} \right), \quad (2.34)$$

where m is the mass of a charge carrier, \vec{v} its velocity, q the charge, n_p the number of charge carriers, \vec{j} the current density, and \vec{A} is the local (magnetic) vector potential. We must have that the superconductor be described by one wave function, but as we have seen earlier, we can add $n2\pi$ (n integer) to the phase without it changing the state. Thus, total change in phase around a hole in the superconductor must be an integer multiple of 2π . In other words, the flux quantization condition states

$$\oint \nabla\phi \cdot d\vec{l} = 2\pi n. \quad (2.35)$$

for integer n . Using the expression for the phase gradient, this becomes

$$\frac{m}{qn_p\hbar} \oint \vec{j} \cdot d\vec{l} + \frac{q}{\hbar} \oint \vec{A} \cdot d\vec{l} = 2\pi n, \quad (2.36)$$

which can be rewritten in terms of magnetic flux. We can simplify this by using the London penetration depth $\lambda_L = \left(\frac{\mu_0 q n_p}{m}\right)^{-1/2}$, which is a constant [35], to get

$$\mu_0 \lambda_L^2 \oint \vec{j} \cdot d\vec{l} + \Phi = n\Phi_0. \quad (2.37)$$

Here $\Phi = \oint \vec{A} \cdot d\vec{l}$ is the flux enclosed by the superconducting loop, Φ_0 is the flux quantum and n is again an integer.

Now we explain how to see Eq. (2.37). Suppose we have a ring-shaped superconductor with zero external magnetic field, Fig. 2.11a. Now if we impose a small magnetic field that generates a certain flux through the ring (Fig. 2.11b), an extra 'compensation' current, known as the screening current, will form in the ring, which follows from Faraday's law. This extra current will generate a magnetic field to counter the flux from our external field, so that the total flux through the ring is still zero.

We can invert the situation, by having some magnetic flux already penetrating the hole in the superconductor (if we cool it down in the presence of an external field). Then we have flux trapping, which means that one or more flux quanta are present in the hole in the superconductor. The currents in the superconductor will then ensure that the flux in the hole is always equal to the trapped flux, even if the external magnetic field is zero, as in Fig. 2.11c.

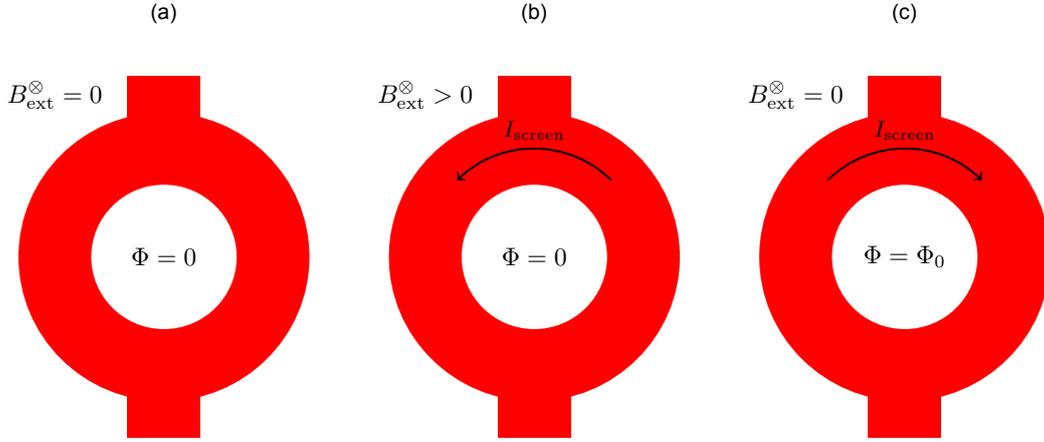


Figure 2.11: a) A superconducting ring (red) cooled down below its critical temperature with zero external magnetic field B_{ext}^{\otimes} will have no flux Φ threading the ring. b) When a weak external field is added, the flux threading the ring will still be zero. To achieve this, a screening current will run through the ring. c) When we apply an external field as we cool down the structure so that it becomes superconducting, we can trap flux in the loop. Then, the screening current will generate flux through the loop even if the external field is zero.

2.2.4 SQUIDS

A Superconducting QUANTUM Interference Device (SQUID) is a device that consists of a superconducting loop interrupted at one (RF-SQUID, where RF is Radio Frequency, and used instead of AC) or two (DC-SQUID) points by a Josephson Junction [9]. The first fabrication of SQUIDS dates back to the 60's [36, 37], and the technology has been developed further alongside the semiconductor industry and with the advent of high- T_C superconductors [9].

SQUIDS tie the concepts of Josephson junctions and flux quantization together. The two junctions form a loop (usually a rectangle) of superconducting material, where the flux penetrating the loop is quantized. An example of a SQUID is shown in Fig. 2.13. The current through the loop cannot take any arbitrary value (as it could in the ideal superconductor in Sec. 2.2.3), because the Josephson junctions have a critical (maximum) current. Suppose we have a current from one side of the loop to the other (via the connections at the top and bottom of Fig. 2.11a, for example), and we apply an external magnetic field. The maximum current through the SQUID will change with the applied field, as plotted in Fig. 2.12a. Effectively, we can change the critical current of the junction by applying a magnetic field.

The modulation of the critical current can also be seen as a modulation of the inductance that the SQUID has. For a DC SQUID with two identical Josephson junctions with each a Josephson inductance $\frac{L_J}{2}$, an external flux Φ_{ext} , the SQUID inductance is given by [38]

$$L_{\text{SQUID}} = \frac{L_J}{\left| \cos\left(\pi \frac{\Phi_{\text{ext}}}{\Phi_0}\right) \right|}. \quad (2.38)$$

Based purely on this equation, we would have a SQUID inductance that diverges for specific values of our applied magnetic field, $\Phi_{\text{ext}} = \frac{\pi}{2}, \frac{3\pi}{2}, \frac{5\pi}{2}, \dots$, as is plotted in Fig. 2.12b. We have only used a simple model to describe ideal SQUIDS, without taking into account things such as the loop inductance,

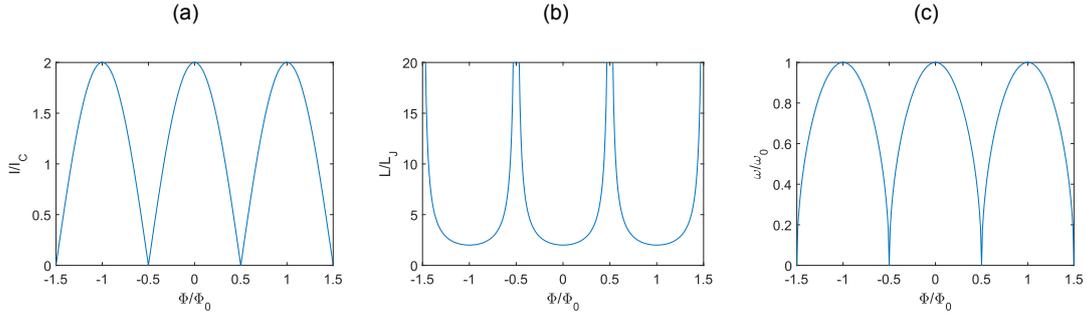


Figure 2.12: a) For a SQUID, the critical current is modulated by the flux penetrating the SQUID loop. The effective critical current is a periodic function of the flux. b) Based on the modulation of the critical current, the inductance of a SQUID is also modulated by the flux penetrating the loop area. c) The resonance frequency of a SQUID embedded in an LC oscillator is also modulated by the flux penetrating the SQUID loop.

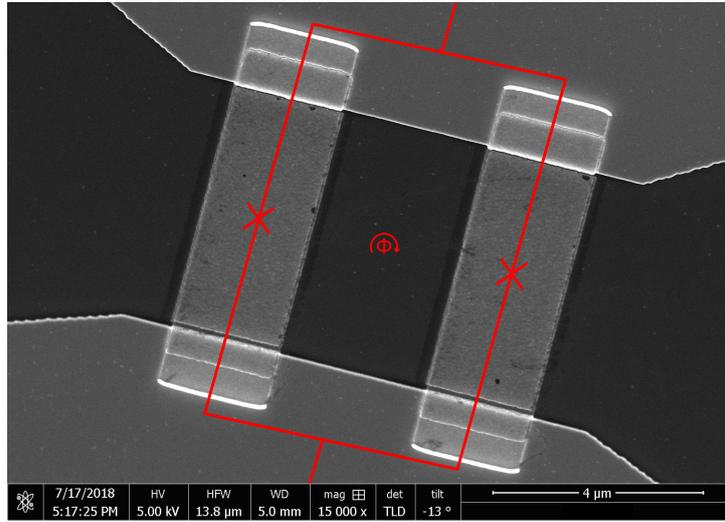


Figure 2.13: SEM image of a SQUID. In black the Silicon substrate, in dark grey two metal electrodes and in lighter grey the two Aluminium Josephson junctions. The flux through the red loop between the two junctions is quantized, and an external field can be used to modulate the SQUID inductance. Overlaid on the image in red is part of the circuit diagram showing two Josephson junctions denoted by crosses, and the flux through the area they enclosed denoted by Φ .

or mutual inductance from the parallel junctions. Qualitatively, this model serves to illustrate the most important function of the SQUID for this thesis, namely that of a tunable inductor.

When combined with a capacitance C , the SQUID forms a resonator with resonant frequency

$$\omega_0 = \frac{1}{\sqrt{C \frac{L_J}{|\cos(\pi \frac{\Phi_{\text{ext}}}{\Phi_0})|}}}. \quad (2.39)$$

The resonance frequency of the SQUID oscillator (also SQUID cavity) is modulated by the flux penetrating the SQUID loop, as plotted in Fig. 2.12c. The fact that we can tune the resonance frequency of an oscillator containing a SQUID by varying the applied magnetic field makes it a valuable structure for many applications in research.

2.3 Parametric amplification

Parametric amplification is an effect where the amplitude of the oscillation of a harmonic oscillator is increased (i.e. it is driven) by a periodically varying system parameter. If the driving frequency is approx-

imately twice the natural frequency of the oscillator, it amplifies an input signal (i.e. the initial oscillation), depending on the phase relation between the two oscillations. An example of this is moving your center of gravity by standing and squatting on a swing to go faster¹. In our case, this input signal is electrical, in the form of an oscillating voltage and current. The phenomenon was first described by Faraday [1], and its description was generalized by Lord Rayleigh [39]. Parametric amplifiers are commonly used in electronics [4] and optics [40].

2.3.1 Parametric oscillation

Parametric oscillators are common across many fields of physics, so parametric oscillation is generally treated in textbooks on dynamics or optics [40]. The main goal of this section is to show how varying the parameter of an oscillator at a specific frequency results in a growing oscillation amplitude, i.e. driving the oscillator. To do that, we will treat the LC oscillator as a parametric oscillator. We can apply the derivation we perform below on an RLC oscillator to include damping, but this makes it rather tedious.

We start from an LC oscillator, and we make the resonance frequency time dependent. We can do this by varying the capacitance, or the inductance (or both). For the derivation, it does not matter which one we vary, so to keep it general we will not specifically choose one. We will just make the resonance frequency itself time dependent,

$$\frac{1}{LC} = \omega_0^2 = \omega_0^2(t). \quad (2.40)$$

Parametric oscillation does not happen for all time-dependent resonance frequencies. To get parametric oscillation, we must have a sinusoidal variation, with a frequency that is approximately twice the natural resonance frequency of our oscillator. So, we describe the time dependence to be small (amplitude α), on top of an average ω_0 , and sinusoidal with frequency 2ω

$$\omega_0^2(t) = \omega_0^2 (1 + \alpha \sin(2\omega t)). \quad (2.41)$$

We can put this into the differential equation describing the current through our LC oscillator. To show parametric oscillation, we must solve

$$\frac{\partial^2 I}{\partial t^2} + \omega_0^2 (1 + \alpha \sin(2\omega t)) I = 0. \quad (2.42)$$

Because the amplitude of the variation of our resonance frequency is small, it will not alter the general form of our solution. We expect the current to be sinusoidal, as in Eq. (2.1), and grow (or decay) over time, slowly with respect to the resonance frequency. We modify the expected solution from Eq. (2.1) by separating the fast (cos and sin) and slowly changing (A and B) parts, and try a solution of the form

$$I(t) = A(t) \cos(\omega t) + B(t) \sin(\omega t). \quad (2.43)$$

Here, we have tacitly assumed the solution has frequency ω , instead of the resonance frequency of the bare oscillator, ω_0 . We can do this, as long as ω and ω_0 are not too different. We will later quantify the allowed difference (i.e. the detuning) between the frequencies.

Now, we substitute the proposed solution into the differential equation and work out the double time derivative. This gives us

$$\begin{aligned} \ddot{A}(t) \cos(\omega t) - 2\omega \dot{A}(t) \sin(\omega t) - \omega^2 A(t) \cos(\omega t) + \ddot{B}(t) \sin(\omega t) + 2\omega \dot{B}(t) \cos(\omega t) - \omega^2 B(t) \sin(\omega t) \\ + \omega_0^2 (1 + \alpha \sin(2\omega t)) (A(t) \cos(\omega t) + B(t) \sin(\omega t)) = 0. \end{aligned} \quad (2.44)$$

We can drop the terms containing $\ddot{A}(t)$ and $\ddot{B}(t)$ because we presumed them to change slowly, so their second time derivative will be small. We also see that by writing out the last term, we get terms that

¹The standing/squatting is important: the parameter we vary is the length of the oscillator. By standing/squatting, we change the distance between our center of mass and the rotation point of the swing, which is determines the length of the oscillator. This is the parameter that we vary to make the oscillation parametric.

contain a product of sines and cosines, which we can rewrite using the product-to-sum relations,

$$\begin{aligned}\sin(2\omega t) \cos(\omega t) &= \frac{1}{2} (\sin(\omega t) + \sin(3\omega t)) \\ \sin(2\omega t) \sin(\omega t) &= \frac{1}{2} (\cos(\omega t) - \cos(3\omega t)).\end{aligned}\tag{2.45}$$

This step contains the crux of the derivation. The sinusoidal time-variation of the parameter of our oscillator at frequency 2ω results in a term at frequency ω , and one at 3ω . Using the assumption that ω is not too different from the resonance frequency of the oscillator, ω_0 , this term effectively serves as a driving force resonant with the oscillator. Intuitively, we would expect the system to oscillate with increasing amplitude if we drive it on resonance. This is the basis of parametric amplification.

To show parametric oscillation more formally than just intuition, we continue with the derivation. We neglect the terms with $3\omega t$, presuming them to be far off-resonance and thus strongly attenuated. Then, we group the terms from Eq. (2.44) by whether they contain a $\sin(\omega t)$ or a $\cos(\omega t)$. We get

$$\begin{aligned}\left(-2\omega\dot{A}(t) - \omega^2 B(t) + \omega_0^2 B(t) + \omega_0^2 \frac{\alpha}{2} A(t)\right) \sin(\omega t) \\ + \left(-\omega^2 A(t) + 2\omega\dot{B}(t) + \omega_0^2 A(t) + \omega_0^2 \frac{\alpha}{2} B(t)\right) \cos(\omega t) = 0.\end{aligned}\tag{2.46}$$

To get this expression to be zero at any time, we must have both the parts in the big brackets be zero, as $\sin(\omega t)$ and $\cos(\omega t)$ will never be simultaneously zero for any t . That is, we get two equations, which are coupled differential equations for amplitudes $A(t)$ and $B(t)$,

$$\begin{aligned}-2\omega\dot{A}(t) - \omega^2 B(t) + \omega_0^2 B(t) + \omega_0^2 \frac{\alpha}{2} A(t) &= 0 \\ 2\omega\dot{B}(t) - \omega^2 A(t) + \omega_0^2 A(t) + \omega_0^2 \frac{\alpha}{2} B(t) &= 0.\end{aligned}\tag{2.47}$$

To solve these equations and obtain expressions for $A(t)$ and $B(t)$, we need to decouple them. We can do this by the following change of variables,

$$\begin{aligned}A(t) &= r(t) \cos(\theta(t)) \\ B(t) &= r(t) \sin(\theta(t)),\end{aligned}\tag{2.48}$$

and we get

$$\begin{aligned}-2\omega \left(\dot{r} \cos(\theta) - r\dot{\theta} \sin(\theta)\right) - \omega^2 r \sin(\theta) + \omega_0^2 r \sin(\theta) + \frac{\alpha}{2} \omega_0^2 r \cos(\theta) &= 0 \\ 2\omega \left(\dot{r} \sin(\theta) + r\dot{\theta} \cos(\theta)\right) - \omega^2 r \cos(\theta) + \omega_0^2 r \cos(\theta) + \frac{\alpha}{2} \omega_0^2 r \sin(\theta) &= 0.\end{aligned}\tag{2.49}$$

Now we have translated the problem into one of finding $r(t)$ and $\theta(t)$. To obtain an equation for $r(t)$, we multiply the upper equation by $-\cos(\theta)$ and the lower equation by $\sin(\theta)$ and add them. We have

$$\dot{r} = \frac{\alpha\omega_0^2}{4\omega} \cos(2\theta)r.\tag{2.50}$$

To obtain an equation for $\theta(t)$, we multiply the upper equation by $\sin(\theta)$ and the lower equation by $\cos(\theta)$ and then add the two. We have

$$\dot{\theta} = -\frac{\alpha\omega_0^2}{4\omega} \left(\sin(2\theta) - \frac{2}{\alpha} \frac{\omega^2 - \omega_0^2}{\omega_0^2}\right).\tag{2.51}$$

The equation for $\theta(t)$ does not depend on $r(t)$, so we can solve it. Linearising θ , we can see that it will move towards an equilibrium value

$$\sin(2\theta_{\text{meq}}) = \frac{2}{\alpha} \frac{\omega^2 - \omega_0^2}{\omega_0^2}.\tag{2.52}$$

With this, we obtain as the equation for $r(t)$

$$\dot{r} = \frac{\alpha\omega_0^2}{4\omega} \cos\left(2\frac{\omega^2 - \omega_0^2}{2\omega}\right) r. \quad (2.53)$$

Which can straightforwardly be solved,

$$r = r_0 e^{at}. \quad (2.54)$$

Where the exponential growth rate a is

$$a = \frac{\alpha\omega_0^2}{4\omega} \cos\left(\frac{2}{\alpha} \frac{\omega^2 - \omega_0^2}{\omega_0^2}\right). \quad (2.55)$$

To recap, we have that $\theta(t)$ becomes a constant, while $r(t)$ grows exponentially. This means that $A(t)$ and $B(t)$ also grow exponentially, with the same rate as $r(t)$, and that the amplitude of our solution $I(t)$ grows exponentially as well. Or, to put it formally,

$$\begin{aligned} I(t) &= A(t) \cos(\omega t) + B(t) \sin(\omega t) \\ &= r(t) \cos(\theta(t)) \cos(\omega t) + r(t) \sin(\theta(t)) \sin(\omega t) \\ &= r_0 e^{at} (\cos(\theta_{\text{eq}}) \cos(\omega t) + \sin(\theta_{\text{eq}}) \sin(\omega t)). \end{aligned} \quad (2.56)$$

So, the current through our oscillator oscillates with a frequency ω , and with an amplitude $r_0 e^{at}$ that grows (or decays) exponentially with time.

Obviously, this becomes un-physical quickly: We cannot have an infinitely large current through our system. A partial solution to this problem can be achieved by including a resistance, which is straightforward but rather tedious, and results in the equation for $r(t)$,

$$r = r_0 e^{(a-\kappa)t}, \quad (2.57)$$

where the κ is the loss rate. The difference between the growth rate a and the loss rate κ determines whether the oscillations keep growing, keep decaying or stay the same. This effect is straightforward: The resistance in a driven RLC circuit damps the oscillations which would continue growing in a driven LC circuit.

As a last comment on this derivation: We have tacitly assumed that ω_0 and ω are not too different. If they are identical, θ_{eq} will be zero, but if they are not (for a certain detuning), it will become non-zero. This will cause the exponential growth constant a to decrease. The condition for parametric oscillation on ω is

$$\omega_0 \sqrt{1 - \frac{\alpha}{2}} < \omega < \omega_0 \sqrt{1 + \frac{\alpha}{2}} \quad (2.58)$$

Outside this range of detuning, we cannot get a growing oscillation intensity. This is one of the limits of a parametric amplifier: We have to drive the oscillator at a frequency 2ω that is not too different from twice the natural resonance frequency ω_0

In this part, we have shown that for an oscillator, varying its parameters at twice the natural frequency results in oscillations at its natural frequency that grow exponentially. In other words, we have shown parametric oscillation. To go from parametric oscillation to (a description of) a parametric amplifier, we need to consider other things, such as for example an input signal (oscillating voltage or current) that is to be amplified. We will do that in the next section.

2.3.2 From a parametric oscillator to a parametric amplifier

To obtain a description of a parametric amplifier, we must add an input signal (oscillating voltage or current) to our differential equation. Initially, we will treat it as an harmonic oscillator where we describe the amplitude of the internal field, and for later steps we can place the harmonic oscillator into our circuit as done in, for example, [41].

We start with a differential equation as in Eq. (2.42), but now we include a signal f_e and some damping in the form of a loss rate $\kappa = \frac{\omega_0}{Q}$. That is,

$$\frac{\partial^2 X}{\partial t^2} + \kappa \frac{\partial X}{\partial t} + \omega_0^2(1 + \alpha \sin(2\omega t))X = f_e. \quad (2.59)$$

We make a switch in variables, from our dummy variable X to the complex variable A and A^* ,

$$\begin{aligned} A &= \frac{\partial X}{\partial t} - i\omega_1^* X \\ A^* &= \frac{\partial X}{\partial t} + i\omega_1^* X \\ \omega_1 &= \omega_0 \left(\left(1 - \frac{1}{4Q^2}\right)^{\frac{1}{2}} - \frac{i}{2Q} \right). \end{aligned} \quad (2.60)$$

These variables, A and A^* represent the complex field amplitudes. We are interested in how they change, and rewriting Eq. (2.59) in terms of the field amplitude gives us

$$\frac{\partial A}{\partial t} = -i\omega_1 A - i\alpha_p^2 \sin(2\omega t) \frac{A - A^*}{\omega_1^* + \omega_1} + f_e, \quad (2.61)$$

where we have taken $\alpha_p = \frac{\alpha}{\omega_0}$. For the systems in this thesis, we can make the high-Q approximation. If $Q \gg 1$, we take $\omega_1 = \omega_0 - \frac{i\kappa}{2}$. Then we can simplify our equation for the complex field amplitude,

$$\frac{\partial A}{\partial t} = -i \left(\omega_0 - \frac{i\kappa}{2} \right) A - i \frac{\alpha_p^2}{2\omega_0} (A - A^*) \sin(2\omega t) + f_e. \quad (2.62)$$

Now we assume that the solution for the intra-cavity field takes the form $A = A_0 e^{-i\omega_s t}$, that our signal is given by $f_e = \sqrt{\frac{\kappa_e}{2}} S_{in} e^{i\omega_s t}$, and we will call $\omega_s - \omega_0$ the detuning between the signal and resonance frequencies.

So far, the description has been of a general parametric amplifier. Here, we will only use this bit of theory for the analysis of our specific amplifier. That is, we are only interested in the transmission of a side-coupled cavity, S_{21} , and we know we measure in a frame rotating with the signal frequency ω_s . Furthermore, for the specific system studied in this thesis, half the frequency of our parametric modulation, ω , is slightly different from ω_s .

Taking all the steps described above, we obtain

$$0 = \left(i(\omega_s - \omega_0) - \frac{\kappa}{2} \right) A_0 - \frac{\alpha_p^2}{4\omega_0} A_0^* e^{i(\omega - \omega_s)t} + \sqrt{\frac{\kappa_e}{2}} S_{in}. \quad (2.63)$$

We can rewrite this to obtain an expression for the cavity field amplitude A_0 ,

$$A_0 = \sqrt{\frac{\kappa_e}{2}} \frac{1}{\frac{\kappa^2}{4} - \left(\frac{\alpha_p^2}{4\omega_0} \right)^2 + (\omega_s - \omega_0)^2} \left(\frac{\alpha_p^2}{4\omega_0} e^{i(\omega - \omega_s)t} + i(\omega_s - \omega_0) + \frac{\kappa}{2} \right) S_{in}. \quad (2.64)$$

We are interested in what we get at the output port of our system. For a side-coupled cavity in transmission, we have

$$S_{out} = S_{in} - \sqrt{\frac{\kappa_e}{2}} A_0. \quad (2.65)$$

Combining these last two equations, we can get an expression for the parameter we can measure,

$$S_{21} = \frac{S_{out}}{S_{in}} = 1 - \frac{\kappa_e}{\frac{\kappa^2}{4} - \left(\frac{\alpha_p^2}{4\omega_0} \right)^2 + (\omega_s - \omega_0)^2} \left(\frac{\alpha_p^2}{4\omega_0} e^{i(\omega - \omega_s)t} + i(\omega_s - \omega_0) + \frac{\kappa}{2} \right). \quad (2.66)$$

So, we have obtained an expression for the transmission parameter, S_{21} of a parametric amplifier that is side-coupled to a transmission line. Qualitatively, we see that we should get a signal oscillating with $\omega - \omega_s$, the detuning between our signal and half the frequency of our parametric pump.

2.3.3 Modes of parametric amplification

Turning from the description of our specific parametric amplifier back to a more general one, we will now describe several types of parametric amplifiers. We will first treat degenerate versus non-degenerate amplification, and later treat the difference between 3-wave and 4-wave mixing, which is the 3-photon versus the 4-photon process.

Degenerate versus non-degenerate

Degenerate versus non-degenerate parametric amplification (Fig. 2.14) refers to the frequency relation between half the pump and the signal. Suppose we pump at a frequency twice the resonance frequency of our cavity, $\omega_d = 2\omega_0$. If the input signal is at exactly half of that frequency, $\omega_s = \omega_0$, we have what is called a degenerate parametric amplifier. In this case, the amplifier is phase-sensitive, the phase relation between the pump and the signal determines if we get amplification or de-amplification [42, 43]. In Sec. 3.5, we simulate a general parametric amplifier and show the phase-sensitivity of the gain for a degenerate parametric amplifier.

Non-degenerate parametric amplification happens when the signal is not exactly half of the pump frequency, $\omega_s \neq \omega = \frac{\omega_d}{2}$, but has a different frequency. We will refer to the difference between the actual frequency of the input signal and half of the pump signal, $\omega_s - \omega$ as the detuning. The non-degenerate case differs in three ways from the degenerate case.

The first difference is that for the non-degenerate case, the amplification (gain) is effectively not phase-sensitive any longer. That is, the gain is still affected by the phase difference between the signals, but the phase difference between two signals at different frequencies is not a constant. The phase difference changes with 2π times the frequency difference (detuning) between half the pump and the input signal every second. If this detuning is large enough, the phase difference changes so fast that the amplifier does not reach the steady-state gain it can reach in the degenerate case. Effectively, the gain in the non-degenerate case is the mean value of the gain in the degenerate case integrated over the phase. As this is a constant, the gain for a non-degenerate parametric amplifier is effectively independent of the phase-difference between the input signal and the pump.

The second difference between the non-degenerate and the degenerate cases is the appearance of an idler at frequency ω_i , detuned on the other side of the resonance so that $\omega_s + \omega_i = 2\omega = \omega_d$. This is shown in the bottom case of Fig. 2.14. If the amplifier amplifies the input signal with a gain \sqrt{G} , the idler will be amplified with respect to the input signal with a factor $\sqrt{G - 1}$ [16].

The last difference between the degenerate and non-degenerate amplifiers lies in the limit to the noise they add to a signal. In the ideal case, a phase-sensitive amplifier does not add any noise to the signal it amplifies, while a phase-insensitive amplifier does add some noise [10]. This lack of noise is one of the main reasons to study parametric amplifiers.

3-wave versus 4-wave mixing

3- and 4-wave mixing refers to the process that a parametric amplifier uses. This is independent of the degenerate versus non-degenerate case. For a 3-wave amplifier, one pump photon ω_d can split into a signal- and an idler wave (signal and idler photon, or two signal photons, in the degenerate case), via $\omega_d = \omega_s + \omega_i$. That is, the pump frequency is twice the resonance frequency of the cavity $\omega_d = 2\omega_0$, as is sketched in Fig. 2.14. For a 4-wave amplifier, two pump photons ω_d can be split into a signal- and an idler wave, via $2\omega_d = \omega_s + \omega_i$. That is, the pump frequency is equal to the cavity frequency, $\omega_d = \omega_0$.

The difference between the 3-wave and 4-wave processes appears small, but it is important for practical reasons. For example, in a 3-wave degenerate parametric amplifier, with the signal at $2\pi \cdot 6$ GHz, the pump must be at $2\pi \cdot 12$ GHz, and we can easily filter the pump from the output signal. For a 4-wave degenerate parametric amplifier, this would be an issue. On the other hand, generating a pump at $2\pi \cdot 12$ GHz could require different electronics than generating a pump at $2\pi \cdot 6$ GHz.

Another difference between 3-wave and 4-wave mixing is that the 4-wave mixing requires a certain

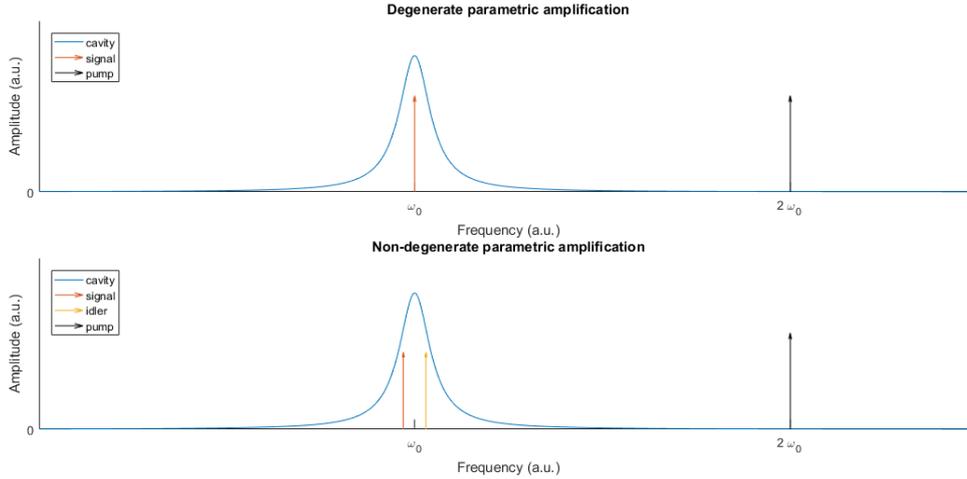


Figure 2.14: Frequency diagram of parametric amplification, degenerate (top) versus non-degenerate (bottom) mode. In both cases, we pump parametrically at frequency that is twice the resonance frequency of the cavity, $\omega_d = 2\omega_0$. For the degenerate case, our signal is at frequency $\omega_s = \omega_0$, while for the degenerate case, it is slightly detuned from that. For the non-degenerate case, we get an additional peak detuned on the other side of the resonance, which is called the idler, such that $\omega_s + \omega_i = 2\omega = \omega_d$.

non-linearity, as it involves a two-photon process. Parametric amplification can only happen if the pump is at twice the resonance frequency of the oscillator. The way 4-wave mixing works is that if the parameter we vary has a large enough non-linearity, the drive at ω_0 can become an effective drive at $2\omega_0$. The difficulty of realizing this non-linearity is dependent on the system specifics. In contrast to the 4-wave process, the 3-wave process does not require any non-linearity at all.

Flux-pumped versus current-pumped

For our specific type of parametric amplifiers, the Josephson parametric amplifier, there is another choice in how to pump the system [44, 16]. It is possible to pump the system by sending in a current (signal) at $2\omega_0$, but it is also possible to vary the flux penetrating the SQUID loop at $2\omega_0$. The current-pump has the advantage of requiring only one transmission line where the flux-pump requires two, but this brings with it the disadvantage of having to separate out the (desired) signal from the pump.

A flux-pumped amplifier requires some extra consideration in design. To amplify as much as possible (achieve a high gain), we want to vary the flux penetrating the SQUID loop by a large amount. This requires positioning the flux-pump transmission line such that it generates the largest magnetic field (and thus flux) at the SQUID. In a current-pumped amplifier, this is not an issue, since the pump is coupled to the cavity in the same way as the signal is.

An overview of these different types of JPAs is given in Fig. 2.15. Here, combinations of degenerate versus non-degenerate, 3-wave versus 4-wave and flux-pumped versus current-pumped are shown. For the relevant systems, the flux penetrating a SQUID loop is denoted by the Φ with an arrow circled around it. The crosses denote Josephson junctions, and the pipes represent a coupling to a transmission line.

2.3.4 JPA parameters

For building a practical JPA, it is useful to look at design parameters from other JPAs used in previous works, an incomplete overview of these is given in Table 2.1. Important typical parameters of an amplifier are the (operating) gain, the -1dB bandwidth, the input power limit and noise temperature. In most of the mentioned papers, the gain stated is not the maximum gain the amplifier can reach. Conventionally, a single operating point is chosen such that the amplifier has a stable gain. That is, around its maximum, the gain might be very sensitive to disturbances, so a better signal-to-noise ratio is reached if a lower

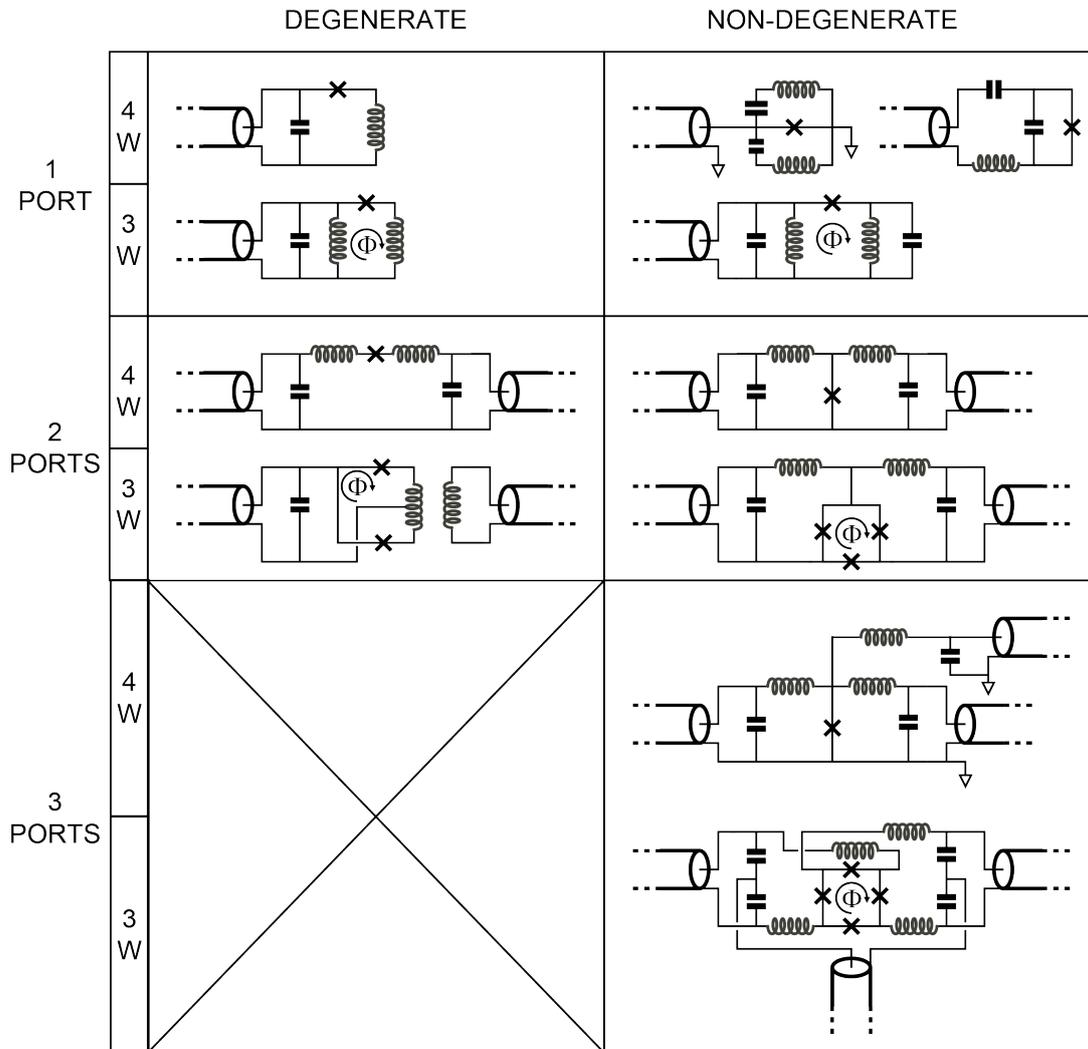


Figure 2.15: Overview of JPAs with different number of ports, different modes of pumping (3-wave and 4-wave) and degenerate versus non-degenerate setups. Flux penetrating a SQUID loop is denoted by the Φ with the arrow circled around it, the crosses represent Josephson junctions [45].

Table 2.1: Overview of JPA parameters in some recent papers

	Gain (dB)	Bandwidth (MHz)	Input power limit (dBm)	Process	Pump	Connection	Mode
[47]	20	30-40	-102	3-wave	current	reflection	degenerate
[48]	20	100	-	3-wave	flux	-	both
[49]	30	5	-93.1	-	flux	transmission	degenerate
[50]	20	10	-124	3-wave	current	transmission	non-degenerate
[51]	20	1	-	3-wave	current	transmission	non-degenerate
[52]	21	2.3	-190	3-wave	current	reflection	-
[53]	20	1	-	3-wave	flux	reflection	non-degenerate
[54]	28	3	-127	3-wave	current	-	non-degenerate
[46]	20	640	-110	-	current	reflection	non-degenerate
[55]	33	1.3	-103	4-wave	current	reflection	degenerate

but more stable gain is chosen.

A JPA has almost by default a rather narrow frequency band in which it can amplify. This is dependent on the quality factor of the SQUID cavity, which we want to be high to achieve a high gain and low losses, but we want it to be low to have a broad bandwidth. Several of the papers mentioned in Table 2.1, e.g. [46], feature specific JPA designs that have a much broader bandwidth than a conventional JPA. A broad bandwidth is desirable for many experiments, as the frequency of the signals of interest might not be known beforehand.

Two types of bandwidth can be mentioned in regards to a JPA. The flux can be tuned with a static magnetic field (biasing field), which changes the frequency of the cavity and thus moves the bandwidth of the amplifier. This is called tunable bandwidth, and can be on the order of GHz. However, this tuning can only be done slowly compared to the frequencies under study (GHz). Usually, we are interested in the instantaneous bandwidth, which is the bandwidth of our amplifier without changing the biasing field. This gives the range of frequencies which can be amplified simultaneously, and this is commonly the limiting factor for experiments. For most of the papers cited in Table 2.1, the instantaneous bandwidth is on the order of MHz on a GHz signal.

The maximum input power of a JPA is also of interest for practical reasons. The current through a JPA is limited, and this also limits the maximum power of the oscillation. If a strong signal enters the JPA, it gets amplified until its amplitude reaches the limit of the amplifier. That is, the gain of the amplifier decreased for higher signal strengths [45]. Typically, the power level where the gain decreases by 1 dB is taken as the input power limit.

The most important parameter missing in Table 2.1 is the added noise, or noise temperature of the amplifiers. The reason for this, is that it is often not reported in the papers, with as underlying reason that the noise temperature is hard to measure. This is treated more in-depth in Sec. 2.4.

2.4 Noise and noise temperature

Quantum noise has been the subject of many studies over the past years (see e.g. [56] for a review), and we will look only in very simple terms at it in this thesis. That is, we will only treat it in terms of the noise temperature of an amplifier, and refer to other works for more fundamental discussions of quantum noise of amplifiers (e.g. [10]).

2.4.1 Thermal noise

Thermal noise is caused by the thermal movement of the charge carriers in a conductor. For high enough frequencies and low enough temperatures, the thermal noise power goes to zero. In the systems used

in this thesis, this zero-temperature approximation doesn't hold. The power spectral density of thermal noise is multiplied with a factor $\eta(\omega)$, given as

$$\eta(\omega) = \frac{\frac{h\omega}{k_B t}}{e^{\frac{h\omega}{k_B t}} - 1}, \quad (2.67)$$

which is generally 1 for electronics at normal temperatures, but which drops to 0 for temperatures near absolute zero and frequencies in the GHz range. For this thesis, we work with $2\pi \cdot 6$ GHz signals at 250 mK, so $\eta \approx 0.5$. Although lower temperatures (10 mK) where $\eta \approx 0$ are quite within reach of modern dilution fridges, the fridge used for this thesis is limited to 250 mK.

This work focusses on amplifiers, for which the added noise is usually expressed as a noise temperature: The temperature where thermal noise would have the same noise power (spectral density) as the noise the amplifier adds to the signal. The noise temperature is given for a certain measuring bandwidth B (Hz) as

$$\frac{P}{B} = k_B T \quad (2.68)$$

with P being the noise power (W) and T (K) the noise temperature. This is also called Johnson-Nyquist noise [57].

2.4.2 Measuring noise temperatures

There are several methods commonly used to obtain the noise temperature of an amplifier, the Y-factor method, the gain-method and the noise-figure-meter method. They are all based on the same principle, namely that of sending in a signal at a known noise temperature and measuring the noise temperature of the output. Based on which parameters of the system under study are known (and to which accuracy), several variations can be used. The noise temperature of the amplifier can be extracted from

$$\frac{P}{B} = G k_B (T_{\text{input}} + T_{\text{amplifier}}). \quad (2.69)$$

where G is the amplifier gain, and the noise temperature of the input, T_{input} , is known.

If the gain of the amplifier is unknown (or not known very accurately), the Y-factor method can be used. It is based on two measurements, where the noise temperature of the input is varied. From the noise temperature of the output signal of each of the measurements, both the amplifier noise temperature and the amplifier gain can be found in principle.

If the gain of the amplifier is known (accurately), the added noise temperature can in principle be extracted from a single measurement where the input noise temperature is known. This is called the gain-method. This can also be done in an automated manner, using a calibrated noise source and a noise-figure-meter function of the spectrum analyser.

2.4.3 Noise temperature of an amplifier chain

For the specific system under study in this thesis, there are several issues with the measurement methods of the previous section. For one, we do not have a single amplifier. Instead, we have an amplifier chain, and we can only measure the noise temperature of the entire amplifier chain. The noise temperature of an amplifier chain is described by the Friis formula [21]. The total noise temperature of the amplifier chain, T_{chain} is determined from the noise temperatures of the amplifiers (T_1, T_2 , etc.) and their respective gains (G_1, G_2 , etc.) as

$$T_{\text{chain}} = T_1 + \frac{T_2}{G_1} + \frac{T_3}{G_1 G_2} + \dots \quad (2.70)$$

Extracting the noise temperature of the first amplifier requires knowing the gain of all amplifiers, and knowing the noise temperature of all other amplifiers.

For our specific amplifier chain, we know the noise temperatures and gains of the other amplifiers from their specifications. This should not be an issue, except that the noise temperature that we are interested in measuring is potentially very low, less than 1 K [58]. This means that any inaccuracy in the noise temperature or gain of the other amplifiers could affect the noise temperature obtained for our amplifier.

3

Design, fabrication and simulation

This chapter details the design and fabrication method of the structures and samples used in this thesis, as well as the measurement set-up and protocols used to characterize them. We also describe issues found in the fabrication process, and ways to resolve them. Finally, we also perform simulations of a general parametric amplifier with an additional Duffing non-linearity.

3.1 Facilities

The fabrication of the samples was done within the Kavli Nanolab, using the facilities of the Van Leeuwenhoek Laboratory, Delft, The Netherlands. The majority of the equipment used for fabrication is summarized below, and shown in Fig. 3.1. A brief description of the equipment function is given, more details can be found in textbooks such as [59].

1. The AC450 sputterer (Fig. 3.1a) is used to deposit metal on a sample chip. To do this, a plasma is created between electrodes in such a way that the ions (usually Argon) are accelerated towards a metallic target. The impact of the ions on the metal releases atoms from the surface of the target, which can then condense on the sample chip.
2. The Raith EBPG5000+ (Fig. 3.1b) is an Electron Beam Pattern Generator (EBPG) used to write patterns on a chip. A layer of electron-sensitive polymer (called a resist) is spin-coated on the chip. Impact of electrons changes the structure of the resist (i.e. the resist is exposed), which can happen in a controlled and localized manner by using a focussed electron beam. Exposed or unexposed parts of the resist layer can then be selectively removed using solvents.
3. The Leybold F1 (Fig. 3.1c) is used to etch away parts of the chip. Here, a plasma is made using SF_6 , O_2 or He gas, and the ions of this plasma are accelerated towards the chip where they can etch the sample using a combination of collisions and chemical reactions. This allows both chemical selectivity and some control over the (an)isotropy of the process. Depending on pressure and plasma power, this can result in vertical side-walls.
4. The Tepla (Fig. 3.1d) functions similarly to the Leybold F1, it also uses the ions from a plasma (O_2 or Ar) to etch away parts of the sample. However, this is mainly used to clean (remove organic material from) a chip.
5. The FEI Nova (Fig. 3.1e) is a scanning electron microscope, which is used to obtain high-resolution images of chips. Its role in fabrication is mainly to provide information on issues that might arise.

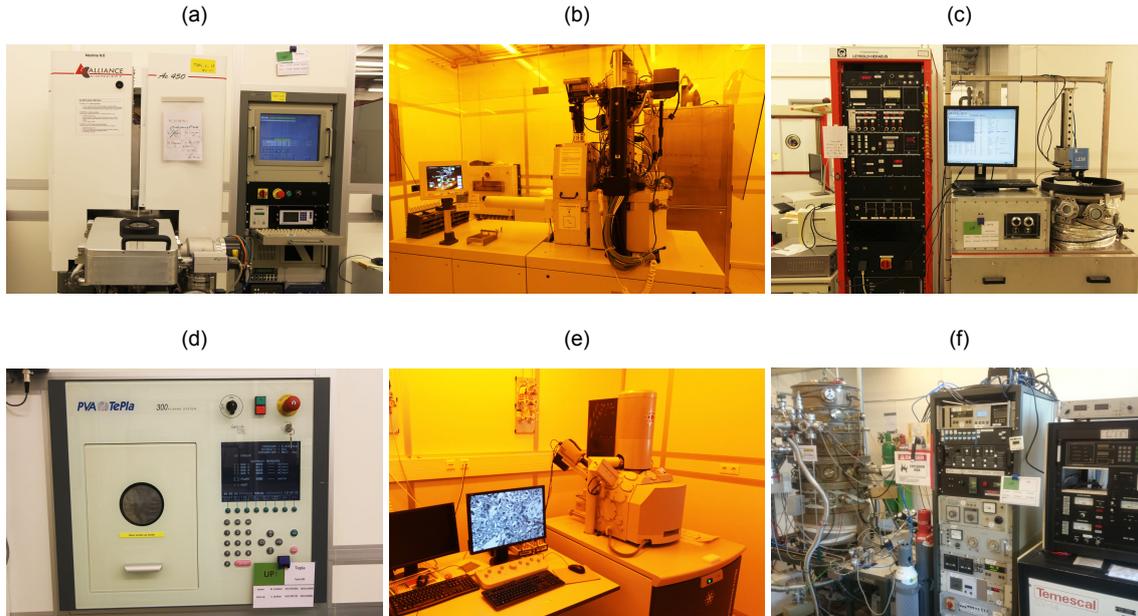


Figure 3.1: Overview of the machines used for fabrication. a)-e) are located in the cleanroom of the Van Leeuwenhoek Laboratory, f) is located outside the cleanroom. a) AC450, b) Raith EBP5000+, c) Leybold F1, d) Tepla, e) FEI Nova. f) UTS a)-e) sourced from [60].

6. The UTS (Fig. 3.1f) is an evaporator, which is also used to deposit metal on a chip. In contrast to the AC450 (sputterer) and evaporator uses a heater or electron beam to heat a metal source. The target chip is far enough away from the source for the atoms to be reasonably collimated. In contrast to the isotropic sputtering (metal atoms hit the sample from all directions), anisotropic evaporation allows good control of the incidence direction of the metal atoms.

Besides this set of equipment, there are multiple smaller machines used for different steps in our fabrication. We will explain in general terms the process of lithography, and the machines used for this.

Lithography is the process of creating patterns on a semiconductor substrate. Our designs are based on a planar, Silicon (Si) substrate which forms the back-bone of our sample. It is a small (usually 10×10 mm²) square of Silicon typically 500 μ m thick, with a layer of Silicon oxide (SiO₂) on the outside. An example of such a substrate, though with other structures already patterned on top, is shown in Fig. 3.2a. Typically, one side of the square is polished and reflective, and the other is not.

The substrate can be coated evenly with a light- or current-sensitive polymer (called a resist). To do this, we place a small amount of the (liquid) polymer in the middle of our substrate, and then spin it around to spread the polymer evenly. This process is called spin-coating. Typically, we have to bake the substrate and polymer so that it is not liquid any more. We can then use the light- or current-sensitivity of the polymer to create the pattern. We selectively expose some areas of the polymer to light or current (this refers to optical and electron-beam lithography, respectively), and use the chemical change that either of these two processes induce in the polymer to selectively remove parts of the polymer layer (this is called developing, as with photographic film). With this pattern in the resist, we can then controllably remove (etch) or deposit metallic or semiconductor parts on the chip. This determines the functionality of our chip.

3.2 Design and fabrication

This section details the designs created in this thesis, and the fabrication thereof. There are three different designs used: The first one is a design for a set of Josephson junctions and SQUIDs for testing the properties of these structures. The second is a design for a parametric amplifier, and the third

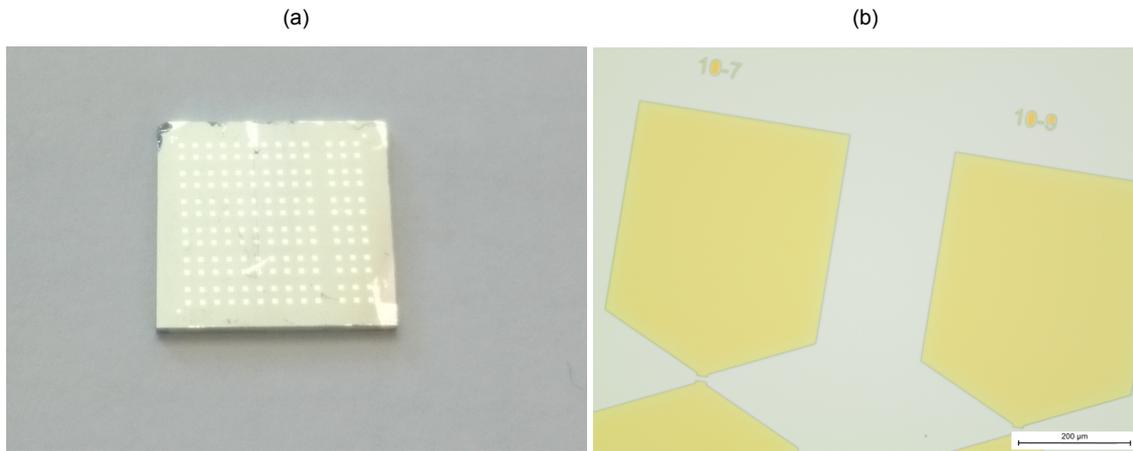


Figure 3.2: a) Photograph of a typical $10 \times 10 \text{ mm}^2$ Silicon chip used in this work. The bright squares are metal contact pads. The other brighter areas are pieces of metal that were left over from fabrication, the dark spots near the top are the edges of the chip damaged during handling. b) Optical microscope image of several tapered rectangular contact pads. They consist of a 100 nm thick alloy of Molybdenum-Rhenium (MoRe). The grey area is the Silicon substrate.

is a different design for a parametric amplifier using a different fabrication method. We will treat the fabrication processes when it is relevant for the design. The last subsection details the Josephson junctions, as they are used in all designs and their fabrication is an important part of this thesis.

3.2.1 Junction and SQUID test design

To test the properties of our junctions and SQUIDs, we need to be able to connect these structures electrically to our measurement equipment. To do this, we use contact pads, which is simply a flat area made of metal that is large enough to connect a wire to, either by hand or by a precision manipulator. In this design, the contact pads are tapered squares made of a 100 nm thick layer of a Molybdenum-Rhenium (MoRe) alloy. An example of a contact pad is shown in Fig. 3.2b.

Each junction or SQUID requires at least two connections, one to send a current or voltage in, and the other to ground. We can place our junction or our SQUID between two contact pads, and characterize it through them. In Fig. 3.2b, we see part of the second contact pad of the pair at the bottom left of the figure, and we can also see the small gap between them where the junction will be fabricated. This way, we require twice as many contact pads as we have junctions (or SQUIDs).

The contact pads we use are large with respect to the size of our substrate, and take some time to pattern, so we would like to reduce the amount of contact pads we need. As one of the contact pads of the junction or SQUID is grounded, we can connect the grounded pads for all the structures we want to test and replace it by a single common ground pad. However, we still need to ensure an electrical connection from all the pads to the junctions, and then from all the junctions to the common ground. We can do this by simply patterning a strip of conductive metal (called a wire, or a line) from the tapered end of the contact pads to the junctions, and then from the junctions to the common ground.

To be able to easily connect to our contact pads from outside our sample, they are spaced symmetrically around the outside edge of the chip. The design shown in Fig. 3.3a features 32 contact pads in four sets symmetrically placed around the chip. One of the connectors of each of the four sets is intended as a common ground, and it is directly connected to a $20 \mu\text{m}$ wide metal strip (a wire) on the chip. The other contact pads taper down to a $20 \mu\text{m}$ wide area where one end of the junctions or SQUIDs makes contact. The other end of the junctions connects to the common ground wire. A close-up of this area of the design is shown in Fig. 3.3b, where a $5 \times 2 \mu\text{m}^2$ junction (red) bridges the gap between the contact pad (left) and ground wire (right). To ensure proper connection, a $2 \mu\text{m}$ overlap is designed between the

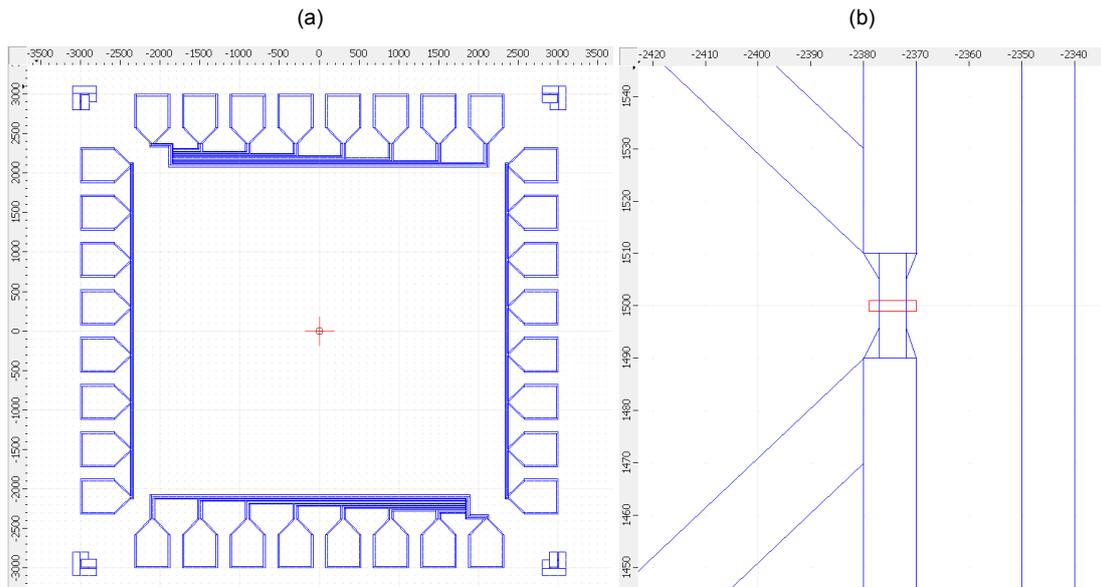


Figure 3.3: a) Overview of the junction test design showing four sets of contact pads with connections between them. The junctions must be oriented in the same direction due to fabrication constraints, this causes two of the arrays to have a different layout. b) Close-up of a single $5 \times 2 \mu\text{m}^2$ junction (red) for DC-measurements. On the left side, the tapered end of the contact pad is drawn, while the right side shows the wire connecting to ground. Units of the figures are in μm .

junction and each of the sides.

3.2.2 Parametric amplifier design

The design for our parametric amplifiers (Fig. 3.4a) resembles the previous design for testing junctions and SQUIDs. It features the same four sets of contact pads with a common ground, but it includes several other components which we will explain.

First, the size of the substrate used for this design is larger. The design still fits in a $10 \times 10 \text{ mm}^2$ area, but should be fabricated on a larger chip and then diced to size to prevent effects from edge- and corner beads. These beads form during fabrication, when we spin-coat our resist, and are unwanted.

While the junction and SQUID test design was built for DC measurements, the parametric amplifier design features several connector fit for GHz-frequency connections (AC connections, also referred to as RF, Radio Frequency, lines). We will mainly operate the amplifier in the 4-8 GHz range, which is in the microwave regime. The connections are referred to as microwave connections, or microwave lines.

There are four microwave connections from the measurement devices to the structures of our parametric amplifier, via the four big contact pads (two on the left side of Fig. 3.4a, and two on the right). They have a tapered square shape similar to the DC contact pads, but are larger, $500 \times 400 \mu\text{m}^2$, and taper down to $10 \mu\text{m}$ width. One of these microwave connections is shown in Fig. 3.4b. The connectors are spaced symmetrically around the design. The tapered end of the microwave connection is connected to the center conductor of a coplanar waveguide. This coplanar waveguide is visible in Fig. 3.4a as the line with the bend in the middle.

Fig. 3.4b also features a dicing marker (large chevron in the top left corner), and a set of $20 \times 20 \mu\text{m}^2$ alignment markers (the three small squares inside the large red square). The dicing markers are used to align the saw used for dicing the chip to size at the end of the fabrication process, the alignment markers are used to align the different layers of the pattern. Both the dicing and the alignment markers are placed symmetrically at the four corners of the design.

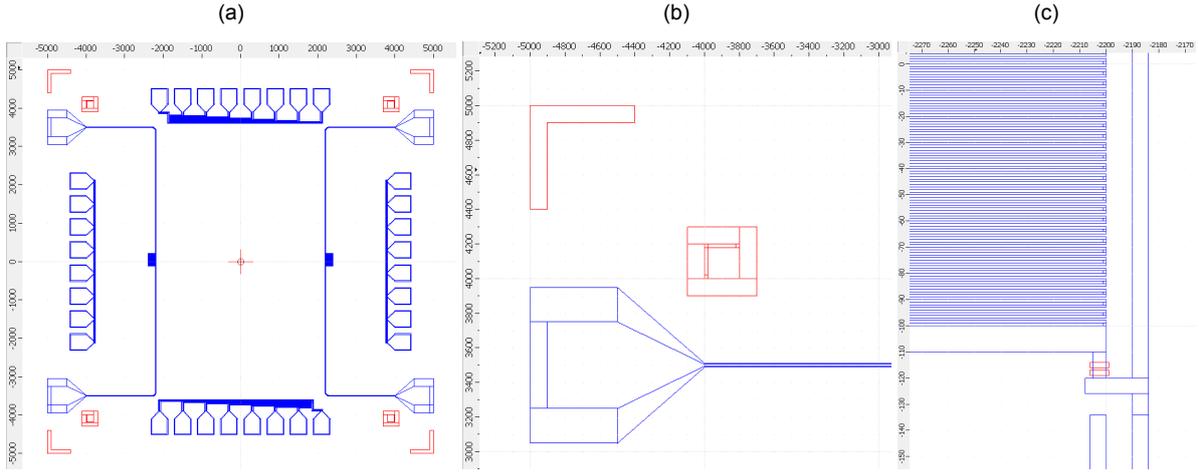


Figure 3.4: a) Overview of the design of the parametric amplifier, showing two sets of two microwave connections, each connected to a single SQUID cavity, the small rectangles near the middle of the design. Also visible are four sets of SQUIDs and junctions, similar to the previous design. b) Close-up on one of the microwave connectors (blue), and a set of red alignment (square) and dicing markers (chevron). c) Close-up on one of the SQUID cavities. The top left shows an interdigitated capacitor stretching out of the image connected on the right side to the signal line. The signal line is connected to ground via a SQUID (red). From below, an approaching flux pump line is also shorted to ground. The scales of the image are in μm .

Another circuit element present in the design of our parametric amplifier is the interdigitated capacitor. This consists of a set of long, rectangular pieces of metal colloquially called the fingers (digits). One half of the set of digits is connected on one end to the center conductor of a coplanar waveguide, and the other half of the set is connected to a ground plane. In this design, the entire area of the chip that appears empty in Fig. 3.4a is in fact a grounded metal plane. If the digits of the structure are connected in an alternating fashion to either of the two sides, they form a capacitor.

The interdigitated capacitor in this design consists of 200 separate $200 \times 1 \mu\text{m}^2$ fingers, half of them are connected to the transmission line and the other half are connected to the ground plane. A part of the interdigitated capacitor in this design is shown in Fig. 3.4c. Here, the right side of the capacitor (the washboard-like pattern) is connected to the center conductor of the transmission line. Half the digits are connected out of view of the figure to the grounded plane.

The interdigitated capacitor forms the C of our LC-oscillator. The L is formed by a SQUID, two $7 \times 2 \mu\text{m}^2$ junctions, which are shown in red in Fig. 3.4c. One end of the junctions is connected to the (end of the) center conductor of a the transmission line, the other is connected to the grounded plane. These two elements together form our SQUID cavity.

Below the SQUID in Fig. 3.4c, we can see the end of a different transmission line. This line is directly connected to the ground plane next to the SQUID. The connection here is also slightly more narrow than the rest of the center conductor. This was done to increase the inductance of this part of the line, so a current through the transmission line generates a strong magnetic field close to the SQUID. This field is then varied over time, and this changes the inductance of the SQUID. This is used to pump the parametric amplifier, and is referred to as the flux pump line.

The parametric amplifier design involves two main layers. The first is a 100 nm thick Molybdenum-Rhenium alloy (MoRe) layer, which is sputtered directly on top of the cleaned Silicon substrate. This layer contains all the contact pads, the transmission lines (coplanar waveguides), the interdigitated capacitors and all the markers. The pattern is etched into the MoRe layer using a reactive ion etcher (Fig. 3.1c). The full recipe containing the parameters of all the fabrication steps is given in Appendix A.

The second layer is the Aluminium layer, which is fabricated in a different way. First, the pattern is created in the resist covering the chip. Then, Aluminium is evaporated onto the entire chip. Most of the Aluminium is then removed by lift-off, a process where the resist is chemically removed and this releases the metal deposited on top. Only the metal that was deposited in the areas where the resist was selectively removed (i.e. the pattern) remains on the sample. The Aluminium pattern layer contains the Josephson junctions.

Impedance matching a coplanar wave-guide

The parametric amplifier design is the first of our designs that contains transmission lines (coplanar waveguides) to transmit GHz-frequency radiation. It is important to pay attention to the design of this circuit element. We want to prevent unwanted reflection of the microwave signal at the interface between the chip and the rest of the cabling, and to do so the coplanar wave-guides in this design must be impedance matched the rest of the cabling, i.e. to $Z_0 = 50 \Omega$.

To do this, we have simulated a straight piece of wave-guide in Avago AppCAD, where we have taken as simulation parameters the substrate height $535 \mu\text{m}$, the thickness of the metal layer $100 \mu\text{m}$, the width of the center conductor $10 \mu\text{m}$, the width of the gap $6 \mu\text{m}$ and the relative dielectric constant of the substrate $\epsilon_r = 11.9$. This gives a characteristic impedance of $Z_0 = 50.1 \Omega$.

Estimating the resonance frequency of the SQUID LC oscillator

In the design phase, it is critical that we pay attention to the resonance frequency of the SQUID cavity. To do this, the design was simulated in Sonnet. To reduce computation time, only a limited area, $400 \times 500 \mu\text{m}^2$, shown in Fig. 3.5a, of the total design ($10 \times 10 \text{mm}^2$) was taken into account. The simulation consisted of a signal transmission line coming from the top, connected to half of the fingers of the Interdigitated Capacitor (IDC, middle), with the other half being connected to ground. The signal line is also connected to one end of the SQUID, which is modelled as an ideal, linear inductor, as shown in Fig. 3.5b. A pump line is shorted to ground just below the SQUID. The simulation area is surrounded by a ground plane.

All the material was modelled as an ideal lossless material. The substrate was modelled as $500 \mu\text{m}$ thick intrinsic Si with a relative dielectric constant of 11.9. The top side was modelled as 1 mm vacuum. It is difficult to simulate the correct inductive behaviour of a SQUID modulated by an oscillating magnetic field. To go around this limitation, we parametrized the SQUID inductance and simulated the circuit behaviour for various values of this parameter.

The result of the simulation that is of interest for us is the magnitude of the reflected signal S_{11} . We have simulated the circuit response between 1 and 15 GHz. The resulting S_{11} traces show a broad dip representing the cavity resonance, which decreases in depth as we increase the SQUID inductance, as expected. The dip magnitude is dependent on the modelling of the losses in the circuit, which are not taken into account accurately in our model.

From the simulation, we observe that to get the resonance frequency of our cavity in the desired range of 4-8 GHz, our SQUID inductance must be between 0.1 and 0.7 nH. This value can be compared to the expected inductance of our Josephson junctions in Sec. 3.2.4.

We are interested in which fraction of the inductance of our SQUID cavity comes from the SQUID. The fraction of the total inductance that is contributed by the SQUID inductance is called the SQUID participation ratio. This determines how much the resonance frequency can change if the SQUID inductance is modulated by a magnetic field. If a significant portion of the total inductance of the cavity comes from the geometric inductance of the circuit, and is thus not changed by the magnetic field, it does not contribute to the parametric amplification. If the SQUID participation ratio is large, the amplitude of our magnetic pump field is small. This is desirable, as unwanted heating of our sample can occur if the pump field amplitude is large.

We can extract the geometric inductance of the SQUID cavity by varying SQUID inductance param-

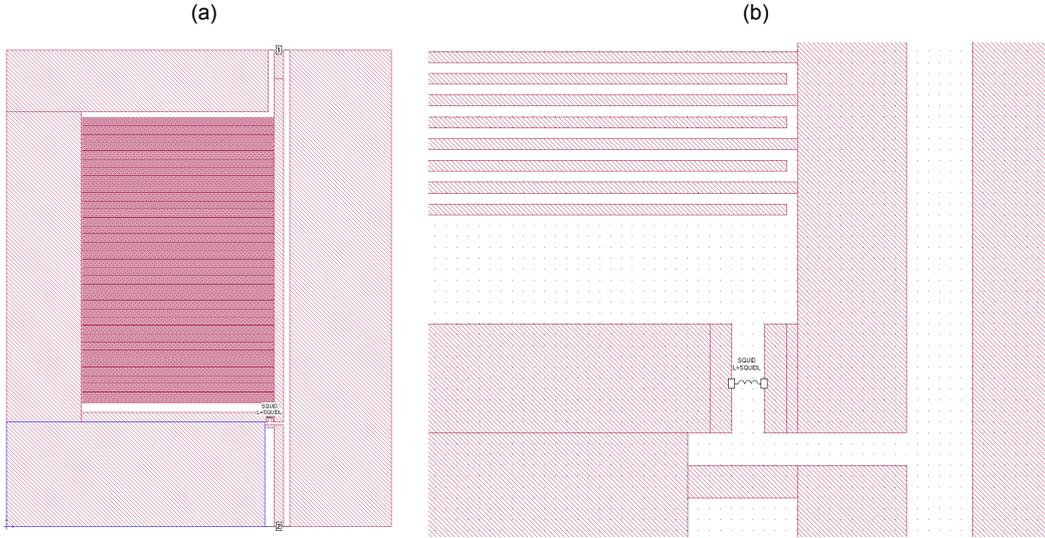


Figure 3.5: a) Design of the parametric amplifier used for simulations in Sonnet. b) Close-up of the area around the SQUID. In both figures, red represents a lossless material on top of a Si substrate. The SQUID is modelled as a lumped element inductor, whose inductance is a parameter to be varied in the simulations.

eter. We plot the inverse of the resonance frequency squared, $1/\omega^2$ against the SQUID inductance L_{SQUID} in Fig. 3.6b. We perform a linear fit, where the crossing with the x-axis (inductance) gives the geometric inductance contribution of the cavity,

$$\frac{1}{\omega^2} = C(L_{\text{geom}} + L_{\text{SQUID}}) \quad (3.1)$$

We obtain value for the geometric inductance L_{geom} of 238 pH. With this value, and the expected SQUID inductance value, we can calculate the SQUID participation ratio.

3.2.3 Transmission cavity design

A third design was used for the parametric amplifier measurements, which was necessary due to machine access issues. Most of the design elements present in the parametric amplifier design, are also present in this one. We will highlight the differences.

The transmission cavity design features a set of seven cavities, which are capacitively side-coupled to the same transmission line. Thus the cavities are measured in transmission, instead of in reflection mode such as the parametric amplifier design of the previous section. Our cavities consist of seven differently-sized interdigitated capacitors, with an equally-sized SQUID. One such cavity is shown in Fig. 3.7a. This figure shows at the top the center conductor of a transmission line, where another (smaller) interdigitated capacitor capacitively couples our SQUID cavity to the transmission line. This capacitor has $17 \times 25 \times 1 \mu\text{m}^2$ fingers for all of the different cavities.

The main interdigitated capacitor providing the bulk of the capacitance of the cavity is in the middle of Fig. 3.7a. This capacitor is symmetrical around a middle connector, features $100 \times 1 \mu\text{m}^2$ fingers. The number of fingers varies between the different cavities. This is to have cavities with different resonance frequencies, so we can separately address them by sending in signals at their resonance frequency.

Each cavity is also inductively coupled to a flux-pump line, similarly to the previous parametric amplifier design. Here, there is only one flux-pump line for all seven cavities, so it is not terminated to ground at each one. The flux-pump line is shown at the bottom of Fig. 3.7a, with just above it the SQUID loop. A close-up of the area around the SQUID loop is given in Fig. 3.7b. This shows at the bottom the flux-pump line, in the center the SQUID loop. The SQUID loop is interrupted by two constriction-type Josephson junctions, which are visible as a narrowing in the line. The SQUID is connected to the ground plane, as

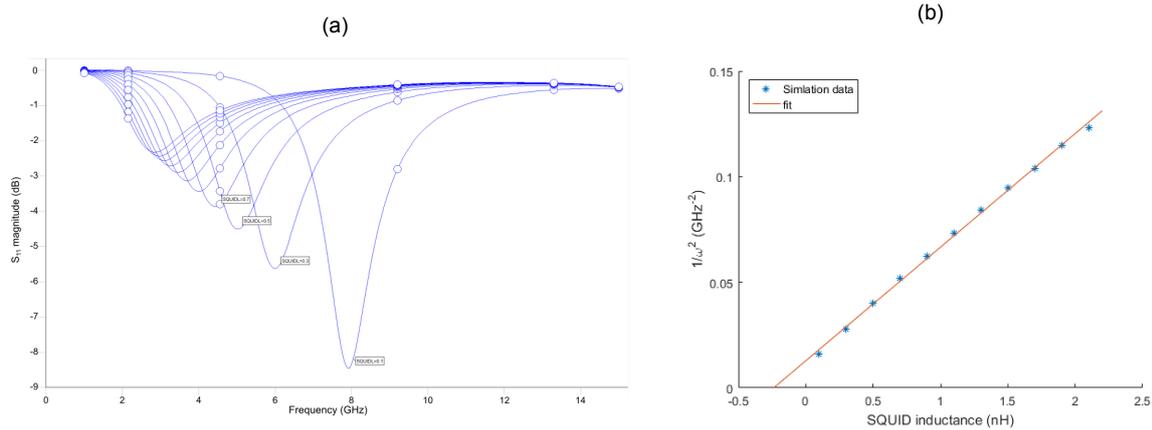


Figure 3.6: a) Reflected signal power S_{11} between 1 and 15 GHz for several values of the SQUID inductance. As expected, for lower SQUID inductance the resonance frequency increases. b) Plot of resonance frequency versus the lumped-element SQUID inductance. The crossing of the fit with the x-axis gives the geometric inductance.

with out previous design, but here this connection is through a wire with bends to create an additional inductance. This is done to lower the cavity frequency.

3.2.4 Josephson junctions

Josephson junctions form the most critical part of the structures used in this thesis. We will treat the fabrication method we use in more detail, and compare it with the conventional method of fabricating our type of junctions.

Fabrication of junctions in general

There are several different methods for fabricating junctions, dependent on the type of junction, the materials involved and the desired geometry. All junction methods share the same principle, they all result in two superconductors being weakly connected. It is possible, for example, to use a crystal grain barrier as a boundary between two superconductors. Or, another example, a junction can be made using only one point contact between two superconductors. In this thesis, we will focus on two fabrication methods based on having a metal oxide as the barrier between the two superconductors.

These methods are the (conventional) Dolan-bridge method [61], and an experimental bridge-less shadow evaporation method ¹ These two methods are suitable for the type of junctions that we are interested in, which are superconductor-insulator-superconductor junctions where the superconductor is Aluminium (Al), and the insulator is Aluminium oxide (Al_2O_3). These junctions are suitable for operation in the mK temperature regime.

Both methods of fabrication are based on the same principle of double-angle evaporation with a shadow. To create the superconductor-barrier-superconductor structure of a Josephson junction, we evaporate metal at two (different) incident angles, and let the sample oxidise in between the evaporations to create the barrier. This method requires a 'shadow', a part where no superconducting material is deposited during the evaporation. This is necessary to separate the electrical connection between the two superconducting parts. Aluminium is a convenient and common material for this because it oxidises into Al_2O_3 , which forms only a suitably thin layer that stops further oxidation.

We will first explain the two fabrication methods, the Dolan bridge method and the bridge-less shadow method, and afterwards discuss their weaknesses and strengths.

¹This idea to use the shadow of the electrodes themselves (instead of a free-hanging bridge) to create the junctions originates from the group of Prof. DiCarlo at TU Delft. However, it is published nowhere yet to the best of our knowledge.

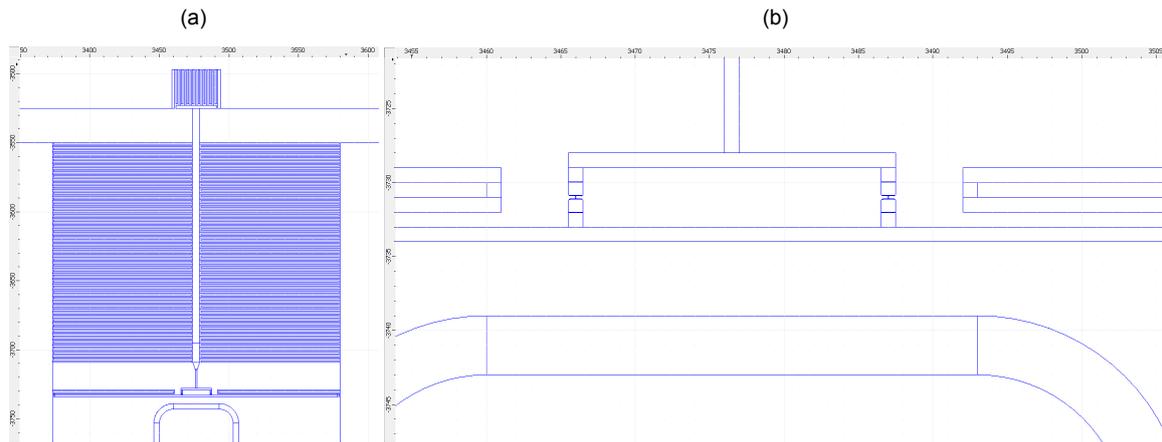


Figure 3.7: a) Overview of one of the cavities in the transmission cavity design. At the top, a coupling capacitor (interdigitated) to the signal transmission line is shown. The large capacitor (interdigitated) in the middle forms the bulk of the capacitance of the cavity. Below that, the SQUID loop is visible. At the bottom, the pump line is shown. b) Close-up of the SQUID loop area of the transmission cavity design. Visible are the two junctions in the SQUID loop. The lines branching off to the left and right are bent to generate extra inductance. The scales in both images are in μm .

The Dolan bridge method

The Dolan bridge method [61] is based on a narrow, free-hanging bridge creating a shadow for the evaporation. This is schematically shown Fig. 3.8, the left process.

- Two resist layers (light and dark blue) are spin-coated on a substrate (gray). A pattern is transferred to the top resist (hatched rectangles)
- Both resists are developed. They have been chosen such that we get an undercut that is sufficient to fully release the middle part of the top resist layer. This is the free-hanging bridge.
- Under a certain angle, the first layer of metal (red) is evaporated onto the sample.
- The outer layer of the metal oxidises (black), and then the other layer of metal is evaporated onto the sample under a different angle.
- The remaining resist is removed, taking the bridge structure with it, and leaving only the metal and oxide on the substrate.

The bridge-less method

In this thesis, we also use a bridge-less method. The shadow is generated by the electrode side-walls. This process is shown on the right side of Fig. 3.8. This method does not require the resist to be on the chip, it can be done completely without it. It is only present in the sketch because it is also present when we perform the evaporation, as we use it for lift-off.

- Two layers of resist (light and dark blue) are spin-coated onto the sample, where electrodes (light gray) are already patterned and present. The electrodes are separated by a gap, where the resist touches the substrate (dark gray). The resist layers are patterned (hatched region) with a slightly larger area than the gap between the electrodes.
- The exposed area is developed. The different resist layers are necessary to generate the required undercut for lift-off.
- Metal (red) is evaporated under the first angle. It connects only to one of the electrodes.
- The first metal layer is oxidized (black), and then the second layer is evaporated at a different angle. This layer connects to the other electrode.
- The resist is lifted off, and the Josephson junction is complete.

Strengths and weaknesses of the fabrication methods

Now we discuss the strengths and weaknesses of the individual methods. The Dolan-bridge method is

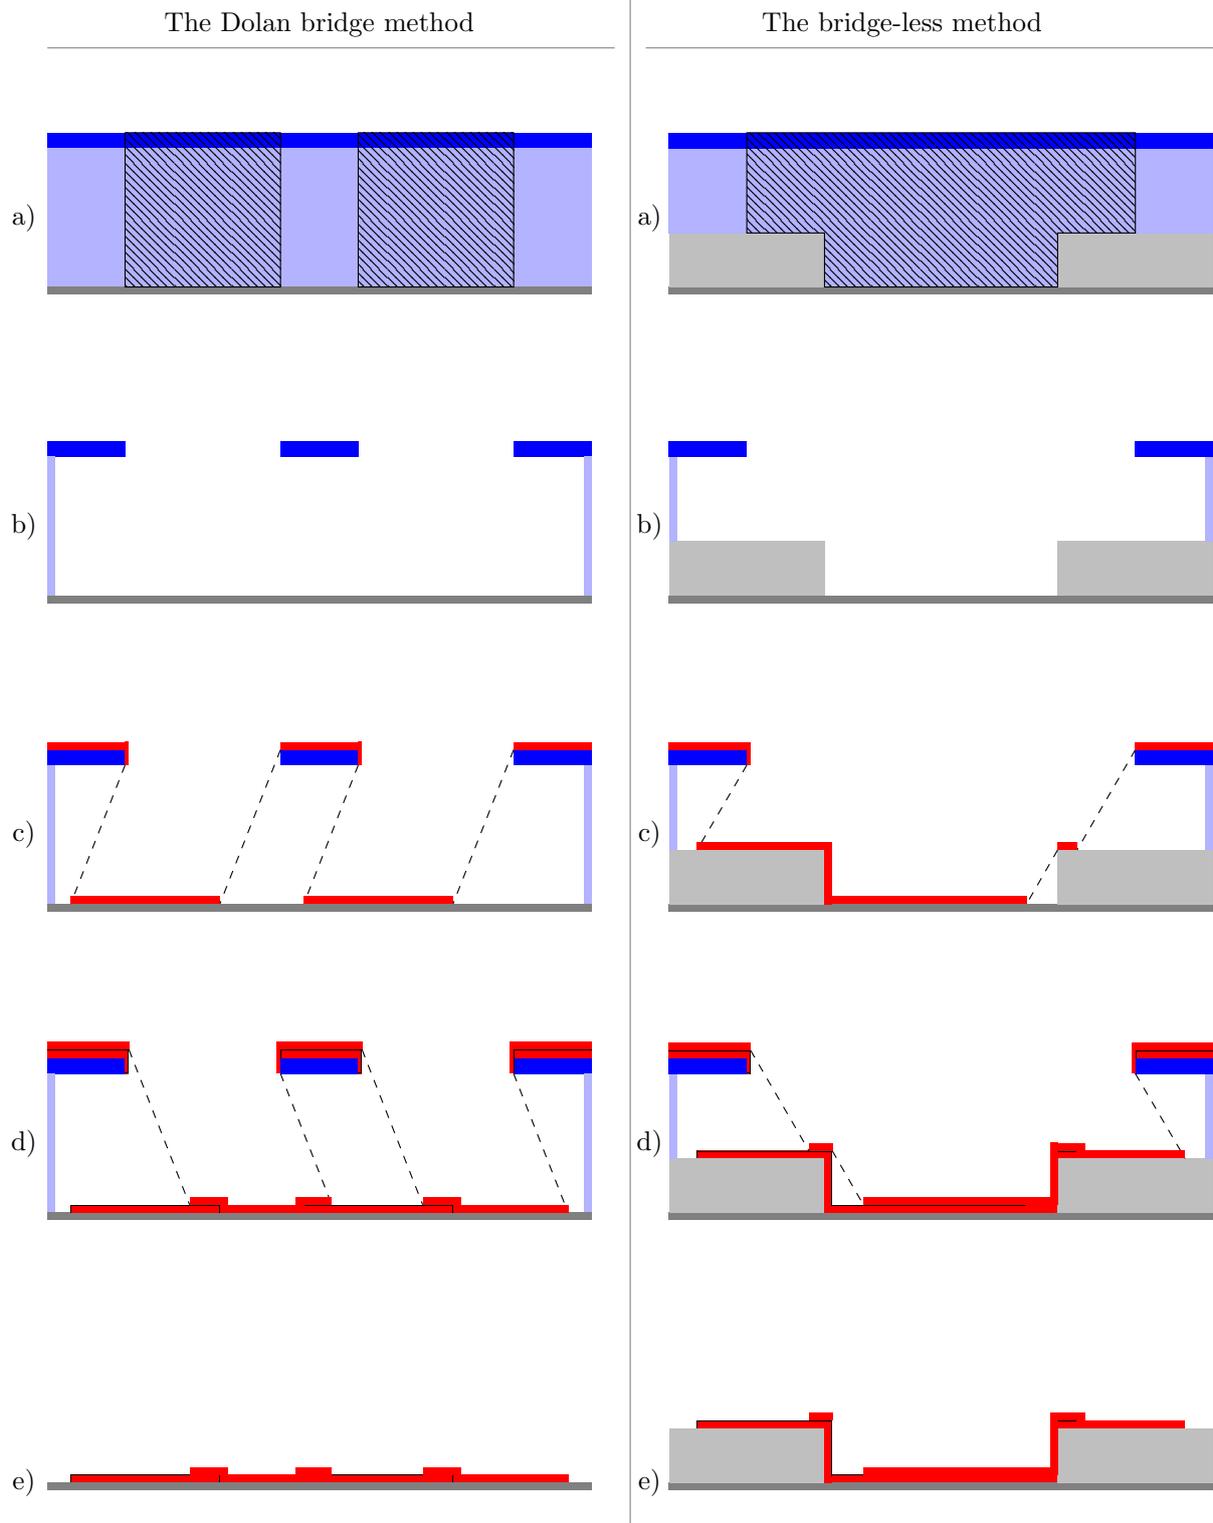


Figure 3.8: Two ways of fabrication Josephson junctions: left part of the image shows the Dolan-bridge, right part shows the bridgeless shadow method. Explanation of the different steps is given in the main text.

already established, so process parameters are well understood and can be controlled, while the bridge-less shadow method is new and process parameters are not well known. However, the Dolan-bridge method is not very flexible regarding the design of the junction size. For the bridge-less shadow evaporation method, the junction area is determined by the gap between the electrodes and the two angles of evaporation, while in the Dolan-bridge method it is determined by the distance between the bridge and the substrate, and the width of the bridge, and also the two angles of evaporation.

This flexibility in designing the junction area is desirable in our case. We want our JPA to be able to amplify even for a relatively large input signal strength. To do this, we need our junctions to have large critical currents. Josephson junctions with a larger area tend to have a larger critical current, though the increase of critical current with junction area diminishes if the area becomes too large [27]. The bridge-less evaporation method allows us to be more flexible in the junction size than the Dolan-bridge method.

Furthermore, the bridge-less shadow method requires (typically) one less fabrication step. With the Dolan-bridge method, the contact electrodes need to be patterned on the end parts of the junction. The contact between the electrodes and the junction is already ensured automatically in the bridge-less method.

Josephson junction design

To be able to design a JPA, we need to know several parameters of our junctions (and SQUIDs). Due to machine access and time constraints, there were no junction test performed for the bridge-less method before a JPA design was created. We have based our design decisions on junctions fabricated with the Dolan-bridge method in a different system than the one used in this work (the UTS). In particular, we have based our design decisions on the critical current density of these junctions, which is $0.7 \mu\text{A}/\mu\text{m}^2$.

With this critical current density, we can calculate the critical current for our junction sizes, under the assumption that they are small enough that the critical current density will not change by too much due to the magnetic field it generates [27]. The smallest junctions designed are $1 \times 1 \mu\text{m}^2$, and should have a critical current of $0.7 \mu\text{A}$. This would give a Josephson inductance of 471 pH . The largest junctions designed are $6 \times 4 \mu\text{m}^2$, should have a critical current of $16.8 \mu\text{A}$ and this would give a Josephson inductance of 20 pH .

3.3 Fabrication issues

Several issues were found and resolved during the fabrication process. As they are an important part of this thesis, but not part of the characterization of a Josephson Parametric Amplifier, they are described here instead of in Chapter 4.

Etch sensitivity

There were issues with the etching step, where the pattern is transferred into the MoRe layer. The etch process is a SF_6 and O_2 plasma etch, which also etches SiO_2 . If the etch is continued for too long, this would prevent the evaporated Aluminium to connect to the MoRe electrodes.

To prevent over-etching the sample, a laser was used to measure the reflection of the surface. MoRe is more reflective than SiO_2 , and the reflection changes in a certain pattern when etching. After the etch is started, the reflection goes up slightly, and then remains constant for most of the etch. Near the end of the etch, it drops down until it stabilizes to a lower value, and that is when the etch should be stopped.

Another issue is present due to the etching process, because the SF_6 and O_2 plasma has a non-negligible etch rate for the resist used. If the etching takes too long, the etch mask can be removed completely and the MoRe layer is etched in unwanted places. This issue was solved by spin-coating at lower speeds, to increase the etch mask thickness.



Figure 3.9: a) One end of the sample is etched successfully. Shown is part of an array of contact pads, the top of which serves as a common ground for the rest. b) The other end of the same array of contact pads (and connections between them) is under-etched, there is still MoRe covering the Si where it should have been removed.

There is a third issue related to the etch step of the fabrication. The etch rate can depend on the location in the reaction chamber of the etcher. This could be due to local variations in the plasma parameters, potentially related to gas flow or potential difference of the charged plates creating the plasma. This sensitivity of the etch process to the location of the chip in the chamber can lead to issues where one end of a sample is fully etched, while the other is not, such as in Fig. 3.9 where one end of a 15×15 mm² chip is etched sufficiently, while the other half is not. There is still MoRe left, though some etching has taken place as the pattern is visible. This issue can be fixed simply by not etching more than one chip at a time, and by ensuring that the sample is in the center of the etching chamber.

Electrode side-wall angle

We have a trio of fabrication issues (this paragraph, and the next two) that share a common origin. When we evaporate Aluminium, we expect the atoms to only hit the sample from one direction, and we assume they do not move too much from the position where they first hit the sample. However, all three issues show that Aluminium does in fact move around somewhat after it hits the sample. We will first comment on each of the three issues separately, and afterwards discuss how we can ensure that Aluminium sticks to where it lands on the sample.

The deposition of the Aluminium for the junctions requires the shadow of the MoRe sidewall. If the sidewall angle with respect to the substrate is rounded by the etching, this prevents the fabrication of junctions. During the evaporation step, the rounded side-walls don't cause a shadow, which leads to the Al simply connecting to both electrodes instead of forming a junction. This was observed, as shown in Fig. 3.10a, where we would expect to see the separation between the Al layers (as visible in the top of the figure) below the side-wall as well.

We have sketched the situation in Fig. 3.10b, where the dark gray represents the substrate, the lighter gray represents the MoRe electrodes, the red is the Aluminium of the junctions and in black we have Al₂O₃ as the insulator separating the layers. The arrows denote the angle of incidence of the evaporation.

This issue could have been solved by evaporating at higher angles with respect to the surface normal, though this uses more Aluminium. Instead, the etch parameters were fine-tuned to etch more an-isotropically. This makes the electrode side-wall closer the substrate normal, and ensures we can have a shadow for our Aluminium evaporation.

Top-layer thickness

There was an issue where the top one of the two evaporated Aluminium layers did not conform to

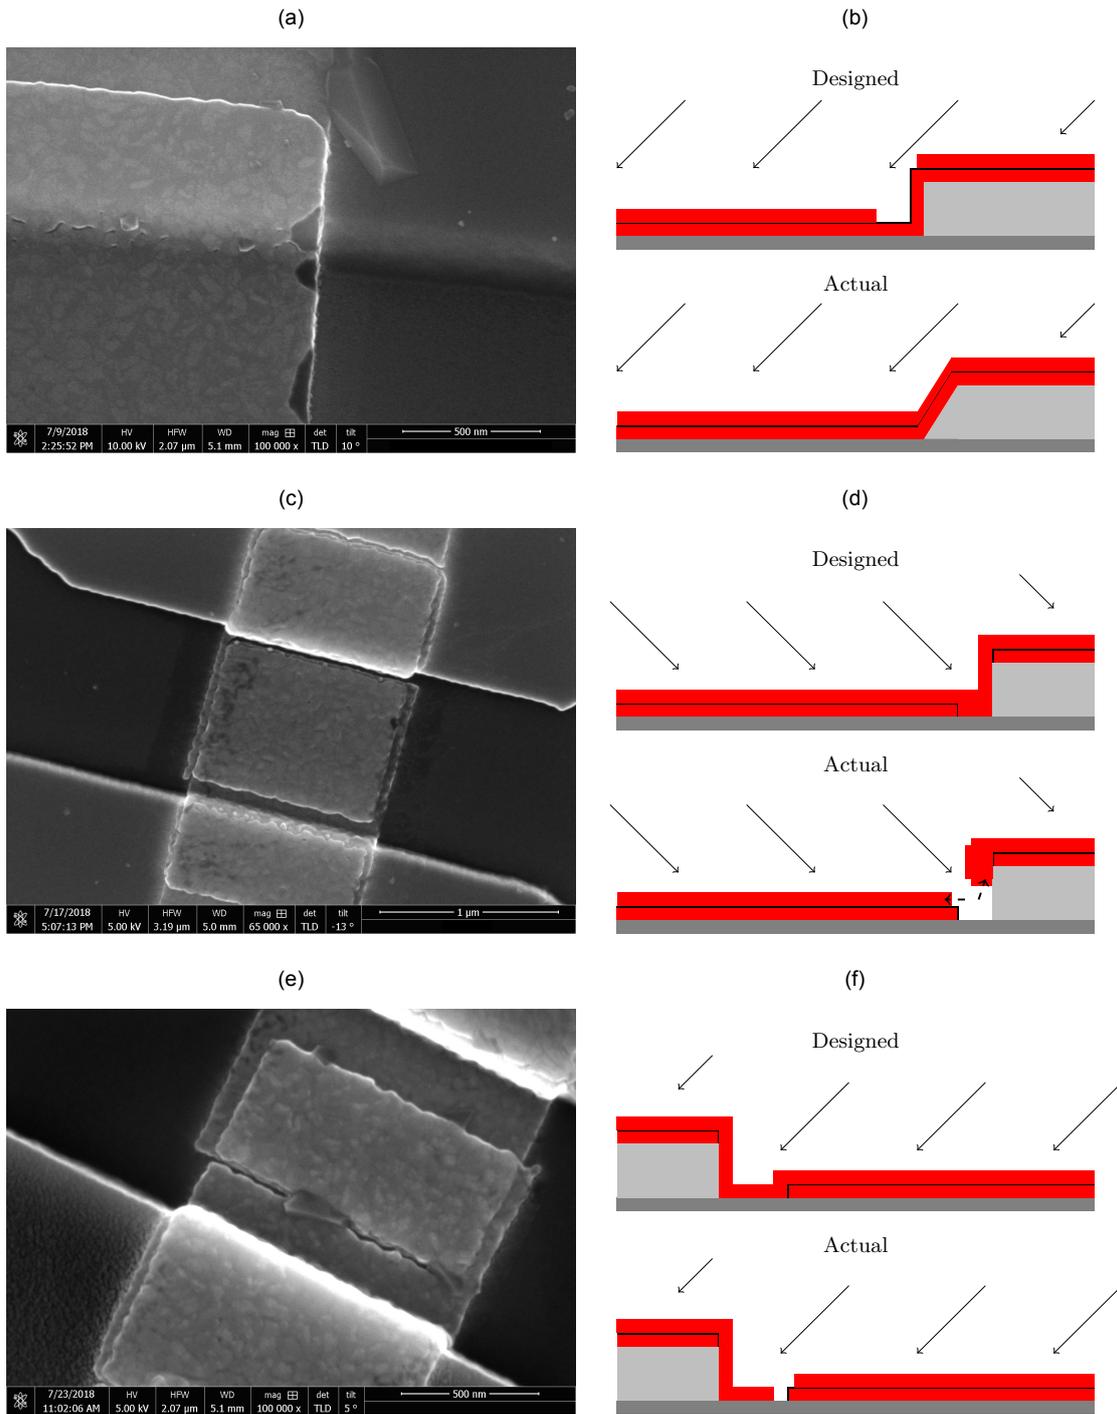


Figure 3.10: a) SEM image of the side-wall of the MoRe electrode, with two layers of Aluminium evaporated on top. Taken at an angle, it is visible that the side-wall is not angled vertically to the substrate. There was no shadow created by the side-wall to disconnect either of the evaporated layers. b) Sketch of the situation in a), where dark gray represents the Si substrate, light gray the MoRe electrodes, red the Al, black the Al_2O_3 , and red the angle of incidence of the evaporated Al. c) SEM image of a junction where the top layer does not conform to the electrode side-wall, at the top of the image both layers are disconnected from the electrode. The two layers are visible due to a slight misalignment during the evaporation step. d) Sketch of the situation in c). The top layer does not conform to the side-wall, and causes both layers to be disconnected. e) SEM image of a junction where the bottom layer causes a shadow for the top layer, so that the top layer is also disconnected. f) Sketch of the situation in e), where the shadow of the bottom layer disconnects the top layer.

the bottom one, and was therefore unconnected to the MoRe electrode on both sides of the junction. The bottom layer is evaporated on flat Si, but the top layer must conform to the gap from the side-wall shadow, as sketched in Fig. 3.10d. The Aluminium layers are visibly disconnected in Fig. 3.10c near the top of the junction structure. We can see that the top layer of the Al has some overhang over the MoRe electrode. This tells us that the Al we expect to have been deposited in the gap between the bottom Al layer and the MoRe electrode has moved from where it has been deposited by the evaporator.

This issue could be fixed by increasing the thickness of the top Aluminium layer, or by decreasing the thickness of the bottom one. Alternatively, Aluminium should move around less after deposition on the Si surface if the temperature is lower. This could be achieved by cryo-cooling the sample via the sample holder.

Bottom-layer shadow

The third issue that has to do with the Aluminium moving around after it hits the sample is that the bottom junction layer can also create a shadow when the top layer is evaporated. This breaks the connection of the top layer, as is shown in Fig. 3.10e and sketched in Fig. 3.10f. This should not be able to happen if the top layer is thicker than the bottom layer, as is the case for these evaporations, but it does happen. This is possible if the Aluminium is able to move around just after it has hit the sample. If the Aluminium atoms prefer to clump together, we get an extra gap at the bottom electrode.

This issue is also related to the evaporation angle of the top layer with respect to the bottom layer, as well as the thickness of this layer. By decreasing the angle of evaporation of the top layer, or increasing its thickness, the issue can be resolved.

To conclude these three paragraphs, we have seen issues partially caused by the Aluminium having the ability to move around somewhat after it hits the sample. We expect movement to be related to the kinetic energy of the incident Aluminium atoms, which could heat up the sample during the evaporation. At higher temperatures, the Aluminium atoms can move around more, so as a solution to these issues, we could try to cool down the sample during evaporation. We would have to ensure a good thermal contact between the sample and the sample holder if we use the sample holder for cooling, but this can be achieved by using, for example, carbon tape between the sample and the sample holder.

MoRe oxidation

MoRe oxidises when left exposed to air, it forms Rhenium-oxide crystals that grow over the course of weeks until they cover the MoRe fully. This is unwanted, as they form an insulating barrier that makes creating a good contact between the sample and measurement devices difficult. The oxide crystals are visible in dark-field microscope images, Fig. 3.11a. The oxidation is slowed by keeping the sample in a Nitrogen-box. Given enough time, the oxidation will happen and be visible to the naked eye, as MoRe will lose its shine. The oxidized crystals can be etched away by dipping the sample in MF321 for 30 seconds, and cleaning with H₂O and IPA baths afterwards. This also etches the MoRe layer somewhat, so it can not be repeated indefinitely.

The oxidation of MoRe is an issue that would prohibit the use of these samples as amplifiers for longer periods of time. Another material is necessary for the superconducting structures, both Aluminium and NbTiN would be suitable options. However, changing to these material would change process parameters and require re-tuning the processes used. Because of that, the choice was made to keep MoRe as the base layer material.

Aluminium bubble issue

There appears to be another issue present on the Aluminium contact pads. That is, there are bright spots visible in a dark-field optical microscope image, Fig. 3.12a, similar to the MoRe oxidation described earlier. Only, Al₂O₃ is a very well known-oxide that forms a thin, smooth layer and does not grow the point-like crystals that MoRe grows.

Instead, we propose that the issue is similar to the one described in [62], where bubbles were shown in

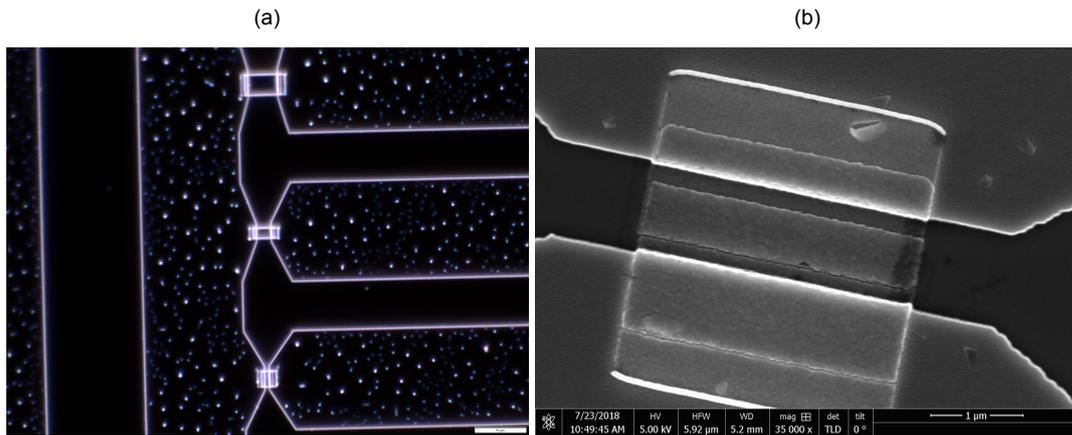


Figure 3.11: a) SEM image of a SQUID. On top of the MoRe electrodes (top and bottom), faint light spots are visible. These are suspected to be the Rhenium-oxide crystals. b) The bottom layer of the junction causes a shadow that breaks the connection of the top junction layer. Some lighter spots are also visible on the MoRe electrode, and another contamination is particularly visible near the top of the Aluminium junction. This could be a Rhenium oxide crystal, but it is also possible that this is a different type of contamination from the evaporation step.

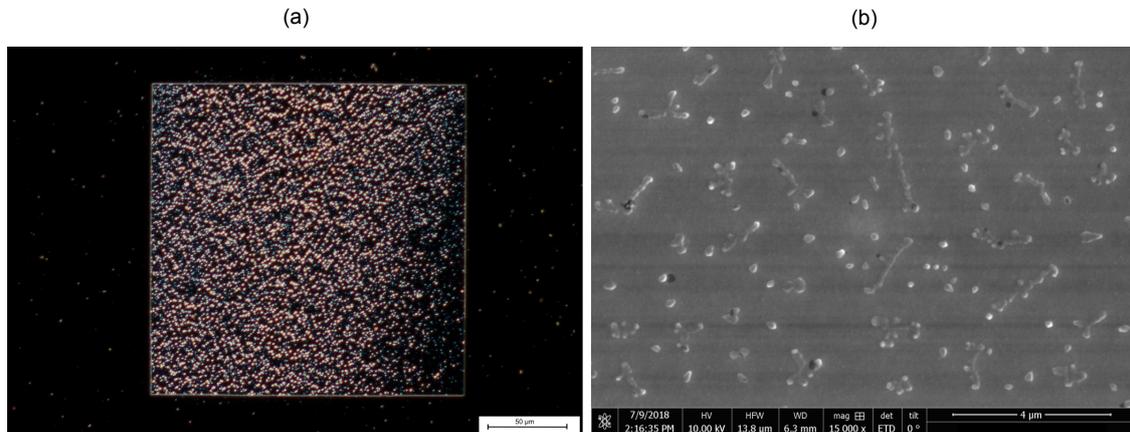


Figure 3.12: a) Bright spots on one of the Aluminium contact pads under dark-field optical microscope. b) SEM image of those same spots on the Aluminium contact pad. They look like (partially) popped bubbles. This could be related to an issue with the etch process, see main text for discussion.

Aluminium films on a Silicon substrate. Though no definitive explanation is offered, it is thought that the bubbles are formed by a contamination on the surface between the Aluminium and the substrate that interacts with one or more of the subsequent steps in fabrication. In this interaction, gas is formed and trapped below the Aluminium. If little gas is present, a bubble forms in the Aluminium, which bursts if more gas is present.

In our case, the Aluminium is not deposited directly on top of Si, as in [62], but on top of a MoRe layer. However, if a contamination is also present on our MoRe as it is on the Si, the effects could be similar, and it would explain the formation of bubbles on the surface of our Al films. Unfortunately, the presence of contaminants on the surface of the MoRe was not checked, and the steps taken in [62] to prevent the formation of such bubbles were not attempted. The issue appears to be present only for large area films, such as our contact pads, as it does not appear to be present in our junctions. Therefore, it is not likely to impact the functioning of the junctions significantly.

Lift-off

The lift-off process initially used consisted of bathing the sample chip in NMP (N-Methyl-2-pyrrolidone)

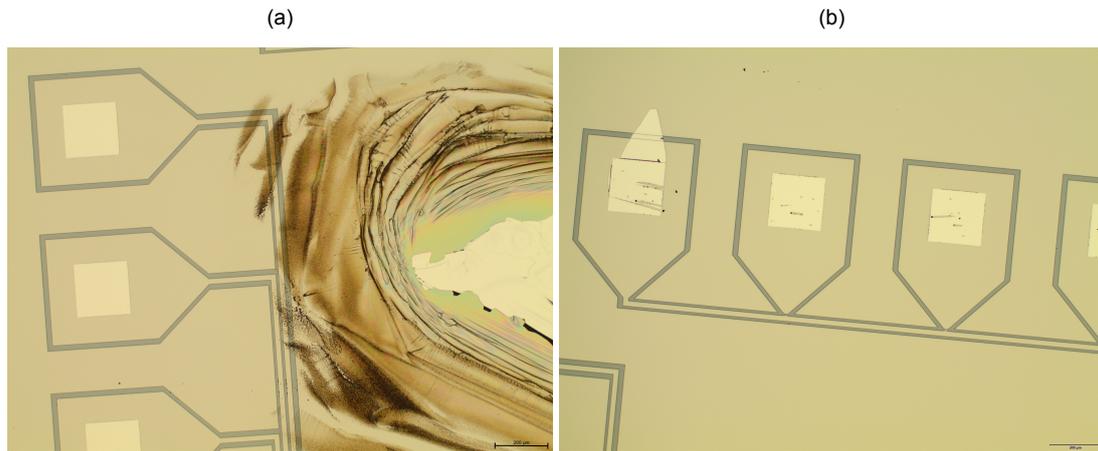


Figure 3.13: Lift-off issues. a) shows a large Aluminium flake that failed to lift-off, with some unidentified brown dirt surrounding it. b) A small flake of Aluminium left connected to the contact pad, where the lift-off locally failed. The scratches on the pads are from probing needles, and unrelated to the lift-off issue.

at 80° for 20 minutes with a stirring magnet. Occasionally, the metal did not lift off until a subsequent bath in Acetone. Even then, large parts of the metal occasionally remained on the chip, such as shown in Fig. 3.13a.

Initially, we thought the problem was the lack of flows in the solution to draw the metal away from the chip when the resist was dissolved. As a solution, a pipette was used while the sample was still in NMP to rinse off the metal. This solved most of the issue, but some smaller flakes were still present after the pipetting, such as shown in Fig. 3.13b. To prevent this, the beaker containing the sample and the 80° NMP was put in an ultrasonic cleaner before and after pipetting.

Conclusion of fabrication issues

We have discovered and solved several issues related to the fabrication of a Josephson parametric amplifier. There were two issues related to the etching step in our fabrication process. Regarding the etch sensitivity, we can conclude that we must closely monitor the etching as it happens. Measuring the reflectance of the sample with a laser from outside the etching chamber is an effective way to do this. We must also monitor the etching parameters (gas flow and pressure, and power) carefully to prevent the electrode side-walls from rounding.

There were also issues regarding the evaporation of Aluminium for the Josephson junctions. The evaporation angle, and thickness of the evaporated layer are important parameters to obtain junctions, so they must be monitored carefully. We could have searched for an optimum in these parameters, but we were limited by access to the evaporator. There is also a contribution common to the mentioned issues from the fact that the Aluminium can move around after its deposition. We suspect we can limit this effect by cooling the sample down during evaporation.

There was an issue regarding the oxidation of MoRe, which is a fundamental issue. This can only partially be avoided by keeping the sample in an atmosphere with little Oxygen between the fabrication step. A way around this issue would be to switch to a different material for the structures (contact pads, waveguides, capacitor, etc.) on our chip. For example, NbTiN (Niobium Titanium Nitride) could be used, as we have access to it in our cleanroom facilities. However, this would mean changing several other steps of the fabrication process, for example the etchant gas, power and pressure would need to be fine-tuned for this new material.

Regarding the Aluminium bubble issue, we have not confirmed that the origin of the issue is the same as that of [62]. They achieve a solution to their problem, but we have not tested to see if it applies in our situation as well. If it does, it involves an extra cleaning step before evaporating the Aluminium, which is

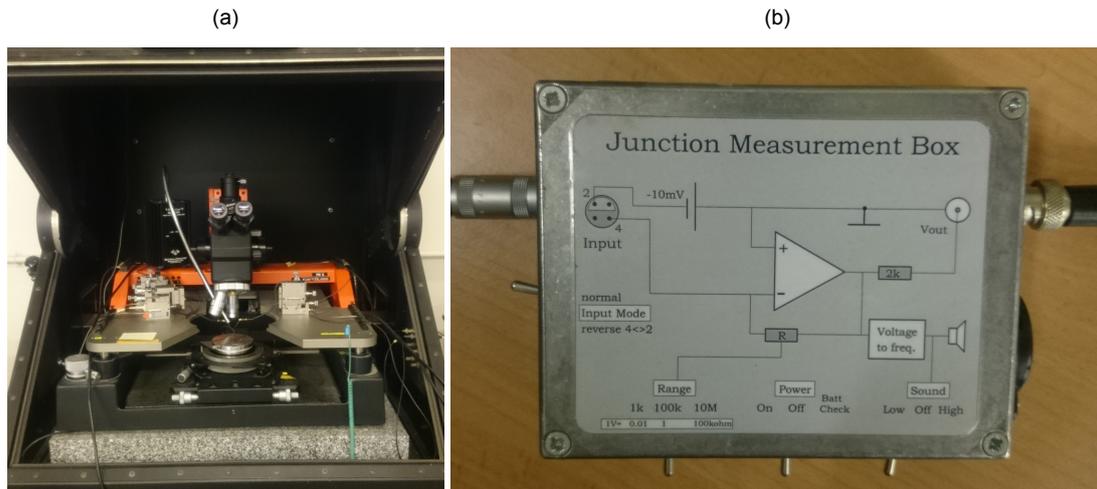


Figure 3.14: a) Probe station used for room temperature DC resistance measurements. At the top is the hood cover, which can be closed to block off light. In the center is the observation microscope, to position the probe needles accurately. Below the microscope is the sample table, where the sample can be held in place by a vacuum pump. To the sides of the microscope are the positioning units, for fine control of the probe needles. b) Junction measurement box, which converts the measured junction resistance into a DC voltage. The circuit diagram is shown on the box itself, with the connection to the junction at the left and the voltage output going to the right. At the bottom is a set of switches, to control the box settings.

not currently possible in the evaporator used. It is possible that other, similar, cleaning processes could solve the issues, but these were not attempted.

There are several avenues open where the fabrication could benefit from tuning of the fabrication process. Alternatively, a shift to a different superconducting material for circuit elements on the chip (NbTiN) might be a better solution in the long term, as we are limited by MoRe oxidation in the practical application of our amplifiers.

3.4 Measurement set-up and protocol

This section details the machines and set-ups used to perform the measurements shown in Chapter 4. In particular, the room-temperature set-up and the Helium cryostat (fridge) will be described.

3.4.1 Room temperature Junction measurements

To measure and characterize the resistance of the Josephson junctions at room temperature, a probe station was used (Fig. 3.14a). The probe needles were connected to a custom-built junction measurement box (Fig. 3.14b). Using a normal multimeter would put too much current through the junction and could damage it, so this measurement box was required. It is custom built to determine the resistance of a structure (Josephson junction) while delivering minimal power to that structure. It measures the resistance of the junction, and outputs a corresponding voltage. This can be converted to a resistance, dependent on the measurement setting.

To perform a room temperature resistance measurement, the sample was put on the sample plate of the probe station, fixed by a vacuum system and brought into the focus of the microscope. The two probe needles were connected to the contact pads of the sample using the probe station's manipulators. Sufficient force was applied to ensure good contact between the needles and the sample. The probe station light source was turned off to decrease the Si photo-conductance. A current was sent through one of the probe needles, it passes through one of the contact pads, through the junction, and through the other contact pad back via the probe needle towards the measurement box.

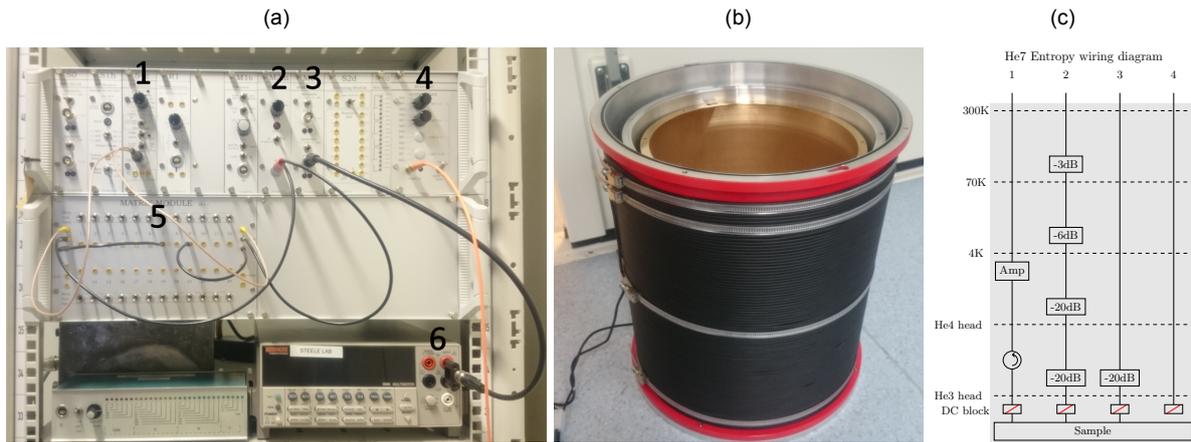


Figure 3.15: a) Setup for DC measurements. 1) I/V source module, 2) Voltmeter module, 3) Output module, 4) Digital-analog converter (DAC) for computer control of the modules 1-3, 5) Matrix module, 6) Digital multimeter. b) The three shielding cans of the Entropy system. The (red) outer can is the only vacuum seal, and around the outside of this can is the magnet (black). The inner two shields (gray and copper-coloured) function mainly as thermal and radiation shields. c) Wiring schematic for the Entropy He3/He4 sorption fridge. The He3 head is the coldest stage at 250 mK.

3.4.2 Fridge

The measurements detailed in this thesis were performed in a He3/He4 sorption fridge (manufactured by Entropy Cryogenics, and therefore referred to as the Entropy) at approximately 250 mK. As is common in these type of fridges, there are multiple stages with decreasing temperatures as one goes further downwards, with thermal shields separating the stages and a vacuum can sealing the outside of the fridge. These shields are shown in Fig. 3.15b, where the outermost (red) shield is a vacuum can, thermal and radiation shield, and the innermost (gray and copper-coloured) cans serve only as thermal and radiation shields. The outermost can also has a magnet consisting of a (black) hand-wrapped wire around the can [63].

Inside the fridge are two sets of cables used for the measurements in this thesis. One is a set of 24 individually addressable wires used for DC measurements, and the other is a set of four cables suitable for transmitting signals in the GHz regime. A schematic of these wires is shown in Fig. 3.15c, where the attenuators and amplifiers on each individual line are shown (detailed in the following sections).

3.4.3 Measurements

Single current-voltage (I-V) curve

We are interested in the relation between the current through and voltage over our structures, as we can use this to determine whether we have a Josephson junction. We put a current through our structure, and measure the voltage. The resultant curve, if plotted, is called an I-V curve. To obtain the I-V curves of junctions and SQUIDs fabricated for this project, a computer-controlled DC 4-point measurement set-up was used, shown in Fig. 3.15a. An I/V-source was used to send a current through the matrix module. The matrix module is a board of switches that allows us to control through which of the 24 DC wires we send a current. A voltmeter was used for read-out. The voltmeter output was sent to a computer-controlled Keithley 2000 digital multimeter.

Inside the fridge, the sample was connected through a 24-pin connector to the 24 DC wires combined into one cable. This cable terminates in the matrix module where the individual wires could be connected to the rest of the measurement set-up.

Magnetic field sweep

A magnetic field was generated using a wire wrapped around the fridge casing, through which a current

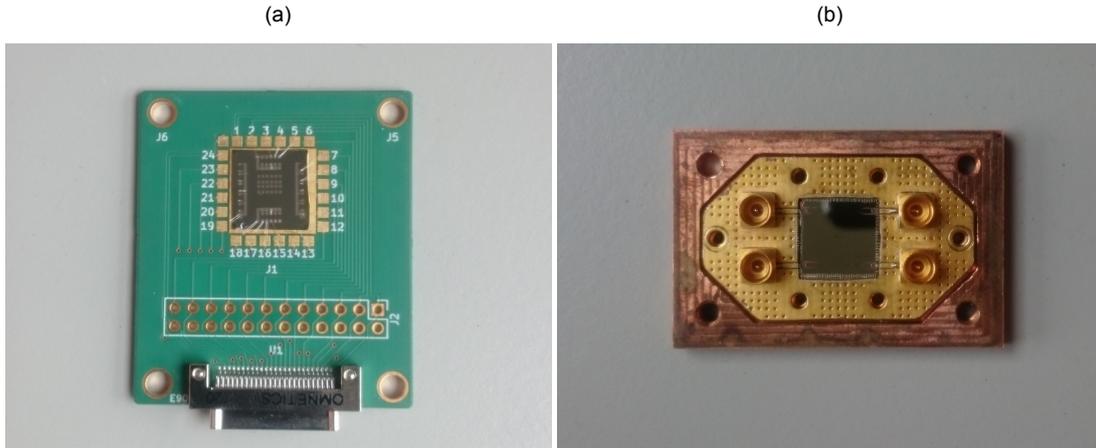


Figure 3.16: PCBs used in this thesis. a) PCB for DC measurements, with 24 ports symmetrically spaced around the inset for the $10 \times 10 \text{ mm}^2$ chip. b) PCB for RF measurements, with four ports

was sent. The current was generated using a Keysight B2901A precision source/measure unit, and controlled from the measurement computer. The maximum current through the wires was 2.6 A, which corresponds to a magnetic field of 0.65 mT.

RF measurements

To operate a JPA, microwave-frequency signals are needed, which require different cabling than DC measurements. The Entropy system has four connections suitable for GHz signals between the outside world and the cold plate, shown schematically in Fig. 3.15c. Line 1 is the signal output line, featuring a High Electron Mobility Transistor (HEMT) amplifier, LNF-LNC4_8C s/n 506Z, with 39 dB gain [64] at the 4K stage. Line 1 also contains a circulator between the He3 and He4 heads, to prevent thermal noise from the HEMT amplifier from reaching the sample via this line. Line 2 is the signal input line, with a total of 49 dB of attenuation. Line 3 is the pump line, with 20 dB of attenuation. Line 4 is unused. All lines also contain a DC block at the He3 head, which prevents a DC (or low frequency) signal from reaching the sample. Outside the fridge, at room temperature, we have connected an additional amplifier to line 1, providing an additional 32 dB of gain [65].

3.4.4 Printed circuit boards

To connect our sample chips to the fridge cabling and the room-temperature measurement devices outside it, two Printed Circuit Boards (PCBs) were used. Both have an inset to connect $10 \times 10 \text{ mm}^2$ chips, and both have internal wires to connect the chips to the fridge cabling. The DC PCB (Fig. 3.16a) has 24 ports symmetrically spaced around the chip inset, and featured a connector for a cable to a matrix board outside the fridge. All ports could be connected to any other via this matrix board, or be set to open or ground. The PCB for GHz measurements (Fig. 3.16b) has four ports in two sets of two on opposing sides of the inset. These can be connected to via SMA-connectors (as shown in the figure), which are common connectors suitable for microwave-frequency signals.

To ensure the chip is connected sufficiently to the PCB, it is glued at the back, both for stability and for thermalization of the chip. The electrical connections are realized through wire-bonds from the contact pads on top of the chip to the pads on the PCB. Some of these wire-bonds can be seen in Fig. 3.16a, they are the small lines stretching radially outward from the center of the chip.

3.5 Simulation of a general parametric amplifier

To gain insight in how a parametric amplifier works, we have simulated a general parametric amplifier. We could use the simulation results to qualitatively study what behaviour to expect from a parametric

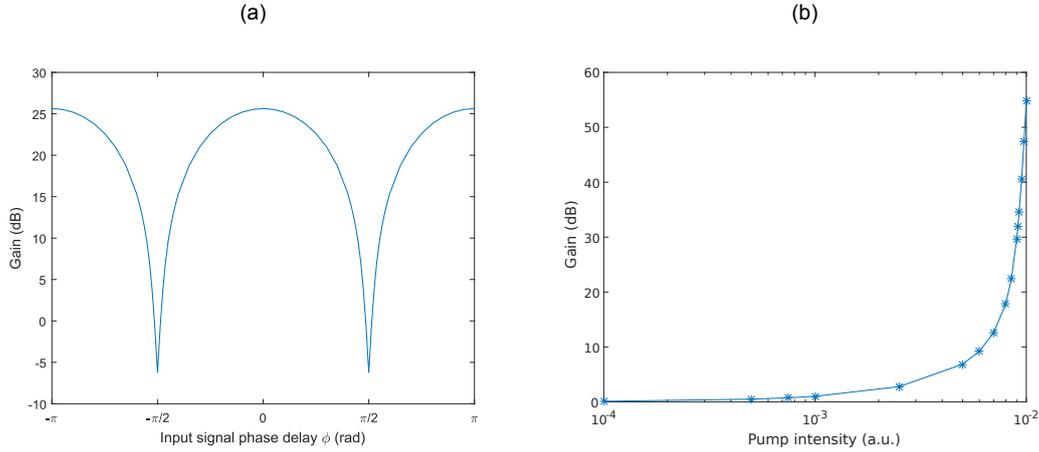


Figure 3.17: a) Gain dependent on input signal phase delay. As expected, the gain is a periodic function ϕ with period π . The maximum de-amplification is 6 dB, related to the fact that this simulation only looks at the internal field of the amplifier [43]. b) Gain dependent on pump strength. As expected, the gain increases if the pump amplitude increases. It appears to diverge above a certain pump strength, but is in fact limited due to the non-linear Duffing term included in our simulation.

amplifier. We have simulated a parametric amplifier with an additional Duffing non-linearity. Earlier theoretical work [66] has shown that a SQUID embedded in a transmission line gains a Duffing oscillator term. We expect this non-linearity to limit the oscillation amplitude.

We have simulated the parametric amplifier using MATLAB. The differential equation used to model the system is

$$\ddot{x} + \kappa \dot{x} + \omega^2 x + \beta x^3 = F(t) \quad (3.2)$$

Where κ is the loss rate, $\omega^2 = \omega_0^2 (1 + \alpha \sin(2\omega t))$ includes the parametric pump as in Eq. (2.41), which shows we are using the 3-wave process. We use β to describe the strength of the Duffing non-linear term and $F(t) = F_0 \sin(\omega t)$ is the driving term (i.e. the signal). This differential equation was numerically solved in MATLAB using the Euler method with 10.000 points per oscillation (natural frequency). To simplify calculations, we have set $\omega_0 = 1$, and taken $\kappa = 0.005$ as the default value (corresponding to a Q-factor of 200).

Amplifier gain versus input signal phase delay

When the input signal is resonant with the natural frequency of the oscillator, the amplifier is phase sensitive. We expect the gain to have a maximum when the phase difference between the signal and the pump is zero², and have a minimum when they differ by $\pm \frac{\pi}{2}$ [42, 44]. We have included the input signal phase delay ϕ as

$$F(t) = F_0 \sin(\omega t + \phi) \quad (3.3)$$

into Eq. (3.2).

The gain is plotted versus the input signal phase delay in Fig. 3.17a, which nicely matches expectations. We see a periodic function with maxima of gain above 24 dB, and minima corresponding to 6dB of signal squeezing [42]³. The periodicity of the gain with ϕ is π , which is also as expected.

Amplifier gain versus pump strength

The gain of a parametric amplifier is controlled partially by the amplitude of the pump. For our amplifier model, we study the relation between these two. For various pump amplitudes, we have simulated our oscillator and calculated the gain. The resulting curve is plotted in Fig. 3.17b.

²This is dependent on whether we choose a sine or a cosine as the pump/input signals.

³The squeezing in this model is limited to 6dB, because we only look at the oscillator amplitude. The output amplitude can be squeezed further, due to interference effects [43].

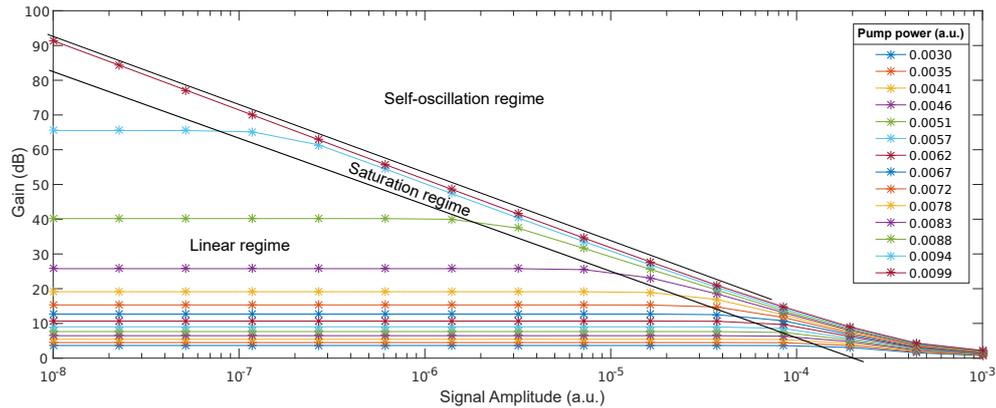


Figure 3.18: Amplifier gain as a function of signal strength for various pump intensities. There are three main regimes: The linear gain regime (the gain is independent of the signal strength), the saturation regime, (the gain decreases with increasing signal strength), and the self-oscillation regime (the oscillator resonates independently of signal strength).

We see that for low pump amplitude, the gain is negligible. When we increase the pump amplitude, the gain increases exponentially. In an ideal amplifier, this would continue without limit, and the gain would diverge [43]. However, in our amplifier, the gain does not diverge. Instead, it is limited by the non-linear Duffing term we have included in our simulation. This term limits the maximum amplitude of the movement of our oscillator, and that is what limits the gain. The gain can be increased by decreasing the input signal strength, and only the total amplitude of the oscillator movement is limited. By decreasing the input signal strength, we can therefore amplify it more and have higher gain.

Amplifier gain versus signal strength

We are interested in the expected gain depending on the input signal strength. The gain also depends on the pump strength, so we sweep over various pump strengths too. The results are plotted in Fig. 3.18. Clearly, three regimes can be recognized [43].

- **Linear gain regime.** Here, the gain is independent of the signal strength. For higher pump powers, the gain increases.
- **Saturation regime.** Here, the gain decreases for increasing signal strength. This is likely due to the oscillation amplitude being limited by the non-linear term. For higher pump powers, the saturation regime sets in at lower signal strengths.
- **Self-oscillation regime.** In this regime, the amplifier oscillates even with zero signal strength. For a linear oscillator, the amplitude of the oscillations would diverge. In this model, the amplitude is limited by the non-linear term.

With our parametric amplifier, or any amplifier in general, we want to be in the linear gain regime. In this regime, we can reliably calculate the strength of the original input signal from the output of the amplifier. For the saturation regime and the self-oscillation regime, this is not possible. For the self-oscillation regime, we could not even reliably say if there was any signal input for a given output.

We see that for any given pump amplitude, we are in the linear regime for low signal power. If we increase the signal power, we will enter the saturation regime. Here, the gain will decrease towards zero. This is related to the maximum oscillation amplitude of our amplifier. The Duffing non-linear term limits this amplitude, so we cannot amplify a sufficiently strong input signal by the same factor as a weaker input signal. This causes the gain to decrease.

Maximum oscillation amplitude versus Duffing non-linearity parameter

Here, support is given to the point made in the previous paragraph, that the maximum amplitude of the

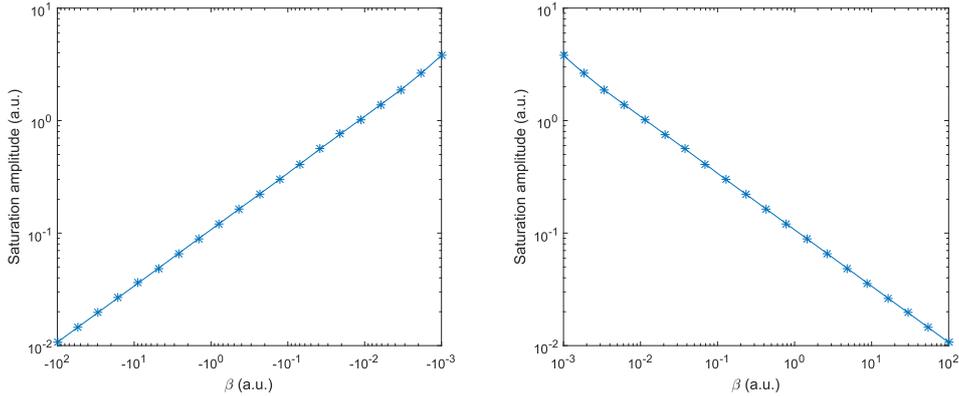


Figure 3.19: Self-oscillation amplitude for various non-linearity parameters. The amplitude goes to infinity as β tends to zero, and it behaves symmetric around the origin.

oscillator motion is limited by the Duffing non-linearity included in the model. We simulated the oscillator for identical pump and signal strengths, and varied the Duffing non-linearity parameter β . We have simulated the oscillator for both positive and negative values of β .

The results are plotted in Fig. 3.19. We can see from the symmetry of the plots that the sign of our Duffing term does not matter for the maximum amplitude of the oscillation. The determining factor for this amplitude is the magnitude of the non-linear term, we see that for larger (absolute) values of β , the maximum oscillation amplitude decreases. For $\beta = 0$, the maximum oscillation amplitude tends to infinity. This is expected, as it is the behaviour of an ideal linear amplifier.

Amplifier gain versus Duffing non-linearity parameter

We have shown the behaviour of the maximum oscillation amplitude for various values of the Duffing term. Now, we study the behaviour the amplifier gain if we were to vary the Duffing parameter. In Fig. 3.20, we simulate and plot the amplifier gain as a function of signal strength for various values of the non-linearity parameter β in Eq. (3.2). The figure shows that for higher non-linearity, the saturation regime sets in earlier and the maximum input signal strength the amplifier can amplify with a constant gain is decreased.

From these simulations of the behaviour of our parametric amplifier for various values of the Duffing non-linearity parameter β , we can conclude that a small β is preferable. This would allow us to reach a higher gain, by allowing us to increase the pump strength further without entering the saturation or self-oscillation regime. Another benefit of having a small β is the possibility to amplify stronger input signals linearly (again, not entering the saturation or self-oscillation regime).

Amplifier bandwidth

An important figure of merit is the amplifier bandwidth. There are several different properties of a JPA parametric amplifier that could be called a bandwidth. The first is called the instantaneous bandwidth, this is the frequency range of signals that can be (simultaneously) amplified. The second is the tunable bandwidth, this is the frequency range over which we can tune the resonance frequency of our SQUID cavity. Typically, the instantaneous bandwidth is the measure used to qualify bandwidth of an amplifier, and this is also the quantity which we will simulate here. The instantaneous bandwidth of a typical JPA is several MHz, as is shown in Table 2.1 for several recent papers, while the tunable bandwidth can be on the order of GHz.

We have simulated our parametric amplifier for various detuning (the frequency shift of the input signal with respect to half the pump), for various pump powers. This is shown in Fig. 3.21. We can see two different qualitative behaviours of the curves in the plot. For low pump strengths, the gain is highest for zero detuning, and decreases if the detuning increases. However, for high enough pump strengths, we reach a plateau in the gain for low detuning, and then a drop.

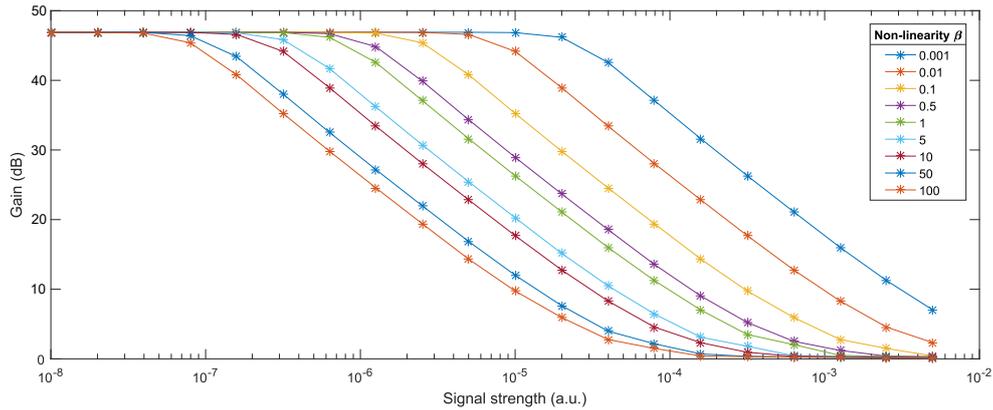


Figure 3.20: Simulation of amplifier gain as a function of signal strength for various values of the Duffing non-linearity parameter β . For higher non-linearity, the saturation regime sets in earlier.

The plateau we see is related to the maximum oscillation amplitude, and the value at which the gain plateaus is governed by the Duffing non-linearity parameter in combination with the input signal strength. The plateau is slanted, with the maximum gain being lower for negative detuning. This appears to be related to the sign of the Duffing non-linearity, as we can change the behaviour to have a maximum gain higher for negative detuning by flipping the sign of the Duffing non-linearity parameter β in our simulation.

The plateau of maximum gain suddenly drops off for large detuning and high pump powers. This appears to be related to a limitation of the simulation. We simulate the oscillator only for a limited (set) time. For some particular parameter combinations, the amplitude does not stabilize in time. For example, for a very small detuning we get a beating of our oscillation amplitude. This can affect the calculation of our gain, and lead to erroneously placed points in the figure. We have attempted to remedy this by extending the time which we simulate the oscillator, but were limited by the time involved in these calculations.

The bandwidth of the oscillator can also be seen from the results of Fig. 3.21. We take the bandwidth to be the range of frequencies where the amplifier gain is within 6 dB of its maximum. We can observe a trend from the curves in the figure. For low pump power, the gain is low and the amplifier bandwidth is broad. As we increase the pump power, the bandwidth decreases, as expected if we assume the product of gain and bandwidth to be constant. But if we increase the pump power further, we break away from this trend, as the gain increases (marginally), and the amplifier bandwidth increases (the peaks broaden in the figure). This is related to the maximum oscillation amplitude, which constricts the gain of the amplifier.

Conclusion of the simulations

We have simulated an ideal parametric amplifier with an additional Duffing non-linearity, with a simple and straightforward method. Qualitatively, the simulation show amplifier behaviour as expected. The gain is dependent on the phase delay between the pump and the input signal, and increases with pump power until it saturates the amplifier. The saturation is governed by the Duffing non-linear term, as it limits the maximum oscillation amplitude. For low pump power, we see a constant gain-bandwidth product. If we increase the pump power so that we can enter the saturation regime, the gain is limited and the amplifier-gain product increases.

There are many avenues left open for a quantitative analysis of the simulations we have performed. There have been significant works by others to simulate parametric amplifiers with more accuracy, and more realistic parameters, for example [46, 45], and we have not attempted to quantitatively link the parameters of our simulation to design parameters of our parametric amplifiers.

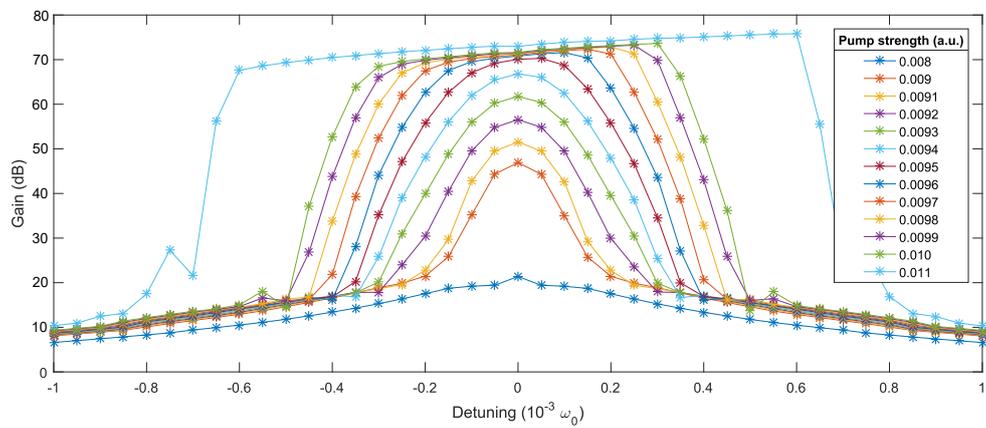


Figure 3.21: Parametric amplifier gain plotted for different detuning (input signal detuned from half the pump frequency), for various pump strengths. The x-axis shows the frequency difference between the input signal and half the pump frequency in units of 10^{-3} the resonance frequency of the oscillator.

4

Results

4.1 Characterization of Josephson junctions

This section details the measurements done to characterize the fabricated Josephson junctions, both at room temperature and at 250 mK.

4.1.1 Room-temperature resistance

An attempt was made to measure the room temperature resistance of each of the junctions. These measurements were performed using a probe station and a custom junction measurement box designed to probe using low power to prevent damage to the junction. Two probe needles were put into contact with the Aluminium contact pads shown in Fig. 4.1a, as the geometry is a 2-terminal one, and the resistance was measured between them.

The results show no clear correlation between the designed junction area and the measured resistance. This is most likely due to the fact that the metal leads are $0.1 \times 20 \mu\text{m}^2$ MoRe. The resistivity of similarly sputtered MoRe thin films is approximately $800 \text{ n}\Omega\text{m}$ [67], or 8.0Ω per square.

We can calculate the expected normal-state (i.e. non-superconducting, such as at room temperature) resistance of the Josephson junctions via the Ambegaokar-Baratoff relation, Eq. (2.29),

$$R_n = \frac{\pi \Delta(T)}{2eI_c} \tanh\left(\frac{\Delta(T)}{2k_B T}\right), \quad (4.1)$$

where we can use the earlier provided estimation of the critical current density, $0.7 \mu\text{A}/\mu\text{m}^2$. Using the smallest ($1 \times 1 \mu\text{m}^2$) and largest ($6 \times 4 \mu\text{m}^2$) junction sizes as limits, we expect critical currents between $0.7 \mu\text{A}$ and $16.8 \mu\text{A}$. If we insert the energy gap for Aluminium, $\Delta(T) = 180 \mu\text{eV}$, we find expected normal-state resistances of 404 and 15.2Ω for the smallest and largest junctions respectively.

We expect to be able to measure the room-temperature resistance of the junctions, as their normal state resistance is comparable to the expected resistance contribution from the MoRe leads. However, no clear trend is observed when plotting the measured resistance versus the designed junction area (which controls the I_c). There is a trend observed of the measured resistance with the designed electric path length, as plotted in Fig. 4.1b. We have taken a linear (Ohmic) fit, in an attempt to subtract the lead resistance and obtain the junction resistance. However, there is a spread of data points around the fit in Fig. 4.1b that is wider for longer path lengths. Subtracting the (fitted) resistance from the data points does not result in an observable trend of resistance scaling with junction area.

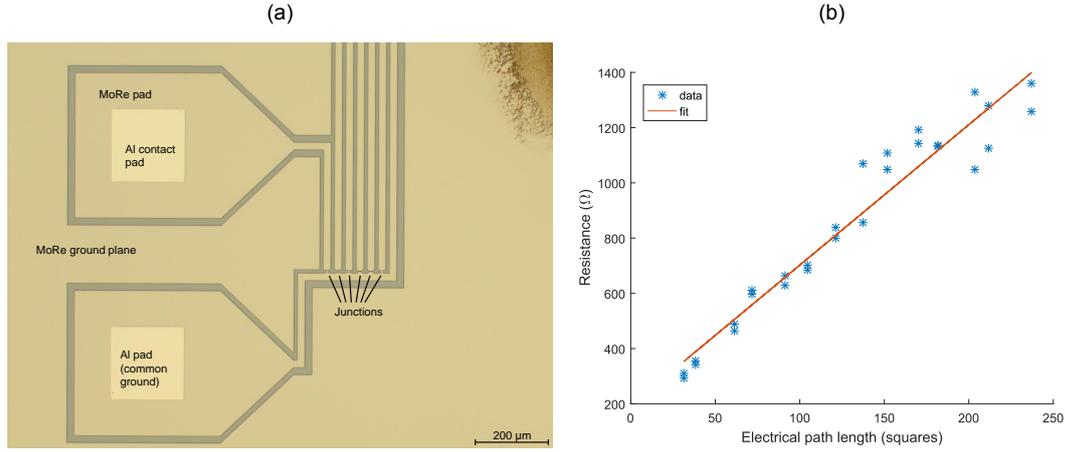


Figure 4.1: **Room-temperature resistance measurements.** a) Optical microscope image of one of the samples. Two Aluminium contact pads on top of MoRe pads are visible, connected through a 20 μm wide MoRe lines to a junction. The lines connecting other junction lead to identical contact pads outside the image. b) Measured resistance versus (designed) electrical path length. The fit is linear, representing Ohmic behaviour, and does not go through the origin due to one or more contact resistance(s) in the system.

Sources of this additional resistance could be the contact between the probe needle and the Aluminium contact pad, or the contact between the Aluminium pad and the MoRe. Also the contact between the Aluminium of the junction and the MoRe lead could lead to the measured spread in the data points of Fig. 4.1b. Potentially, larger-scale variations in, for example (see Sec. 3.3), the thickness of the MoRe over the chip, or small impurities or contaminations from fabrication can cause variations in the resistance.

The photo-conductance of the Si substrate could also affect the resistance measurements, but for the measurements in Fig. 4.1b the microscope light was turned off. The magnitude of the photo-conductance can be estimated from a sample where no junction was deposited between the MoRe electrodes. Over a distance of several microns, the resistance of the Si is 20-50 kΩ with the microscope lamp, and MΩ without the microscope lamp. Thus, we do not expect the photo-conductance to influence these measurements significantly.

The slope of the fit can be used to estimate the resistivity of the MoRe. Based on the designed geometry ($0.1 \times 20 \mu\text{m}^2$ leads), the resistivity can be obtained for different samples. From the measured resistance R , and the designed length L of the path and cross-section A , we can calculate the resistivity according to

$$\rho = R \frac{A}{L}. \quad (4.2)$$

The results are shown in Table 4.1, where both the offset resistance and the MoRe resistivity are tabulated. The resistivity is less than the expected value of approximately 800 Ω [67] measured for thin-films in our lab, though the uncertainty (based in the 95% confidence interval from the fits) is considerable. This could be explained if the MoRe layer thickness is different than the designed 100 nm. Based on profilometer measurements, the depth of the gaps between MoRe electrodes (where the Si substrate is exposed) is increased by a small factor, which suggests that the samples were over-etched slightly. This would explain a calculated resistivity that is higher than the literature value, while we have a resistivity that is higher than the literature value. We attribute the discrepancy between the calculated and expected value of the resistivity to differences in sputtering parameters.

#	Offset resistance (Ω)	MoRe resistivity ($n\Omega m$)
1	439 (± 322)	477 (± 183)
2	344 (± 486)	524 (± 272)
3	193 (± 52)	508 (± 52)

Table 4.1: Tabulated MoRe resistivity and offset resistance for various samples. Uncertainties give the 95% confidence interval.

4.1.2 Critical current at 250 mK

One of the samples from the previous section (# 3 from Table 4.1) was loaded into the Entropy fridge and cooled down. The I-V curves of the single junctions were taken as detailed in Sec. 3.4.3. We see two qualitatively different behaviours, and we will treat them separately in the following paragraphs.

Critical-current like behaviour

At first glance, three of the junctions show the behaviour expected from a Josephson junction, but a more detailed look shows that this is not the case. In Fig. 4.2a, we have plotted the I-V curves of these three junctions, denoted by the numbering of the DC lines they are connected to, that appear to have a jump in the measured voltage around 130 μA (denoted with a 1 in the figure). While we expect such a jump from the critical current of the Josephson junctions, its magnitude should reflect the designed junction size (13-14: 9.6 μm^2 , 13-16: 23.2 μm^2 , 13-18: 2.4 μm^2) based on a constant critical current density, which it clearly does not. The normal state resistance is approximately 30 Ω based on the slope of the I-V curve, which would correspond to a critical current of 10 μA for an Aluminium junction, based in Eq. (2.29). From this, relation, we can calculate the $I_c R_n$ product, which is a material-dependent constant. From the measured I_c and R_n , their product is 3.9 mV, where only 0.28 mV is expected for an Aluminium junction. One of the junctions also has a much larger jump in the I-V (denoted by a 2 in the figure) at 170 μA .

The voltage jump is a critical current of some structure on the sample, as the sample is superconducting at lower applied currents, and not superconducting at higher currents, where we see a normal-state resistance. All three junctions that show this behaviour are connected to one of the common ground contact pads, number 13. However, other junctions connected to this pad show qualitatively different behaviour, as in Fig. 4.2b. Therefore, we hesitate to ascribe this difference in behaviour to an issue with the contact pad, or the wire leading from the common ground pad to the junctions.

A potential origin of the voltage jump at 130 μA is that it is the critical current of a constriction in the Aluminium of the junction. However, the current through this constriction would have to generate a magnetic field higher than the critical field of Aluminium, 0.01 T. If we approximate the constriction as a wire,

$$B = \frac{\mu_0 I_c}{2\pi r} \quad (4.3)$$

it would have a radius r of 2.6 nm, which is much smaller than the the designed junction sizes, The designed junction dimensions are: The bottom layer 32 nm thick, the top layer 74 nm thick, with widths and lengths ranging between 1 and 6 μm .

The disparity between the designed size of the Aluminium junction and the apparent size of the constriction (if it were in Aluminium) could be caused by the Aluminium diffusing or partially de-wetting during evaporation. This would result in the evaporated Aluminium atoms moving around before settling on the substrate. This could result in constrictions in the Aluminium that limit the critical current to 130 μA . The voltage jump at 170 μA could be caused by a similar but slightly larger constriction.

However, this explanation does not take into account that we obtain a very similar critical current of 130 μA for three different junctions. It is unlikely that the Aluminium clumping together on the sample during the evaporation process would result in three different constrictions with identical size, and no constrictions apparent in the other junctions. It would appear that if the cause of the critical current is a constriction, it is somewhere else in the sample than the junctions.

We would suspect MoRe lead of the ground wire connected to pad 13, as this is the one connection that the three junctions share. However, measurements of other junctions also connected to this wire shown no voltage jump at 130 μA . These junctions did show the other type of I-V behaviour, as in Fig. 4.2b, and it is possible that the voltage jump did not appear clearly in the measurement that we have done for this range of applied current. We conclude that we have not found a conclusive origin of the critical-current like step in the voltage that appears at 130 μA for three of the junctions on our sample.

Shapiro step-like behaviour

Several of the tested structures show a qualitatively different type of I-V curves, as shown in Fig. 4.2b. Here, several steps are present in the I-V before the junction appears to enter a normal (non-superconducting) state and the I-V curve becomes linear, at currents approximately ten times smaller than of the I-Vs shown in Fig. 4.2a.

Our initial guess to the origin of the steps was that they are Shapiro steps [29]. If the junction is driven with an AC current at frequency ω , you would see steps in the I-V curve at

$$V = n \frac{\hbar\omega}{2e}, \quad (4.4)$$

where n is an integer. Based on the voltages of the steps, approximately 270 μV , the associated driving frequency would be $2\pi \cdot 130$ GHz. We send only a DC current through the junctions, and another source of current of this frequency would have to be located close to the junctions, as the cabling used for DC measurements is not a good conductor at those frequencies. It is unlikely that we would have such a high-frequency noise in our system at sufficient amplitude to generate these Shapiro steps.

Another possibility is that the steps are Fiske steps [68]. However, to excite Fiske steps, we would require a magnetic field to be applied (in-plane) to the junctions. Here, we do not apply such a field. We can rule out a stray field from other parts of the sample in a subsequent experiment (Fig. 4.3a), where we do apply a magnetic field with a strong in-plane component and vary the field strength. If the steps were Fiske steps, they would change in the size of the voltage jump with the applied magnetic field [68].

A variation of the Fiske steps, called zero-field steps [69], could be possible. We are in the right junction size-range (junction size comparable to the Josephson length), and structural variations in the critical current density of the junction could be possible. However, we can also rule out the zero-field steps with the same argumentation as above, as zero-field steps would also be affected by a magnetic field in the same way that Fiske steps are. We do not see the voltage jump change size with the applied magnetic field, so we conclude that the steps are not zero-field Fiske steps.

A third possibility for the origin of the steps is that they are photon-assisted tunnelling steps [70]. If we were to attribute two neighbouring steps to this, the frequency of radiation required would be $2\pi \cdot 65$ GHz. Though this is not as high a frequency as for the Shapiro steps, we measure only using DC and have no microwave-frequency lines connected to the sample. It is therefore unlikely that the steps are photon-assisted tunnelling steps.

A fourth option for the possible origin of these steps is that they are related to the thermoelectric effect. If we have connections between two different types of metal in the wiring of our fridge, we could have a thermocouple voltage. Depending on the temperature difference between the two connections, this voltage can be on the order of mV. We have not studied the wiring specifics of the fridge in sufficient detail to be able to reject this explanation. It is possible that the thermocouple voltage biases the junction, and makes it radiate at a certain frequency. Because we have several junctions close together, this radiation could affect the other junctions nearby, and create Shapiro steps.

We have searched for explanations for these steps, and have found a reasonable explanation in the thermal voltage from connections between different metals at different temperatures in the form of the

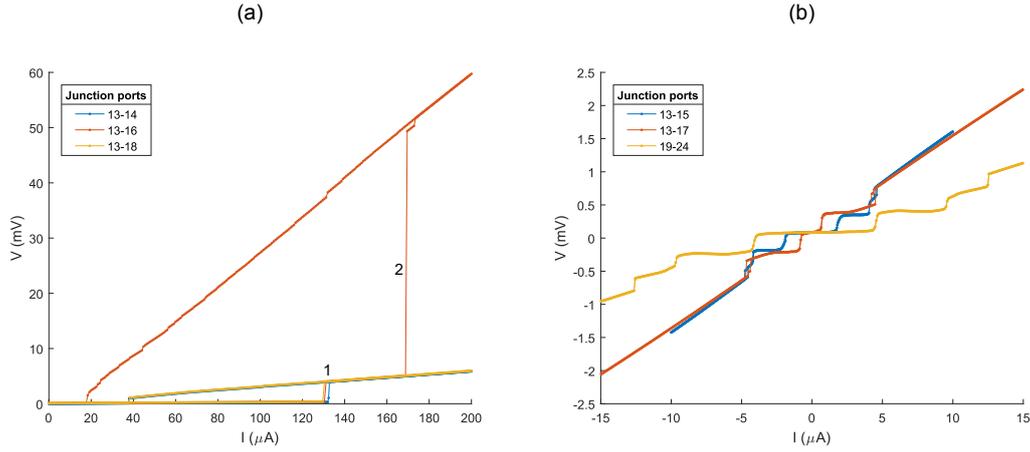


Figure 4.2: **I-V measurements of junctions and SQUIDs.** a) I-V curves from three different junctions. All junctions have a jump in the measured voltage at 130 μA , despite different junction sizes, and one junction has a jump at 170 μA . The numbers in the legend denote contact pads on either side of the various junctions. b) I-V characteristics for two single Josephson junctions (13-15 and 13-17) and a SQUID (19-24) showing steps of approximately 270 μV .

fridge cabling. However, we have not studied the fridge cabling in sufficient detail to be certain about this explanation.

4.1.3 SQUIDs in a magnetic field

To check whether the fabrication of the SQUIDs had succeeded, A magnetic field was applied to check for modulation of the critical currents. The I-V curves of the SQUIDs were taken while varying the applied magnetic field, the results are plotted in Fig. 4.3a and show dV/dI of one of the SQUIDs.

Fig. 4.3a shows four vertical lines, corresponding to steps in curve 19-24 of Fig. 4.2b, that are not modulated by the applied magnetic field, and two lines (also in Fig. 4.2b) that are modulated. If we attribute the modulating line to the critical current of the SQUID, we can extract the apparent size of the SQUID loop to check if it matches to the design. However, the addition of a bracket to the fridge set-up (which was added for ease of mounting) causes the sample to be mounted at a 90° angle with respect to the expected orientation. For this reason, the previously out-of-plane magnetic field becomes an in-plane magnetic field, as shown in Fig. 4.3b.

Despite this experimental error, we still see some modulation of one of the voltage steps with the magnetic field in Fig. 4.3a. The difference in applied (in-plane) magnetic field between the peaks of the modulation is 0.41 mT, and if one modulation period represents one flux quantum moving in or out of the SQUID loop we can determine the apparent loop area, $A = \frac{\Phi_0}{B}$, where Φ_0 is the flux quantum ($2.067 \cdot 10^{-15}$ Wb) and B the applied magnetic field. For this SQUID loop, the apparent area is $5.2 \mu\text{m}^2$ which is significantly smaller than the designed area of $32 \mu\text{m}^2$. If we attribute this difference to a rotation, and we can determine how much off from the perfectly in-plane field we appear to be. Here, we appear to have an angle difference of 9.4° . For similar B-field sweeps involving other SQUIDs on the sample, similar results are reached. These are tabulated in Table 4.2.

One immediately noticeable thing about the data tabulated in Table 4.2 and the fridge shown in Fig. 4.3b is the discrepancy between the apparent angle and visible rotation angle of the bracket. If the bracket, to which the sample is bolted flatly, was 9° out of plane, it should be visible by eye, and it is not. Thus, this is not the reason of the magnetic field appearing to have an out-of-plane component at our sample.

There are several things that could be (partially) responsible for the out-of-plane component of the magnetic field at our sample. Because the solenoid is around the outside of the fridge, and not infinitely long, and the sample is not exactly in the middle of the solenoid, we could get a radial component of

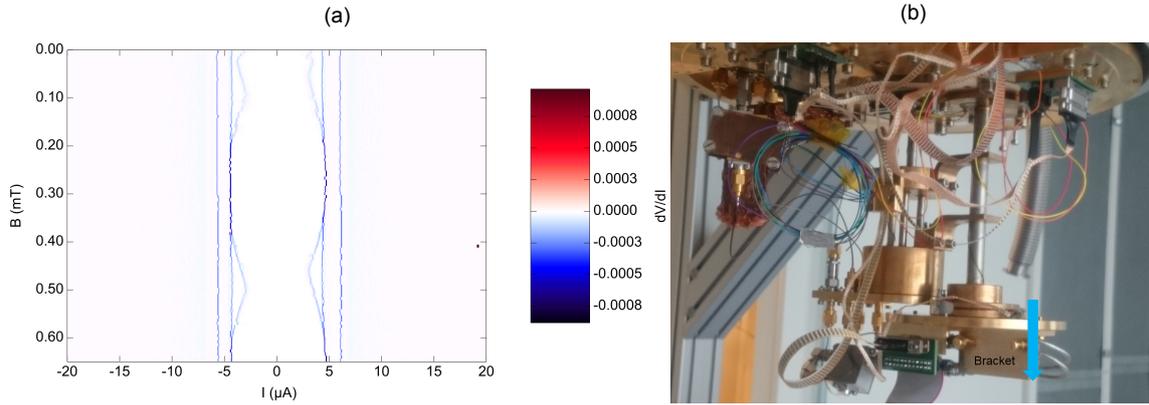


Figure 4.3: **SQUID behaviour in a magnetic field.** a) The I-V curve of a SQUID is modulated by the applied magnetic field in a periodic way. However, the magnetic field is applied in-plane with the SQUID loop, instead of the conventional out-of-plane manner, due to the addition of a new bracket (added for ease of mounting) in the fridge. b) Entropy fridge set-up with in blue the direction of the applied magnetic field, and the bracket shown where the sample is mounted so that the magnetic field is in-plane.

B-field for one period (mT)	Apparent area (μm^2)	Designed area (μm^2)	Apparent difference angle ($^\circ$)
0.41	5.2	32	9.4
>0.65	<3.2	24	<7.7
>0.65	<3.2	18	<10.2

Table 4.2: Tabulated B-field periodicity, area calculated from the periodicity and designed area, and the angle of orientation of the SQUID based on the discrepancy between the two.

the magnetic field. If our sample is 2/3 of the solenoid radius away from the center, and 2/3 of (half the) solenoid length away from the middle, the ratio of radial to vertical (parallel to the solenoid axis) fields is 0.07, which is enough to explain an angle difference of about 5° .

Other components localized near the sample (e.g. superconducting wires) or on the sample itself (superconducting ground plane covers most of the chip, as in Fig. 3.4a) could modify the local field, potentially enhancing out-of-plane components. To gain insight in how this affects our sample would require simulations and knowing the location of components in the fridge very well.

Another effect might be the calibration of the magnet itself. This was done in another work [63], but it was measured only inside the outer shield of the fridge, and not inside the innermost shield. If the magnetic field generated by the coil is different from the expected value, the apparent angle also changes.

A fourth effect that contributes to the apparent angle difference could be an overestimation of the effective SQUID-loop area. In general the area through which the flux penetrates is not strictly the area surrounded by the superconductor. Some flux could also penetrate the edges of the superconductor, which increases the area. For the designed area, a line through the middle of the junctions was taken, and a similar distance to the actual edge of the superconductor was taken for the ends of the SQUID. This is a very rough estimation, and could give an error of several μm^2 in the designed area in Table 4.2.

We have so far mentioned several effects that could cause the apparent difference angle to be approximately 10° . Without further study, we cannot be sure which effect contributes what to this figure, but this is beyond the scope of this work. For now, we will consider this apparent angle difference not a major unexplained issue.

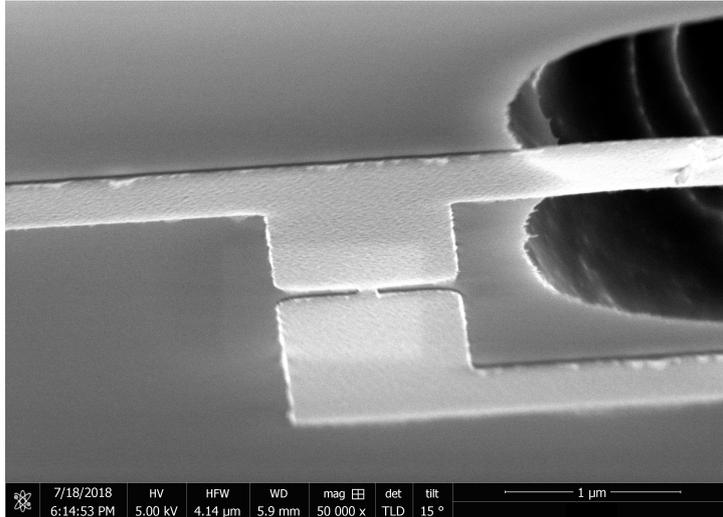


Figure 4.4: **Constriction Josephson junction.** SEM image of one of the constriction type Josephson junctions present on the chip used for the measurements of Sec. 4.2. The image shows part of the SQUID loop, where at the top right a part of the loop is released, and forms a mechanical element that is free to move.

4.2 Parametric Amplification

For parametric amplification, a different sample (with the design of Sec. 3.2.3) was used. Due to machine access issues, this sample was fabricated with a different type of Josephson junction, one with narrow constriction. One such junction is shown in Fig. 4.4. This figure also shows part of a free-hanging mechanical element (top left, where the substrate is etched away) as this design was created for a different project in the research group.

First, we characterize the SQUID cavities. We start by determining their resonance frequencies, and how this changes when we apply a magnetic field (i.e. we bias the SQUIDs). We also characterize their behaviour when we apply microwave signals through the signal and the flux pump line. After this, we will show that the SQUID cavities can be operated as a Josephson parametric amplifier. We will operate the JPA in 3-wave mode, and attempt to characterize the gain, bandwidth and noise temperature. Finally, we will operate the JPA in 4-wave mode, and attempt to characterize the gain and bandwidth as well.

4.2.1 SQUID cavities

To find the cavity resonance frequencies, we send in a microwave signal at a certain frequency, and measure how much of it is transmitted. By changing the frequency of the signal (i.e. performing a frequency sweep), we can detect the cavity resonance frequencies. We have done a frequency sweep between 3 and 8 GHz. The results are shown in Fig. 4.5, where all seven cavities are visible as sharp dips in the transmission. Several cavity parameters are tabulated in Table 4.3. Cavities 1-5 had a released mechanical part, and higher quality factors, while cavities 6 and 7 had no movable parts and a lower quality factor. Related to that, the external loss rate, κ_e , was higher for cavities 6 and 7. We have found the Q_e and κ_e from fitting the cavity. Ideally, we would have found the internal quality factor Q_i as well, but this is difficult to fit accurately if Q_i is very different than Q_e .

Noticeable in Fig. 4.5 is the background, which is not constant. Several dips are visible in the transmission, but they are too low in frequency to be higher modes of our cavities.

4.2.2 Biasing the SQUIDS

To drive the SQUID cavity parametrically, we modify the resonance frequency via the flux threading the SQUID loop. The change in resonance frequency per change in flux through the loop, $\frac{\partial\omega}{\partial\Phi_{\text{ext}}}$ is not a constant, there are certain ranges of flux where the resonance frequency is sensitive to a small change

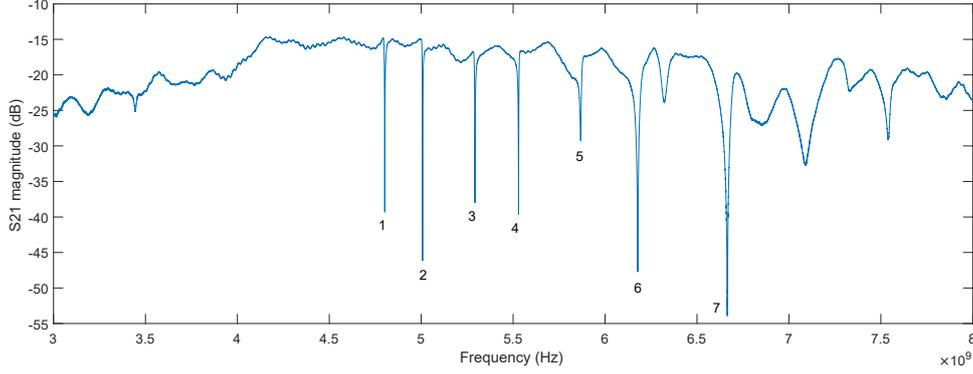


Figure 4.5: **SQUID cavity characterization.** Frequency sweep between 3 and 8 GHz showing all 7 cavities of the sample, denoted in the figure. The background is not flat, which is most likely due to the cable resonances. Based on the frequencies involved, none of the other dips is likely to be a higher mode of our cavities.

Cavity #	Center frequency (2π -GHz)	Q_e	κ_e (2π -MHz)
1 (released)	4.80	680	7.05
2 (released)	5.01	797	6.29
3 (released)	5.29	722	7.33
4 (released)	5.53	703	7.87
5 (released)	5.87	495	11.86
6	6.19	129	47.98
7	6.70	162	41.36

Table 4.3: Overview of SQUID cavity parameters

in flux. We can see this from the curve (flux arc) shown in Fig. 2.12c. To drive our parametric amplifier, we want the change in frequency per change in flux, $\frac{\partial\omega}{\partial\Phi_{\text{ext}}}$, to be as big as possible. To achieve this, we apply a constant, external magnetic field to 'bias' the SQUID to one of these points on the flux arc, where a small modulation of the flux invokes a large change in the resonance frequency of the cavity.

In Fig. 4.6a, we can see the curve of the resonance frequency with the applied field, which is the flux arc. If there is inductance present from the SQUID loop itself (instead of only from the SQUID junctions), the flux arc changes shape from the ideal case (Fig. 2.12c) and the different arcs can overlap. The SQUID can jump between different flux arcs if perturbed, which is visible in Fig. 4.6b. This can cause hysteretic behaviour, if the cavity jumps to a different flux arc due to flux noise, it will follow it if the magnetic field is reduced to below where the cavity jumped. It can be made to jump back by lowering the field further.

The biasing is automated via a Python script. The magnetic field is initialized at zero, and the resonance frequency of the cavity is measured. In steps of 0.5 mA, the current through the magnet is increased until the resonance frequency has reached 0.998 times the resonance frequency at zero field, which is denoted by the green dots in Fig. 4.6b for cavities 6 and 7 respectively. The current applied to the magnet was between 30 and 40 mA. The results of a typical SQUID-biasing procedure are shown in Fig. 4.6a. Starting from zero applied field at the top of the image, the applied current is increased in steps until the cavity resonance frequency has reached 0.998 times its zero-field value at the bottom of the image.

4.2.3 SQUID cavity power limits

If we put a large amount of power into our SQUID cavity, either by having a strong signal or a strong pump, the behaviour of the SQUID cavities can become different from the behaviour for small signal and pump powers. We can, for example, observe this from the behaviour dip in the S_{21} transmission ampli-

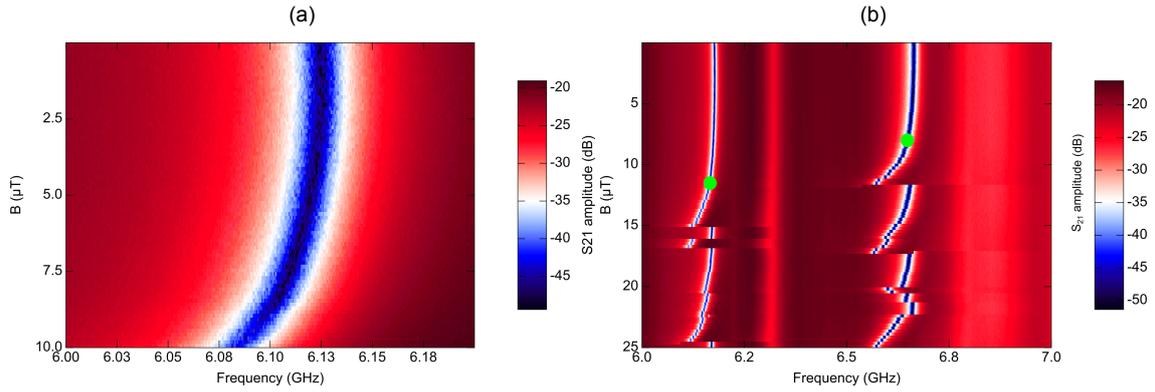


Figure 4.6: **SQUID cavity behaviour in magnetic field.** a) SQUID biasing. In steps, a larger current is applied to create a stronger magnetic field. When the resonance frequency of the SQUID reaches 0.998 times the resonance frequency at zero applied field, the biasing process is finished. b) Overview of cavity 6 and 7 for changing magnetic field. Several jumps are clearly visible for larger magnetic fields.

tude, which becomes asymmetric for high input signal power, as is visible near the bottom of Fig. 4.7a. Here, cavities 6 and 7 change their transmission behaviour for signal powers (generator output, not power-on-chip) of 9.4 dBm and 11.9 dBm respectively. That is, the dip in the transmission becomes smaller and behaves asymmetrically in frequency, as can be seen from the line-cuts of Fig 4.7c.

From the line-cuts shown in Fig 4.7c, we see that for input signal power of 0.5 dBm (blue line), the dip in transmission is roughly symmetric. There is some asymmetry, particularly visible if the leftmost and rightmost part of the line-cut are compared), but we attribute this to the background being non-constant. For higher signal powers, 13.5 dBm (red) and 15.0 dBm (yellow), the cavity becomes asymmetric, and more shallow.

We see some more cavity behaviour that deserves comment in Fig. 4.7a. The cavity frequency is seen to increase slightly with increasing signal power. The change in frequency between -60 dBm and -10 dBm signal power is comparable to the cavity linewidth. We must take note of this if we want to apply a pump tone resonant with the cavity for different signal powers.

At -10 dBm and -9 dBm pump power, both cavities have a jump in their resonance frequency. We attribute this to them jumping from one flux arc to the next. If we apply a strong input signal, this could cause a stronger varying flux and the SQUID cavity might jump to a different flux arc. If and when the SQUID cavity jumps depends on how far it is biased, with further-biased SQUID jumping at lower signal strengths.

We have also applied a pump with varying strength, resonant with the cavity, to measure if the SQUID cavity shows similar behaviour as when we apply an input signal. The results of this measurement are plotted in Fig. 4.7b. Unlike the previous measurement, we do not see the resonance frequency of the SQUID change with applied pump strength. We do see a jump in the resonance frequency of the SQUID, for an applied pump power of 1.6 dBm.

At applied pump strengths above this 1.6 dBm, we see two peaks in the cavity transmission that we do not see for lower pump strengths. Line-cuts for applied pump strengths above (red) and below (blue) this change are shown in Fig. 4.7d. From the location of the dip in the transmission, we see that the cavity resonance frequency has jumped by several MHz. The two peaks are related to the pump tone we apply. The higher-frequency peak is the pump frequency we apply, the lower-frequency peak is not expected.

We have investigated the origin of this additional peak. It is not related to the microwave source used for pumping, as the peak remains when we use a different microwave source. It is also not related to

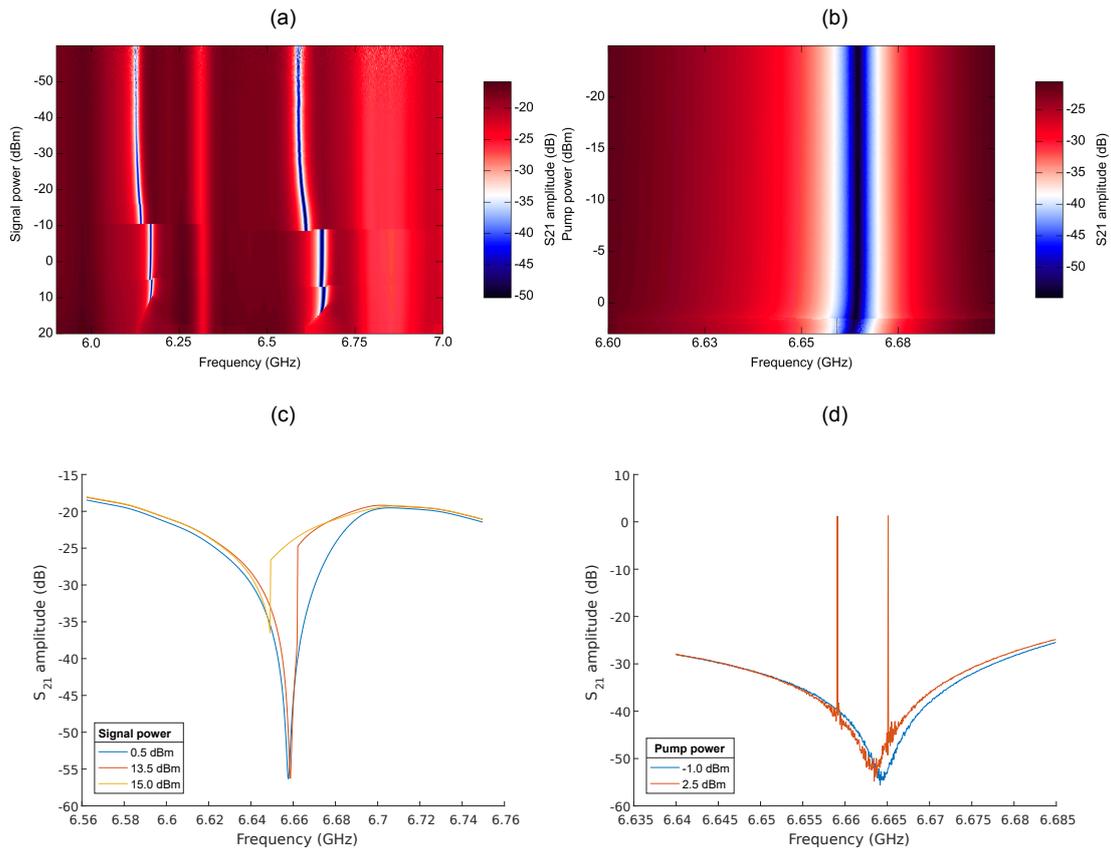


Figure 4.7: **SQUID cavity behaviour for pump and signal power.** a) Cavities 6 and 7 for increasing signal power. At 9.4 dBm and 11.9 dBm respectively, the cavities become non-linear. b) Cavity 7 for increasing pump power. At 1.6 dBm, there is a jump in cavity frequency, and two high intensity peaks become visible. The higher-frequency one is at half the pump frequency, the other is 6 MHz lower in frequency. c) Line-cut of SQUID cavity 7 for several signal powers. For low power the cavity dip is symmetric, but for higher powers we see it become more and more asymmetric. d) Line-cut of SQUID cavity 7 for various pump powers. For low pump power, the cavity is symmetric. It suddenly jumps a sufficient pump power, and we see two clear peaks which we attribute to the pump.

the specific frequency of the cavity, as it also happens if we pump at 8 GHz (where there is no cavity). It is possible that it is related to the synchronization of the Vector Network Analyzer (VNA) clock and the microwave source clock, but then the additional peak should occur 5 MHz below the pump peak, but it is very clearly 6 MHz below the pump peak.

4.2.4 Transmission amplitude versus pump strength

As first sign that our amplifier shows amplification, we look into the behaviour of the cavity when we apply a varying pump power. We expect that the depth of the dip in transmission of the cavity would decrease. We do not vary the input signal strength, but if there is amplification, the strength of the cavity output would increase if compared to no amplification. So, we would expect the depth of the transmission dip to decrease.

We have performed this experiment by applying a fixed-strength signal, and varying the strength of the pump. We have plotted the transmission of the cavity over a range of 50 MHz in Fig. 4.8a for different applied pump strengths. However, we do not see a significant change in the depth of the cavity transmission dip. We have tried different signal powers, but were unable to find a parameter range of applied input signal strength and pump strength where we found a noticeable difference in cavity dip depth.

It is possible that we have not found the right parameter regime (pump strength, input signal strength)

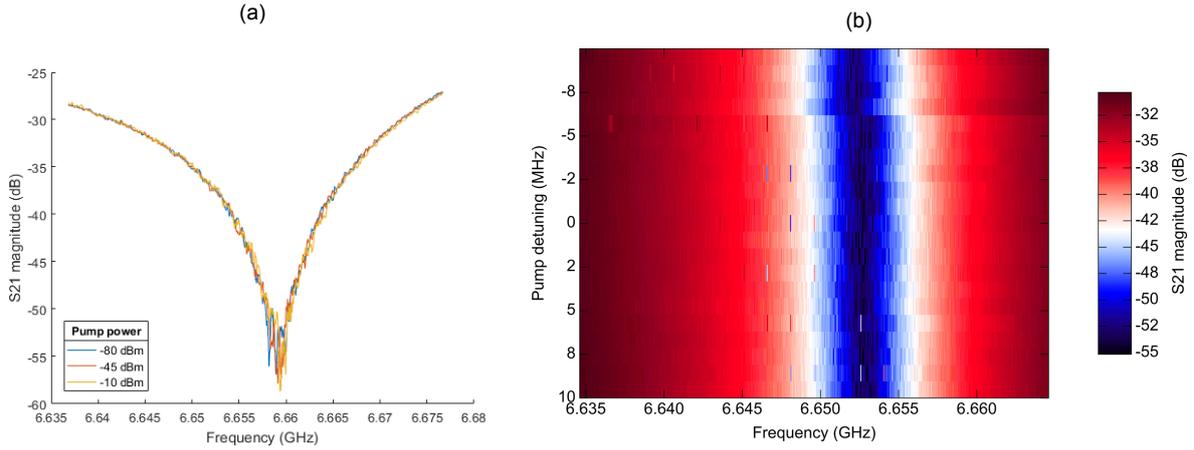


Figure 4.8: **Pumped SQUID cavity.** a) Sweeping the pump power does not measurably change the cavity signal. b) Pump detuning sweep. One every three traces shows the pump, but there are now four pump peaks. This is most likely due to aliasing.

to clearly see the change in cavity transmission. It is also possible that our experimental parameters (measurement integration time, frequency resolution) were such that the amplifier gain is difficult to distinguish from the noise. We could rule out the frequency resolution affecting our measurements by changing the frequency of the applied pump.

We have swept the detuning between (half) the pump frequency and the cavity resonance frequency, and measured the effect it has on the transmission. The results are plotted in Fig. 4.8b, and several things are noticeable. The first is that we see peaks in the cavity transmission that we can attribute to the pump, as they follow the detuning from the cavity frequency. Secondly, we notice that we see not one but four different peaks, in two sets of two. The peaks are separated by 1.5 MHz, the sets are separated by 3.0 MHz. The outer peaks match the detuning of the peaks measured in the previous section for high pump powers, but we are at lower pump powers (-8 dBm versus +1.6 dBm in the previous section).

Another noticeable thing is the fact that we only see the peaks periodically, only one in every three traces shows the peaks. This we can attribute to the pump having a very low linewidth compared to the frequency resolution of our measurement. Due to a mismatch of the pump and frequency steps, only one in every three traces has the pump exactly at one probe frequency point. The issue that the frequency resolution of the measurement is much larger than the linewidth of the pump could cause the measurement of Fig. 4.8a to not show any variation with pump strength.

4.2.5 Gain versus phase delay

In the previous section, our measurements have not shown any amplification, which could be due to the spectral resolution of our measurements. Here, we describe an experiment where we should be able to see whether our amplifier amplifies regardless of the spectral resolution of our measurement. To show parametric amplification, we study the system output while varying the phase of our input signal. We expect a gain-phase relation as shown in Fig. 3.17a.

However, we do not have a phase-shifter in our set-up. We can use a workaround: If the signal is slightly (0.1 Hz, in this case) detuned from half the pump signal, the phase difference between the signal and half the pump varies slowly. In time (every 10 s, in our case), we shift the phase by 2π , so we obtain in time similar curves as Fig. 3.17a. We have chosen the frequency difference of 0.1 Hz to allow an integration time of 1 s per measurement point to reduce the visible noise.

Simple amplifier model

One such measurement is shown in Fig. 4.9a (blue dots), where we have included a fit to a simple theoretical model is also plotted (orange line), with the model being

$$y = a \cdot \sqrt{\left(\frac{\sin(b\pi t + \phi)}{1-d}\right)^2 + \left(\frac{\cos(b\pi t + \phi)}{1+d}\right)^2}, \quad (4.5)$$

where a , b , d and ϕ are the fitting parameters, y is the linearised magnitude of the S_{21} signal and t is the time. a represents a vertical offset of the signal from 0 due to attenuators/amplifiers in the line and the cavity transmission dip, b represents the beating frequency, d represents the amplitude of the oscillations, and ϕ is a phase offset. We found $b = 0.096$ Hz for all fits, which is sufficiently close to the expected 0.1 Hz given the high frequency of the involved signals (6 GHz range).

An initial measure of the gain of the amplifier is described by the parameter d . For the maxima and the minima of the phase-gain curve, the fit equation reduces to $\frac{1}{1-d}$ and $\frac{1}{1+d}$. These describe the maximum amplification and deamplification [42, 44].

We have repeated the measurement of Fig. 4.9a for various combinations of input signal and pump power. We have calculated the amplifier gain for each of these measurements, and show the results in Fig. 4.9b. The JPA shows little gain for low signal powers, even for high pump powers. For higher signal strengths, there appears to be a peak in the gain, after which it drops off. The peak height appears independent of the pump power, but the signal power for which the gain peaks is affected by the pump strength. This plot does not conform to expectations from the simulation of a parametric amplifier, nor to what is expected from literature. We expect the gain to be constant for low signal power, and decrease for stronger signal powers, conform Fig. 3.18.

A partial explanation for this discrepancy may lie in the resonance frequency of the cavity being dependent on signal- and pump power. We bias the SQUID and determine the resonance frequency only for one signal power strength. Thus, for other combinations of signal and pump power, the resonance frequency may have shifted and we pump and measure off-resonant. This would lead to a significant decrease in gain.

Another partial explanation could be that the simplistic model used above does not describe our system well. We have side-coupled cavity measured in transmission, so the amplified signal emitted from our cavity could interfere with the directly transmitted signal. We also have a background of the cavity transmission dip, where we fit the data as if we only had the amplifier.

Involved amplifier model

We have used a more involved model of our parametric amplifier, which takes into account the fact that we have a side-coupled cavity measured in transmission, to see if the amplifier gain does conform to the simulations.

To take into account the background signal in the transmission spectrum, we have measured the transmission spectrum without a pump being sent in. Using these data, we have used a fitting script that separately fits the background and the cavity, and then fits them together using the initial parameters found in the from the separate fits to refine the results. This script is publicly available [71].

From this fitting script, we obtain the parameters of the background, and of the cavity. The fitting script fits a function of the form

$$S_{21} = (b_1 + b_2\omega + b_3\omega^2) \left(1 - \frac{\kappa_e e^{i\theta_0}}{\kappa_e + \kappa_i + \omega - \omega_0}\right) e^{i(b_4 + b_5\omega)} \quad (4.6)$$

Where the fit parameters are b_1 , b_2 , b_3 describing the amplitude variation of the background, b_4 and b_5 describing the phase variation of the background, and θ_0 describing the phase shift between the amplified signal emitted by the cavity with respect to the directly transmitted signal. We also obtain the

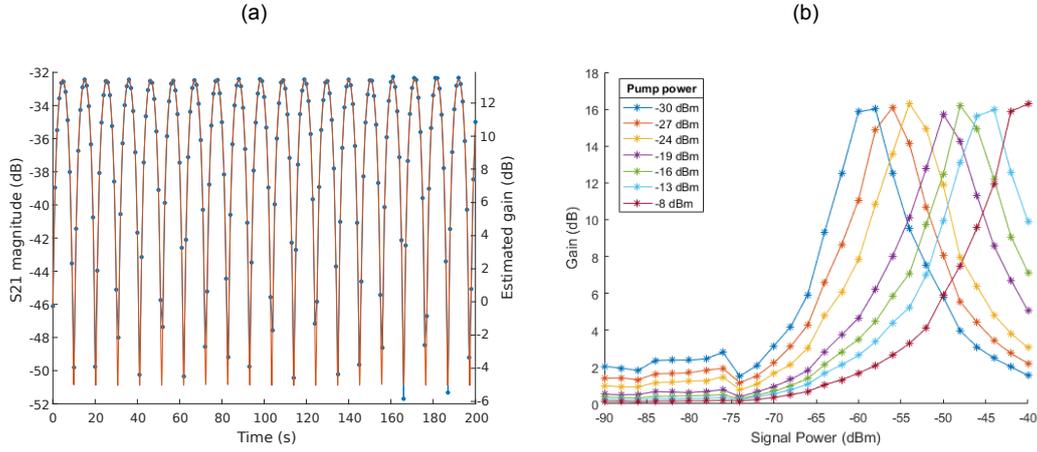


Figure 4.9: **Amplifier gain versus phase delay.** a) Measurement of S_{21} magnitude versus time. Blue dots are measurement, orange line is a fit to theory. We have estimated the gain of the amplifier from the fit to the theory curve. b) Overview of gain for various combinations of pump and signal power.

cavity parameters, which are the resonance frequency ω_0 and the internal and external quality factors, Q_i and Q_e respectively. We can calculate the internal and external loss rates, κ_i and κ_e , from these. One limitation is that if the difference between Q_i and Q_e is larger than an order of magnitude, we cannot determine the larger of the two accurately. In our system, we do not know the Q_i accurately for this reason.

With these background parameters found for our system, we can subtract the background from our total transmitted signal and end up with only the part of the transmission that is due to the cavity. We can then use the model derived in Sec. 2.1.3 to fit to our cavity signal to the transmission with the background subtracted,

$$S_{21,\text{subtr}} = 1 - \frac{e^{i\theta_0} \kappa_e}{\frac{\kappa^2}{4} + (\omega_s - \omega_0)^2 - a_c^2} \left(\frac{\kappa}{2} + i(\omega_s - \omega_0) + a_c e^{i(a_\omega)t + a_\phi} \right) + a_0. \quad (4.7)$$

We have denoted the parameters of our fit as a_c representing the amplitude, a_ω representing the frequency difference between the applied signal at ω_s and half the pump frequency ω_d , and a_ϕ a phase-offset. Lastly, we have found that an additional offset was necessary, which we have denoted as a_0 . This offset is a complex constant.

We have plotted the data of our measurement with the background subtracted in Figs. 4.10a and 4.10b. We have plotted the real and imaginary parts of S_{21} separately for clarity, as plotting them in a single plot would result in overlapping circles, as they describe several rotations through the complex plane. This is the oscillating behaviour that is visible in the plots, where we can more clearly separate the data (blue dots) and fit (orange lines). The fits do not look smooth, which is because we have calculated their values only for the data points and interpolated between them. This does not qualitatively affect the behaviour of the fit in the plot.

We have used this model to fit our measurements for various signal- and pump strengths. We have plotted the fit parameter of the amplitude, a_c , for these traces in Fig. 4.11. We see that for increasing pump strength, the amplitude fit parameter value decreases. For increasing signal strength, it increases. We have not yet been able to relate this amplitude parameter to the gain of the amplifier. It is directly related to the strength of the parametric pump as defined in Eq. 2.59,

$$a_c = \frac{\alpha_p^2}{4\omega_0} = \frac{\alpha^2}{4\omega_0^3}, \quad (4.8)$$

where we expect α to be of the order 10^{-3} based on our simulations. The parameter α in our simulations is normalized to the resonance frequency of the oscillator. Still, we would expect a_c to be of the order

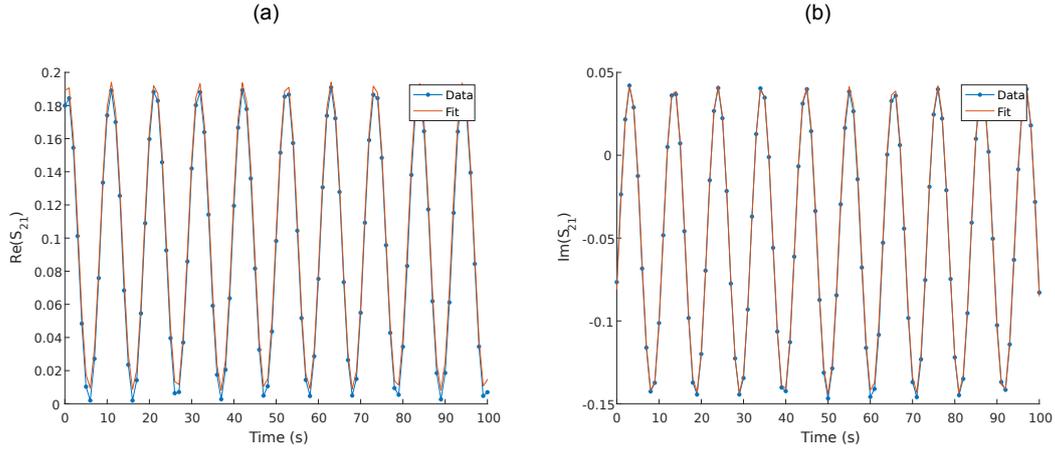


Figure 4.10: **Amplifier gain versus phase delay, refined model.** a) Plot of the real part of S_{21} over time (blue dots) and a fit (orange line) according to the model described in the main text. b) Plot of the imaginary part of S_{21} over time (blue dots) and a fit (orange line) to the same model. The value of the fit is calculated at the data points and interpolated between them, which causes the fit line to not be smooth.

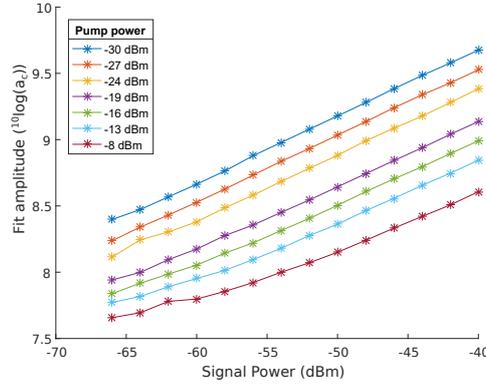


Figure 4.11: **Fit amplitude versus signal power.** Amplitude parameter of our fit to the data for various signal and pump strengths. The trends shown here do not follow the expected trend, and the value of the parameter is different from what is expected based on the simulations.

10^{-15} instead of 10^8 as we find from the fits. We conclude that something in either our fit or our analysis of the fit parameters with respect to the simulation parameters is wrong. Therefore, we have not been able to draw conclusions about the gain from this model.

4.2.6 Noise Temperature

One of the most important parameters of an amplifier is the noise temperature. In this subsection, we describe the measurements performed to gain insight in the noise temperature of our amplifier.

Because we have a chain of amplifiers (instead of only our single JPA) and cannot access the sample directly without warming up the fridge, we are severely limited in our choices of how to measure the noise temperature. For example, using the Y-factor method (Sec. 2.4) would require us to be able to control the noise temperature of the input signal. This could be achieved for example by connecting a resistor at the input of the amplifier, with a heater and a thermometer, neither of which are things we have.

Another option would be the gain method, this would require us to know the temperature of the input signal very accurately, which could be achieved with a resistor and a heater, or even by looking at the frequency splitting of a qubit, but it would also require us to know the gain of the amplifier very accurately.

A third option to gain an estimation of the added noise temperature of the JPA is to have a switch in parallel to it, so we could bypass it and calibrate out the rest of the amplifier chain. However, we do not have that option, as we lack the hardware of such a switch.

The only (seemingly) available avenue is to take the input noise temperature to be equal to the temperature of the fridge (250 mK), which is not an unreasonable approximation if we do not send in a signal (and the attenuation of our input line is properly done). If we have some gain in our amplifier, this should amplify the noise in our cavity above the rest of the noise.

To do this, we must first determine whether we can actually measure the thermal noise of our system. A limiting factor might be the HEMT amplifier, or alternatively the noise floor of the spectrum analyser. The rated noise temperature of the HEMT is 2.3 K [64]. With the equation for Johnson noise,

$$\frac{P}{B} = k_B T \quad (4.9)$$

where we use $B = 1$ Hz, we can find out the power level we expect, which is $P = -195$ dBm. The noise floor of our spectrum analyser (Displayed Average Noise Level, DANL) is rated at -146 dBm at 7 GHz. With the 39 dB gain from the HEMT (High Electron Mobility Transistor), and the 32 dB gain from the room temperature amplifier, the thermal noise should be visible, but attenuation in the lines might be an issue. We check it by giving no input signal or pump, and measuring the noise level in the frequency region we are interested in (approx 300 MHz width centred around our cavity frequency of 6.16 GHz). If we toggle the HEMT on/off, we should see the noise level of the amplifier change (it should be higher if the HEMT is on). If this is the case, and the power detected is significantly above -146 dBm, we can detect the thermal noise.

Then we try to measure the cavity thermal noise. This should be visible as an increase in noise from the cavity photons. Compare the energy of a photon at 6.165 GHz ($\hbar\omega = 4.08 \cdot 10^{-24}$ J) to the thermal energy at 250 mK ($k_B T = 3.45 \cdot 10^{-24}$ J), if there are some extra photons in the cavity, we should see it rise above the rest of the spectrum (where there is no cavity).

This was done, and the results are shown in Fig. 4.12. Here, we see the measured output signal for various pump powers. For no (and very low) pump powers, the noise is expected to be thermal noise. There appears to be a non-constant background, this could be a property of our sample. For higher pump powers, we expect to see the cavity noise rise faster than the noise at other frequencies. This can be observed in Fig. 4.12, particularly the trace for -10 dBm pump power has a region where it is elevated above the other traces. We also see an increase in output power that appears to be constant over the measured frequencies. This could be attributed to the pump increasing the local (noise) temperature of the sample.

For the highest pump powers, we see a narrow peak. We could attribute this to the pump leaking into the transmission line from the flux line, except that the pump was operated at twice the frequency we measure here (we operate the JPA in the 3-wave process). Thus, we attribute this peak to be the gain by the JPA. It is rather narrow in bandwidth (<1 MHz), but it is unlikely that the pump would leak in any significant amount at half the frequency, and it is possible that the gain is approximately 1 for a range of pump powers, and then it suddenly increases.

To make the increased cavity noise more visible, we can try to subtract the background. We do this by subtracting the (linear) amplitude of the zero-pump signal from all other traces. For some traces, this results in negative numbers due to the noisiness of the signal. To cope with this, we give everything a slight offset. This does not qualitatively affect our results.

For no pump, the temperature corresponding to the noise level (if we go by the specifications of the amplifiers) is 340 mK without the pump and 360 mK with the pump. This is a very small difference, considering the pump strength is sufficient to have a noticeable gain, and that should affect the noise contributed by the HEMT and the RT amplifier.

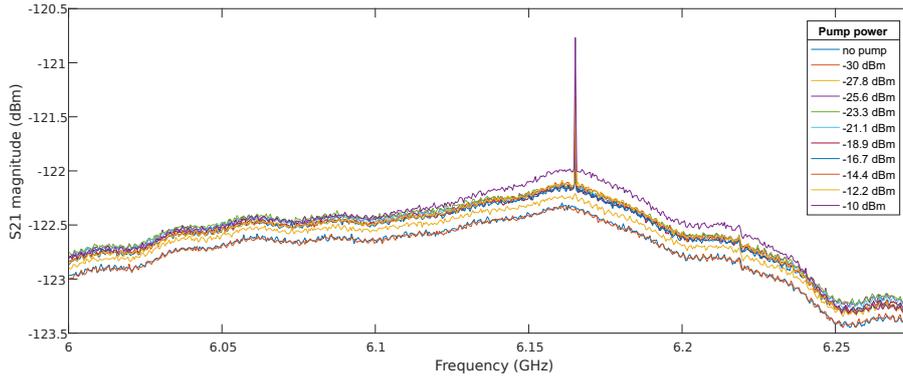


Figure 4.12: **Cavity noise.** Measurement of output from the amplifier for zero input, for various pump powers. For low pump powers, the S_{21} magnitude follows the same trend, but with an offset. For higher pump powers, the cavity noise becomes visible. The sharp peak is due to the gain of our amplifier.

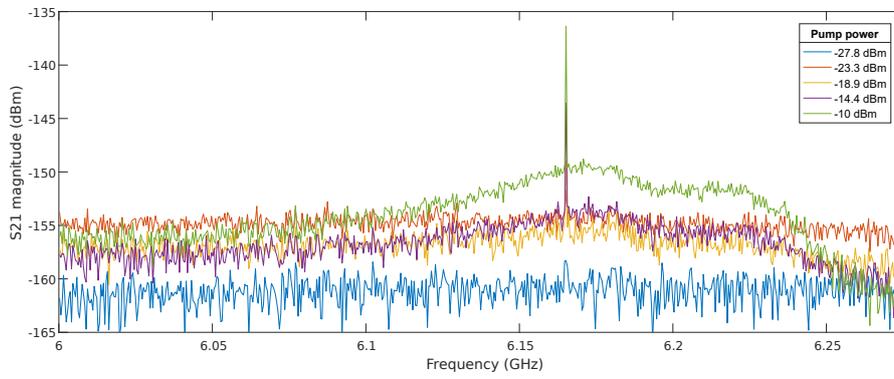


Figure 4.13: **Cavity noise with subtracted background.** Output from the amplifier for zero input, for various pump powers, with the background removed. From this, the added noise from the cavity for higher pump powers (and thus higher gains) is clearly visible. Without pumping, the cavity noise temperature corresponds to approximately 340 mK, with pumping the temperature increases to 360 mK.

We have not obtained the noise temperature of our amplifier from this measurement. During the analysis of the data, we had discovered that the sharp peak was actually due to the gain of the amplifier instead of due to the pump, as we pump at $2\omega_0 \approx 12.3$ GHz. To obtain the noise temperature of the amplifier, we should have measured the output signal for no input signal (i.e. the noise) at exactly the frequency of the amplifier. From the noisiness of this output, we can extract the noise temperature of the amplifier. From the height of the gain peak with respect to the background (outside the cavity), we can extract the gain. However, no measurement was performed at half the pump frequency for a prolonged time (to obtain the noisiness of the amplified signal).

4.2.7 4-wave mixing

So far, we have only operated the JPA in 3-wave mode. In this section, we describe the results from operating the JPA in 4-wave mode. To observe 4-wave mixing, the JPA was pumped at resonance (frequency ω_0), and a weak signal was sent in with a slight detuning from the pump frequency. We varied this detuning over 1 MHz, which is shown in Fig. 4.14a. Several line-cuts are taken from this data set, and shown in Fig. 4.14b.

In Figs. 4.14a and 4.14b, we see the pump at resonance (big peak in the center), the signal at a frequency denoted by the detuning Δ from the pump peak and on the other side of the pump, the (much smaller) idler is visible at the same detuning Δ . The presence of this idler, and the fact that it moves

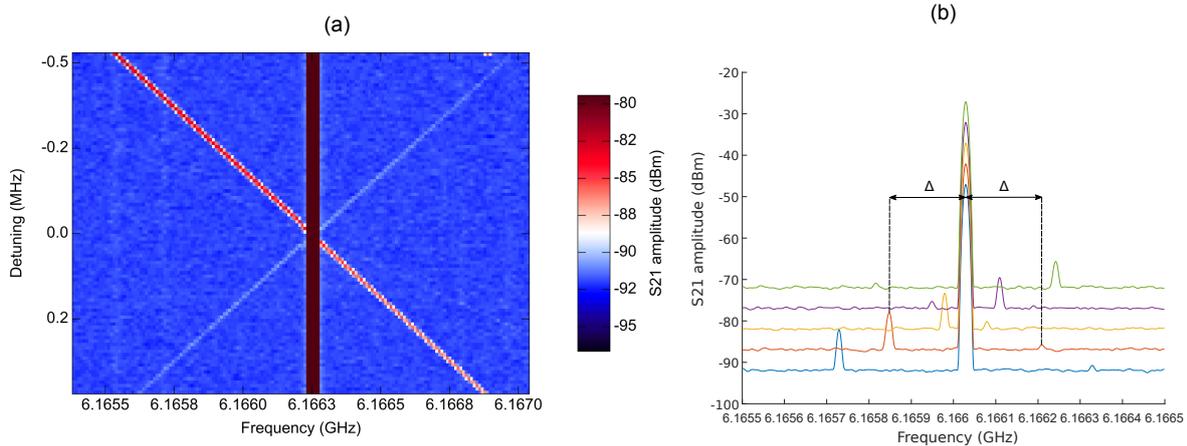


Figure 4.14: **4-wave mixing.** a) Probe signal detuning sweep over 1 MHz. We see the probe signal going from top left to bottom right, and the smaller idler going from top right to bottom left. There is a tiny line going the same direction as the probe, but at twice the detuning from the pump. There are also multiple side-bands of the pump, visible as vertical lines at the left side of the plot. b) Spectrum showing the 4-wave mixing process. The cavity is pumped resonantly (the big peak, frequency ω_0) and a weak probe is sent in with detuning Δ , and an idler is created detuned with Δ on the other side of the pump. Different curves are for various detunings, and offset with 5, 10, 15 and 20 dBm in increasing order.

opposite to the input signal with respect to the pump, show that we can operate our JPA as a 4-wave mixer. The asymmetry between the signal and idler peak heights is due to the fact that we measure cavities coupled in transmission. Part of the signal does not enter the cavity and is directly transmitted, while no such thing happens for the idler peak, as we do not apply a tone at the idler frequency.

Extracting the gain of our JPA from the data shown so far of 4-wave mixing is tricky. Ideally, the signal and idler peaks would be amplified by a factor of \sqrt{G} and $\sqrt{G-1}$ respectively [16], and we could extract our gain G from the ratio of the peak heights. However, as our amplifier works transmission mode, we see as the signal peak both the directly transmitted part that bypassed the cavity, and the amplified part that entered the cavity. These could interfere, so our cavity contribution can not be singled out easily.

We can make an estimation of the range of our gain, however. The upper bound of the gain is found if we take the strength of the noise background (-92 dBm in Fig. 4.14b) as the signal strength. The idler strength is then $\sqrt{G-1}$ times the signal strength, and we can read out both these strengths from the graph. This gives an upper limit of the gain of 7.1 dB. The lower limit is reached by attributing the entire signal peak as detected to the direct transmission. This would give a lower limit to the gain of 2.3 dB. If we take the nominal input signal strength, -89 dBm, we end up with a gain of 4.3 dB. We repeat this calculation for all traces of the data set of Fig. 4.14b and obtain an average (nominal-signal-strength) gain of 4.6 dB over the range of our measurement.

We would like to know how the JPA gain changes for various pump and signal strengths. The pump strength sweep is plotted in Fig. 4.15a. The uncertainty in the determination of the gain is quite considerable, no clear trend can be observed. For pump strengths outside the plotted limits, no idler could be observed. For lower pump strengths, the idler was lost in the noise, while for higher pump strengths, the pump side-bands prevented accurate determination of the idler peak height.

The pump side-bands are shown in Fig. 4.15b for the pump strengths relevant for Fig. 4.15a. The average separation between the pump side-bands is 131 kHz. For high pump powers, the side-bands are at least 50 dB smaller than the pump, while for lower powers the difference in power between the pump and the side-bands diminishes. This is in reasonable accordance to the specifications of the microwave source used to generate the pump tone [72]. The side-bands at higher frequencies are 3 dB smaller than the side-bands at lower frequencies. The spectral width of the side-bands is comparable to the spectral width of the main peak. The strongest side-bands are those detuned 790 kHz from the

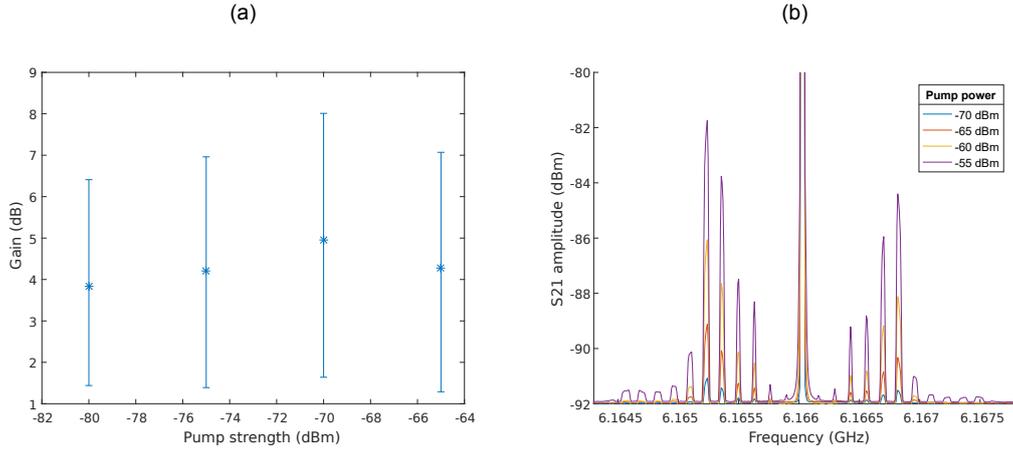


Figure 4.15: **4-wave amplification gain and pump sidebands.** a) Gain of our amplifier in 4-wave mode for various pump strengths, and an estimated signal strength of -89 dBm at the sample. Error bars are calculated as upper and lower limits as described in the main text. b) Side-bands of the pump for various pump powers. The side-bands are separated by 131 kHz, around the center pump peak, with the strongest side-band detuned by 790 kHz.

main peak, this is probably related to the internal filters in the microwave source.

We have also calculated the gain for various signal strengths, as plotted in Fig. 4.16a. As in the above parts, the uncertainty in the calculated gain is quite considerable. The upper and lower errorbars come from calculating the gain when taking the input signal to be equal to the noise, and when taking the input signal to be the full signal peak (respectively). However, for low signal strengths, the noise (-92 dBm) is higher than the expected signal strength. For these, the data point plotted is the average between the upper and lower limits.

We see from Fig. 4.16a that the gain goes down (even though the uncertainty in our gain determination goes up). This is what we would expect, if the amplifier saturates. However, we do not expect this to happen here, as the sum of all inputs to the cavity together is less than the power at which the cavities become non-linear (Sec. 4.2.3).

Instead of the cavities becoming non-linear, there is another process happening, which can be seen in Fig. 4.16b. Here, the input signal is detuned from the pump, and the detuning is varied. This results in the red line from the upper left to the lower right of the figure. However, there is also a white line at twice the detuning from the pump as the probe signal, in the same direction. This could be related to an interference between the probe signal ω_{probe} and the pump ω_{pump} , where the process is $2\omega_{\text{probe}} = \omega_{\text{pump}} + \omega_{2\Delta}$, where the 'twice-detuned' signal is denoted with $\omega_{2\Delta}$. The strength of this peak increases for signal strength, while the idler strength diminishes (causing the gain to drop by our calculations in Fig. 4.16a).

In fact, if one looks very carefully at Fig. 4.16b, there is also a very faint triply-detuned peak visible when the detuning is small. This could also be related to an interference between the signal and the pump.

Finally, to gain insight in the gain and bandwidth of our JPA in 4-wave mode, we determine the height of the signal peak for various detunings. This is shown in Fig. 4.17. The gap near 6.121 GHz is due to the pump peak overlapping with the signal/idler peaks, so their height cannot be determined accurately. The bandwidth appears to be 10 MHz. Note the frequency difference between the measurements of Fig. 4.14b and 4.17. They were taken from different cavities on the sample chip.

The signal peak strength could also be tracing out the cavity transmission, instead of the amplifier transmission. However, we have seen earlier in the characterization of our cavities that the linewidth of the

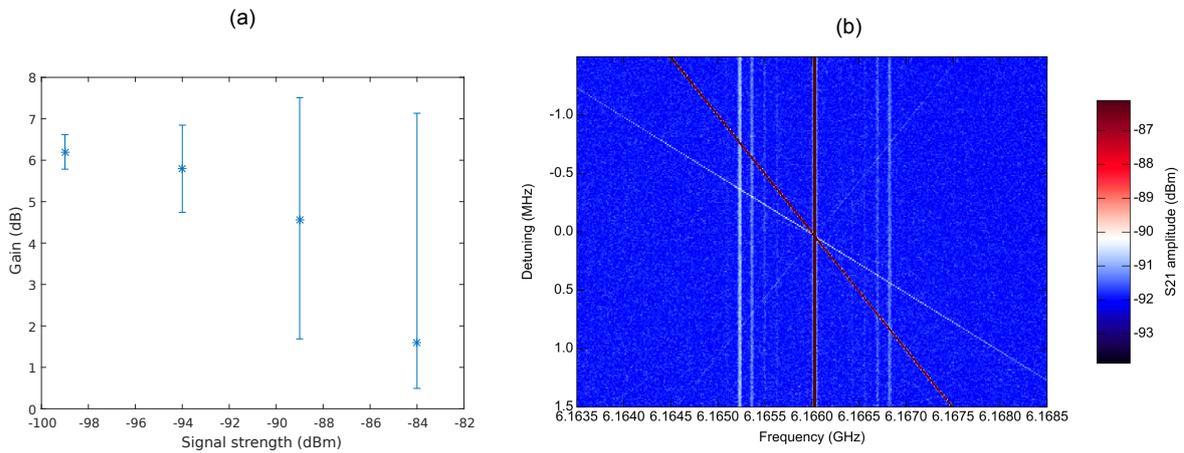


Figure 4.16: **4-wave amplification gain and signal-pump interference.** a) Gain of our amplifier in 4-wave mode for various signal strengths, for a pump strength of -70 dBm. The error bars are calculated as the upper and lower limits as described in the main text, except when the nominal signal strength is lower than the average (-92 dBm) noise level. Then, the nominal gain is taken to be the average between the upper and lower limits. b) For higher signal strengths, an additional peak at twice the detuning of the signal with respect to the pump is seen (white line from top left to bottom right). The idler (from top right to bottom left) is visible, but weaker than for lower signal strengths.

cavities is on the order of 50 MHz. Due to the difference of this number with the measured linewidth of the dip in the signal peak strength of Fig. 4.17, we attribute the linewidth to the bandwidth of the amplifier.

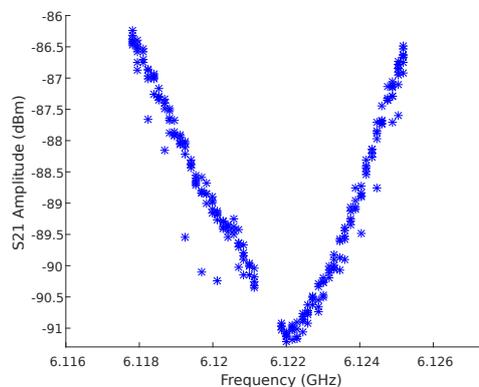


Figure 4.17: **4-wave amplifier bandwidth.** Plot of the signal peak strength versus the detuning traces the cavity. The gap near the middle is where the pump peak is. The width of the dip shows the bandwidth of our amplifier, approx 10 MHz if we extrapolate the trends linearly.

5

Conclusion

5.1 Josephson junctions

5.1.1 Fabrication

We have attempted to use a new method to fabricate superconductor-insulator-superconductor Josephson junctions. This bridge-less shadow method works based on the shadow generated by pre-patterned electrodes connecting the Josephson junction. This method allows more flexibility in the design of junctions, and does not require a suspended bridge of resist, which could collapse or be susceptible to damage.

We have found and solved several issues in the fabrication procedure of these Josephson junctions. The fabrication procedure has two important steps that require tuning, the etching step and the evaporation step. Our etch process step is susceptible to over-etching, but this problem can be avoided by monitoring the reflection of the sample with a laser. The etching of the MoRe layer provides a predictable curve in the reflectance. The second issue related to the etching step is the electrode side-wall angle. If the etch is isotropic, the side-walls of the electrode are rounded and can not form a shadow required to create the junction in the evaporation step. This can be solved by fine-tuning the etching parameters, or partially solved by evaporation at a higher incidence angle.

We have also found fabrication issues related to the evaporation of the Aluminium for the junctions. Firstly, the two layers of the junction must conform to the electrode side-walls, and if the layers are too thin, they might be disconnected. This can be avoided by evaporating a thicker layer of Aluminium. The second issue relates to this, because if the bottom layer is too thick, it could also create a shadow to disconnect the top layer. This can be avoided by making the top layer thicker than the bottom layer, of evaporating the top layer at a lesser angle. These issues share a common factor in that Aluminium can move around somewhat after deposition on the sample. This could be avoided by cooling the sample down during evaporation.

We have two other issues that required tuning of specific fabrication steps. The evaporated Aluminium films of the contact pads show (partially) popped bubbles. We suspect this is related to one of the later steps in the fabrication process causing a gas-producing reaction with contaminants present at the interface of the MoRe and the Aluminium. These bubbles do not appear to be present in small-area films, such as our Josephson junctions. If the origin of the bubbles is as we expect, the issue can be avoided by taking additional cleaning steps as described in [\[62\]](#).

The second issue pertaining to other steps in the fabrication is the with lift-off. The initial procedure

for lift-off was not always (fully) successful, and left parts of the sample still covered in Aluminium. Using an ultrasonic cleaner, and a pipette to blow away the Aluminium flakes using 80°, proved sufficient to remove the Aluminium from the sample completely.

The last, and most fundamental issue with the fabrication is the usage of Molybdenum-Rhenium as a base material for the parametric amplifier. This material grows Rhenium-oxide crystals under normal atmospheric conditions, which can affect the conductivity of the material. This process was not an issue for our samples, as they were kept mostly under a nitrogen-atmosphere when not in use, and the formation of Rhenium-oxide crystals happens over a time-span of days. However, for practical applications of the JPAs designed and fabricated in this thesis, the oxidation process is undesirable.

A clear path to avoid this oxidation issue is to switch to a different material for the base layer of the JPA design. For example, NbTiN (Niobium Titanium Nitride) could be used, or Aluminium. Both are materials accessible in our cleanroom facilities. However, switching to a different base layer material would involve fine-tuning the process parameters of all the steps in the fabrication process. Due to time and machine access constraints, we have chosen not to switch to either of the two mentioned materials as alternative to MoRe. However, for future projects that strive to create a practical JPA, both NbTiN and Al are viable options.

5.1.2 Characterization

One of the goals of this thesis was to characterize the Josephson junctions we have fabricated using the bridge-less shadow method. We have first attempted to measure their resistance at room temperature, Sec. 4.1.1, but we must conclude that we have not succeeded in measuring the room-temperature resistance of the Josephson junctions. We have found a correlation between the measured resistance and the designed electrical path length. The resistivity of our MoRe we can derive from this appears lower than the literature value, though the uncertainty is considerable. The geometry of the junctions and contact pads is not suitable for determining the junction resistance at room temperature, mostly due to the resistance in the leads. This could be overcome by using a four-terminal geometry very close to the junction, but due to the sheet resistance of approximately 8.0 Ω/sq , it would require careful calibration of the sheet resistance to see the junction resistance.

From Sec. 4.1.2, we find that the fabricated structures have to qualitatively different types of behaviour in the I-V curve. One type appears in three junctions that share a common critical current of 130 μA . However, this critical current is not affected by the junction size, and is too large for the Aluminium junctions following the Ambegaokar-Baratoff relation. From the lack of resistance below this current, it is a critical current, and we suspect it is related to a different structure on our sample. If this structure was a cylindrical Aluminium constriction, it would be approximately 2.6 nm in radius, which is much smaller than any of the expected junction size. We have attempted to find a conclusive origin for this phenomenon, but have been unable to find a suitable one.

The second type of behaviour in the I-V curve is a series of steps. We have proposed multiple explanations for the origin of these steps. They could be Shapiro steps, but would require driving at a frequency of 130 GHz to create steps of 270 μV . We do not send in such a frequency. They could also be (zero-field) Fiske steps, but do not behave as Fiske steps would in a magnetic field. They could be photon-assisted tunnelling steps, but this would require sending in a signal at a frequency of 65 GHz, which we do not.

We suspect the origin of these steps to be the thermoelectric effect. Having different cabling materials connected at different temperatures can create a thermal voltage, which might be sufficient to cause one of the Josephson junctions to emit radiation at the frequency required to induce Shapiro steps in a nearby junction. As our junctions are close together, we think this is the most likely explanation for the step-like behaviour.

In Sec. 4.1.3, we have measured SQUIDs with a critical current that is modulated by an applied magnetic field. Despite experimental error causing most of the magnetic field to be applied in-plane instead

of out-of-plane, the modulation is clearly visible. There are also $270 \mu\text{V}$ steps observed that are not modulated by the magnetic field, and appear in both SQUIDs and in single junctions.

From these results, we can conclude that we have fabricated Josephson junctions. Characterization of the junctions at room temperature requires a different geometry, and characterization at 250 mK was limited due to the Shapiro steps. Due to machine access issues, it was not possible to continue fabrication of the junctions.

5.2 JPA

5.2.1 Simulation

We have simulated a general parametric amplifier with an additional Duffing non-linear term. The goal of these simulations was to gain qualitative insight in how the parameters used in simulation (loss rate and quality factor of the oscillator, the phase and strength of the input signal and strength of the pump, and the sign and size of the Duffing parameter) affect the behaviour of the parametric amplifier. Although the model was simple and the simulation method not particularly optimized for accuracy, the qualitative behaviour of the simulated parametric amplifier matched expectations.

We have simulated the gain of the amplifier dependent on the phase difference of the input signal and the pump, and found the expected π -periodic function. We have found the gain to increase exponentially with pump strength, up to a point where the maximum amplitude of oscillation is limited by the Duffing term. We have found the expected three different regimes for the gain depending on signal strength.

For low input signal strength, the gain is linear. For higher signal strengths, the gain saturates. If the pump strength is above a certain value, the oscillator enters the self-oscillation regime, and oscillates independently of the input signal. The signal amplitude at which the saturation regime starts is determined by the Duffing term, for higher non-linearity the regime sets in at lower signal strength.

We have also studied the bandwidth of the parametric amplifier. For low pump strength, the gain is low and the bandwidth is large. For higher pump strengths, the gain increases and the bandwidth decreases, as expected from an amplifier with a constant gain-bandwidth product. However, for sufficiently high pump strength, the amplifier enters the saturation regime, which causes the gain-bandwidth product to increase by increasing the bandwidth. For a sufficiently high pump amplitude, we enter the self-oscillation regime, where we see a flat gain over a certain bandwidth. The slope of this gain plateau is determined by the Duffing term.

5.2.2 Characterization

SQUID cavities

We have a sample with seven SQUID cavities, five of which have a released mechanical element and two which do not. These cavities are the ones with a center frequency of 6.19 and 6.70 GHz respectively, and external quality factors Q_{ext} of 129 and 162. We have shown that we can bias the SQUID cavities to 0.998 times their zero-field resonance frequency in an automated, repeatable and controllable manner. For higher fields, the cavities can jump between different flux arcs. This can be remedied by going down in field to make the cavities jump back to their original arc.

We have studied the limits in terms of pump- and signal power of our SQUID cavities. The cavities can take 9.4 and 11.9 dBm signal power (generator output) before their transmission dip becomes asymmetric. The cavities are prone to jump between flux arcs for higher signal powers if they are biased. With respect to pump powers (in the 3-wave process, so pumped at twice their resonance frequency), we see peaks related to the pump at the resonance frequency, and 6 MHz detuned from that. These appear for pump powers of 1.6 dBm (generator output).

We have so far not been able to find a good explanation for this second peak. We see it appear for

sufficiently high pump powers regardless of cavity, biasing, microwave source, clock-synchronization between the detector and the source or frequency used in measuring.

Transmission amplitude versus pump strength

We have attempted to measure the gain of our amplifier directly from the depth of the dip in the cavity transmission. We do not see the cavity transmission change significantly for experimental parameters (signal and pump strength) or measurement parameters (frequency resolution, integration time). The range of pump strengths over which we have measured, -80 to -10 dBm, is sufficiently large that we expect this is not the issue.

To check if the frequency resolution of our measurements was the cause of the lack of visible gain of our amplifier, we have performed a measurement where we detuned half the pump frequency from the cavity resonance, and varied the detuning. In this measurement, we see four peaks in the transmission that change in frequency matching the change in detuning. We suspect that these four peaks share their origin with the two peaks seen in earlier measurements.

The set of four peaks only shows up periodically (that is, in the first out of every three measurements). The fact that these peaks do not show up in the other two of every three measurements leads us to suspect that the frequency resolution of our measurement is the limiting factor here. This could then also explain why we did not see transmission change with pump power.

Gain versus phase

To determine the gain without being affected by the frequency resolution of our measurement, we can measure the amplifier gain by measuring how the transmission changes when we change the phase difference between the input signal and (half) the pump. We do this by sending in the input signal as a slightly (0.1 Hz) lower frequency than half the pump. This way, we get a 2π phase shift every 10 seconds.

We have analysed the data obtained via this method with two different models. The first is a rather straightforward model that gives the gain of a parametric amplifier as a function of the phase [42]. This model fits very well to the data, but the gain (16 dB maximum) we can derive from this does not behave as expected. A possible reason for this is that the resonance frequency of the cavity shifts with input signal power, which can affect the measured transmission amplitude that we fit our model to. The second reason is that we measure a side-coupled cavity in transmission, and this is not taken into account in this model.

To remedy both points, we have characterized background signal and the change in cavity resonance frequency and quality factor for all the combinations of pump and signal powers used. We have used this characterization to subtract the background from the cavity signal. Furthermore, we have refined our model based on the derivations done in the theory chapter, Sec. 2.3.2, for a side-coupled cavity measured in transmission.

We have re-analysed the data from these measurements with this new model. It also provides a good fit to the measured data, but we were unable to translate the fitting parameters to the gain of our amplifier. We have attempted to do so via the derivation of the model, but there is a large difference between the obtained fit parameters and what is expected for a parametric amplifier based on the simulations done in Sec. 3.5.

Noise Temperature

The main benefit of using a JPA over other types of amplifiers is the lower noise temperature. We have attempted to characterize the noise temperature of the fabricated amplifier, but were quite limited in ways to achieve this accurately. There are several reasons for this. The first is that we want to know the noise temperature of the first amplifier in our chain of amplifiers. However, without accurate calibration of the other amplifiers in the chain, and knowledge of the gain of our amplifier, we are very limited in determining the noise temperature of the amplifier.

Another limiting factor was due to the set-up available in the fridge. To determine the noise temperature of the amplifier, we need to know accurately the noise temperature of the input to the amplifier. We could either measure this, using a qubit or using a switch to bypass our JPA, or control this, using a resistor with a heater. However, we have none of these options available, as the hardware is not present in our fridge.

We have attempted to determine the noise temperature of the JPA using the thermal noise of our fridge. By pumping the amplifier to achieve a sufficient gain, the cavity noise should be amplified, which we have observed (340 mK outside the cavity, 360 mK inside the cavity). We also saw a peak with a small linewidth (<1 MHz) that at first appeared to be an artefact of the pump. However, we operate the JPA in 3-wave mode, i.e. we pump at 12.3 GHz and measure at 6.15 GHz. It is unlikely that this peak is caused by the pump leaking into our output line. Instead, we attribute it to the gain of our parametric amplifier.

We could have determined the noise temperature of our amplifier, had we measured at the amplified frequency for a certain amount of time. By comparing the noise at this frequency to the noise at different frequencies, we can calculate the difference in noise temperature. However, we have only come to this conclusion during the analysis of the data, when no further measurements could be performed.

4-wave mixing

We have shown that our JPA can also be operated in 4-wave mixing mode. For several combinations of signal and pump powers, we see a gain of between 2 and 6 dB. The amplifier is limited in the pump strength by the pump side-bands that become stronger than the idler peak we use to characterize the gain. In signal strength, the JPA is limited by what appears to be an interference process between the signal and pump tones, which causes a peak with twice the detuning from the pump as the signal. For this case, the idler disappears almost completely.

We have characterized the amplifier bandwidth by plotting the intensity of the signal peak with respect to the detuning. The bandwidth we obtain from this is 10 MHz, which is more than the bandwidth observed in the previous section, but sufficiently less than the cavity linewidth that we do not attribute it to that. As we operate the JPA in a different mode and at different pump and signal strengths, we do not see a major issue in the difference between the bandwidth measured here (10 MHz) and earlier (1 MHz). Both values are comparable to the bandwidths of other JPAs used in research applications (Table 2.1).

We conclude that the JPA can be operated in 4-wave mode, but only in a limited regime of signal- and pump strengths. This could be remedied by changing the design to enhance amplification, for example by operating it in reflection mode.

Outlook

There are several avenues available for follow-up research. There are some fabrication issues remaining that were not (fully) resolved, which could be the subject of a future project. An important factor for this is that MoRe oxidises when left exposed to air, which could affect the functionality of an amplifier. Other superconducting materials could be considered as a base material for the amplifier, such as NbTiN or Aluminium.

We have also described fabrication issues related to the bridge-less shadow method of fabricating Josephson junctions. A major issue here appears to be the Aluminium moving around after deposition. We have suggested that this could be avoided by cooling the sample down during the evaporation, but have not actually done so. This could also be the subject of a future research project, albeit more materials science-oriented than this project.

In doing the analysis of the measurement data and writing up this thesis, we have found several measurements that could have been performed to better describe the characteristics of our JPA. For example, we could have improved our description of the gain of the amplifier as a function of the phase difference between the input signal and half the pump. We also could have determined the noise temperature of

our amplifier from measurements at half the pump frequency and at a sufficiently different frequency. From the noise amplitude of both measurements, we could have found the noise temperature of the amplifier, and from the difference in average amplitude, we could have derived the gain of the amplifier. All these measurements can be done in a follow-up project.

There are also several design aspects that could be changed to benefit the functioning of the JPA, such as the size of the SQUID loop, the position and length of the flux pump line and the critical current of the junctions. Taking into account the previous points, a logical step would be to re-design the JPA and fabricate it. With the measurement methods described here, including the improvements detailed above, it should be possible to better characterize a JPA.

Acknowledgements

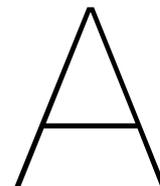
I want to thank first of all Ines, Daniel and Gary for their supervision, for their training in matters related to the clean room and the labs, for their efforts in explaining the relevant theoretical concepts and for keeping the project on track and me focussed. Without them, this project would not even have started, let alone be finished in a somewhat timely manner.

I would also like to thank the rest of the Steele lab, for their help and discussions when I needed to do something new or was stuck on a particular problem.

Further, I'd like to thank my fellow master students in QN for their camaraderie during lunch- and coffee breaks, for their collegiality regarding complaints about machines and samples not working, or having to write down the work done in a neat and organized manner.

I also want to thank my family and friends, for being the people I could contact every half an hour when there was nobody else in the cleanroom on weekends.

No animals were harmed in the creation of this thesis



JPA recipe

For completeness, the latest iteration of the full recipe to create Josephson Parametric Amplifier is given. The work for this thesis was performed in the cleanroom of the Kavli Nanolab, Delft. Care has been taken to denote the steps of the recipe such that they could be applied by others or in other labs.

Step 1: Clean the chips.

Start with a $15 \times 15 \text{ mm}^2$ pre-diced square chip of $500 \text{ }\mu\text{m}$ Si with polished native oxide on top. Put the chip in an acetone bath for 30 s, and rinse afterwards. Repeat this two more times, and then repeat it three times with IPA instead of acetone.

Step 2: MoRe layer.

Sputter 100 nm MoRe on the sample (in our case, using 100 W power, 10 μbar pressure and 4 min 30s). Then, spin-coat the chip with AR-P 6200.13 at 4000 rpm, and pre-bake for 3 min. Pattern the MoRe electrodes, transmission lines and MoRe layer using an e-beam with a $180 \text{ }\mu\text{C}/\text{cm}^2$ dose.

After the exposure, develop the resist using a Pentylacetate bath (60 s), followed by a MIBK:IPA 1:1 bath (60 s) and an IPA bath (60 s). Then, etch with 12.5 sccm SF_6 and 10 sccm O_2 with 10 μbar pressure and 50 W RF power. In our system, the etching takes approximately 3 min, and the etch progress can be monitored with a laser. After the etching, strip the remaining resist using a hot (80°C) PRS3000 bath for approximately 1 hour, followed by an IPA rinse.

Step 3: Al junctions.

First, we make a double layer of resist to get the required undercut to make the junctions. Spin-coat the chip with PMGI SF7 at 4500 rpm, and bake it for 15 min at 180°C . Then, spin-coat PMMA 950K A3 at 5500 rpm, and bake it for 15 min at 175°C .

Make the pattern for the junctions, using an e-beam with $2200 \text{ }\mu\text{C}/\text{cm}^2$ dose. Then, develop the exposed chip with a MIBK:IPA 1:3 bath for 70 s, an IPA bath for 70 s, an MF321 bath for 10 s, followed by two separate H_2O baths of 15 s and finally another IPA bath of 20 s.

Then, perform an additional cleaning step before the evaporation. Clean it with an O_2 plasma, 200 W for 30 s with 200 sccm O_2 . Our system has a Faraday cage, which is important if this recipe is followed exactly (otherwise, the O_2 plasma damages the sample). Then, perform an HF dip to clean the surface. We used a BOE (Buffered Oxide Etch) 7:1 dip for 30 s, followed by two H_2O baths of 30 s each.

Next, the evaporation of the junctions. This is the most critical step, and it must be performed quickly after the HF dip to prevent oxidation. After loading the sample, we pump down to 10^{-7} mbar , and evap-

orate some nm of Ti to trap the remaining gases to the wall of the chamber. We evaporate the first layer of Al, 35 nm at a 45° angle. We let the sample oxidise for 10 minutes in an O₂-rich atmosphere at 0.5 mbar, and evaporate the second layer of Al, 45 nm at again a 45° angle.

After unloading, we must lift-off the remaining Al from the sample. We do this with a bath of NMP at 80° C for 20 minutes, with a pipette to create extra flow. We follow this by a Acetone bath (30 s) and an IPA bath (30 s).

Step 4: Dicing and mounting.

The last step involves dicing the chip to size to expose the microwave connectors near the side, and to mount it to on the PCB for connections to the fridge cables. This was done in our case by glueing the chip, and wire-bonding the connecting pads.

Bibliography

- [1] M. Faraday. On a peculiar class of acoustical figures; and on certain forms assumed by groups of particles upon vibrating elastic surfaces. *Philosophical Transactions of the Royal Society London*, 121:299, 1831. doi:[10.1098/rstl.1831.0018](https://doi.org/10.1098/rstl.1831.0018).
- [2] T. G. Roer. *Microwave electronic devices*, volume 10 of *Microwave and RF Techniques and Applications*. Springer US, 1st edition, 1994. ISBN 978-0-412-48200-7. doi:[10.1007/978-1-4615-2500-4](https://doi.org/10.1007/978-1-4615-2500-4).
- [3] E.D. Reed. The variable-capacitance parametric amplifier. *IRE Transactions on Electronic Devices*, 6(2):216, 1959. doi:[10.1109/T-ED.1959.14473](https://doi.org/10.1109/T-ED.1959.14473).
- [4] C.S. Aitchinson, R. Davies, and P.J. Gibson. A simple diode parametric amplifier design for use at S, C and X band. *IEEE Transactions on Microwave Theory and Techniques*, 15(1):22, 1967. doi:[10.1109/TMTT.1967.1126364](https://doi.org/10.1109/TMTT.1967.1126364).
- [5] R. A. Baumgartner and R.L. Byer. Optical parametric amplification. *IEEE Journal of Quantum Electronics*, 15(6):432, 1979. doi:[10.1109/JQE.1979.1070043](https://doi.org/10.1109/JQE.1979.1070043).
- [6] Hamradioacademy.com. What is a parametric amplifier, March 2018. URL <https://hamradioacademy.com/amateur-extra-what-is-a-parametric-amplifier/>. Obtained 09-07-2018.
- [7] Spectra Physics. *TOPAS Prime automated OPA*. Newport Corporation, 3635 Peterson Way, Santa Clara, CA 95054, USA. Obtained 09-07-2018.
- [8] B.D. Josephson. Possible new effects in superconductive tunneling. *Physics Letters*, 1(7):251, 1962. doi:[10.1016/0031-9163\(62\)91369-0](https://doi.org/10.1016/0031-9163(62)91369-0).
- [9] J. Clarke and A.I. Braginski, editors. *The SQUID handbook: Volume I*. Wiley-VCH, 1st edition, 2004. ISBN 3-527-40229-2. doi:[10.1002/3527603646](https://doi.org/10.1002/3527603646).
- [10] C. M. Caves. Quantum limits on noise in linear amplifiers. *Physical Review D*, 26(8):1817, 1982. doi:[10.1103/PhysRevD.26.1817](https://doi.org/10.1103/PhysRevD.26.1817).
- [11] P.K. Tien. Parametric amplification and frequency mixing in propagating circuits. *Journal of Applied Physics*, 29(9):1347, 1958. doi:[10.1063/1.1723440](https://doi.org/10.1063/1.1723440).
- [12] R. Bradley, J. Clarke, D. Kinion, L. J. Rosenberg, K. van Bibber, S. Matsuki, M. Mück, and P. Sikivie. Microwave cavity searches for dark-matter axions. *Reviews of Modern Physics*, 75:777, 2003. doi:<https://doi.org/10.1103/RevModPhys.75.777>.
- [13] L. Zhong, S. Al Kenany, K. M. Backes, B. M. Brubaker, S. B. Cahn, G. Carosi, Y. V. Gurevich, W. F. Kindel, S. K. Lamoreaux, K. W. Lehnert, S. M. Lewis, M. Malnou, R. H. Maruyama, D. A. Palken, N. M. Rapidis, J. R. Root, M. Simanovskaia, T. M. Shokair, D. H. Speller, I. Urdinaran, and K. A. van Bibber. Results from phase 1 of the HAYSTAC microwave cavity axion experiment. *Physical Review D*, 97:092001, 2018. doi:[10.1103/PhysRevD.97.092001](https://doi.org/10.1103/PhysRevD.97.092001).
- [14] N. Bergeal, F. Schackert, M. Metcalfe, R. Vijay, V. E. Manucharyan, L. Frunzio, D. E. Prober, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret. Phase-preserving amplification near the quantum limit with a Josephson ring modulator. *Nature*, 465:64, 2010. doi:[10.1038/nature09035](https://doi.org/10.1038/nature09035).
- [15] B. Abdo, K. Sliwa, S. Shankar, M. Hatridge, L. Frunzio, R. Schoelkopf, and M. Devoret. Josephson directional amplifier for quantum measurement of superconducting circuits. *Physical Review Letters*, 112:167701, 2014. doi:[10.1103/PhysRevLett.112.167701](https://doi.org/10.1103/PhysRevLett.112.167701).
- [16] M. Hatridge, R. Vijay, D.H. Slichter, J. Clarke, and I. Siddiqi. Dispersive magnetometry

- with a quantum limited squid parametric amplifier. *Physical Review B*, 83:134501, 2011. doi:[10.1103/PhysRevB.83.134501](https://doi.org/10.1103/PhysRevB.83.134501).
- [17] B. Yurke. Squeezed-state generation using a Josephson parametric amplifier. *Journal of the Optical Society of America B*, 4(10):1551, 1987. doi:[10.1364/JOSAB.4.001551](https://doi.org/10.1364/JOSAB.4.001551).
- [18] B. Yurke, L. R. Corruccini, P. G. Kaminsky, L. W. Rupp, A. D. Smith, A. H. Silver, R. W. Simon, and E. A. Whittaker. Observation of parametric amplification and deamplification in a Josephson parametric amplifier. *Physical Review A*, 39(5):2519, 1989. doi:<https://doi.org/10.1103/PhysRevA.39.2519>.
- [19] M. H. Devoret and R. J. Schoelkopf. Superconducting circuits for quantum information: an outlook. *Science*, 339(6124):1169, 2013. doi:[10.1126/science.1231930](https://doi.org/10.1126/science.1231930).
- [20] C.K. Alexander and M.N.O. Sadiku. *Fundamentals of electric circuits*. McGraw Hill, 5th edition, 2013. ISBN 978-0-07-338057-5.
- [21] D.M. Pozar. *Microwave engineering*. John Wiley & Sons, Inc., 4th edition, 2012. ISBN 978-0-470-63155-3.
- [22] M0UKD Amateur Radio Station Information Page. End fed half wave antenna coupler, February 2016. URL <https://m0ukd.com/homebrew/baluns-and-ununs/end-fed-half-wave-antenna-tuned-coupler-efhw/>.
- [23] R.N. Simons. *Coplanar waveguide circuits, components, and systems*. John Wiley & Sons, 2001. ISBN 0-471-16121-7.
- [24] H.W. Ott. *Noise reduction techniques in electronic systems*. John Wiley & Sons, 2nd edition, 1988. ISBN 0-471-85068-3.
- [25] C.P. Wen. Coplanar waveguide: A surface strip transmission line suitable for nonreciprocal gyromagnetic device applications. *IEEE Transactions on Microwave Theory and Techniques*, MTT-17(12):1087, 1969. doi:[10.1109/TMTT.1969.1127105](https://doi.org/10.1109/TMTT.1969.1127105).
- [26] P.W. Anderson and J.M. Rowell. Probable observation of the Josephson superconducting tunneling effect. *Physical Review Letters*, 1963. doi:[10.1103/PhysRevLett.10.230](https://doi.org/10.1103/PhysRevLett.10.230).
- [27] T. van Duzer and C.W. Turner. *Principle of superconducting devices and circuits*. Prentice Hall, 2nd edition, 1998. ISBN 0132627426.
- [28] A. Barone and G. Paternò. *Physics and applications of the Josephson effect*. Wiley, 1982. ISBN 0471014699.
- [29] S. Shapiro. Josephson currents in superconducting tunneling: the effect of microwaves and other observations. *Physical Review Letters*, 11(2):80, 1963. doi:[10.1103/PhysRevLett.11.80](https://doi.org/10.1103/PhysRevLett.11.80).
- [30] P. Seidel. Josephson effect in iron based superconductors. *Superconductor Science and Technology*, 24:043001, 2011. doi:[10.1088/0953-2048/24/4/043001](https://doi.org/10.1088/0953-2048/24/4/043001).
- [31] Ambegaokar. V. and A. Baratoff. Tunneling between superconductors. *Physical Review Letters*, 10(11):486, 1963. doi:[10.1103/PhysRevLett.10.486](https://doi.org/10.1103/PhysRevLett.10.486).
- [32] F. London. On the problem of the molecular theory of superconductivity. *Physical Review*, 74(5):562, 1948. doi:[10.1103/PhysRev.74.562](https://doi.org/10.1103/PhysRev.74.562).
- [33] B.S. Deaver and W.M. Fairbank. Experimental evidence for quantized flux in superconducting cylinders. *Physical Review Letters*, 7(2):43, 1961. doi:[10.1103/PhysRevLett.7.43](https://doi.org/10.1103/PhysRevLett.7.43).
- [34] R. Doll and M. Näbauer. Experimental proof of magnetic flux quantization in a superconducting ring. *Physical Review Letters*, 7(2):51, 1961. doi:[10.1103/PhysRevLett.7.51](https://doi.org/10.1103/PhysRevLett.7.51).
- [35] M. Tinkham. *Introduction to Superconductivity*. McGraw-Hill Inc., 2nd edition, 1996. ISBN 0070648786.
- [36] R.C. Jaklevic, A.H. Silver, and J.E. Mercereau. Quantum interference effects in Josephson tunneling. *Physical Review Letters*, 12(7):159, 1964. doi:[10.1103/PhysRevLett.12.159](https://doi.org/10.1103/PhysRevLett.12.159).
- [37] A.H. Silver and J.E. Zimmerman. Quantum states and transitions in weakly connected superconducting rings. *Physical Review*, 157(2):317, 1967. doi:[10.1103/PhysRev.157.317](https://doi.org/10.1103/PhysRev.157.317).
- [38] M. Devoret and A. Roy. Introduction to quantum-limited parametric amplification of

- quantum signals with josephson circuits. *Comptes Rendus Physique*, 17(7):740, 2016. doi:[10.1016/j.crhy.2016.07.012](https://doi.org/10.1016/j.crhy.2016.07.012).
- [39] J.W. Strutt. On the mainenance of vibrations by forces of double frequency, and on the propagation of waves through a medium endowed with a periodic structure. *Philosophical magazine*, 24(147): 145–159, 1887. doi:<https://doi.org/10.1080/14786448708628074>.
- [40] G. Grynberg, A. Aspect, and C. Fabre. *Introduction to Quantum Optics and Lasers*. Cambridge University Press, 1st edition, 2010. ISBN 9780521551120.
- [41] H. A. Haus. *Waves and fields in optoelectronics*. Prentice-Hall series in solid state physical electronics. Prentice-Hall Inc., Englewood Cliffs, New Jersey 07632, 1984. ISBN 0139460535.
- [42] D. Rugar and P. Grütter. Mechanical parametric amplification and thermomechanical noise squeezing. *Physical Review Letters*, 67(6):699, 1991. doi:[10.1103/PhysRevLett.67.699](https://doi.org/10.1103/PhysRevLett.67.699).
- [43] M. A. Castellanos Beltran. *Development of a Josephson parametric amplifier for the preparation and detection of nonclassical states of microwave fields*. PhD thesis, University of Colorado, 2010.
- [44] T. Yamamoto, K. Inomata, M. Watanabe, K. Matsuba, T. Miyazaki, W.D. Oliver, Y. Nakamura, and J.S. Tsai. Flux-driven Josephson parametric amplifier. *Applied Physics Letters*, 93:042510, 2008. doi:[10.1063/1.2964182](https://doi.org/10.1063/1.2964182).
- [45] A. Roy and M. Devoret. Quantum-limited parametric amplification with Josephson circuits in the regime of pump depletion. *ArXiv Preprint*, 2018. URL <https://arxiv.org/abs/1801.10115>.
- [46] T. Roy, S. Kundu, M. Chand, A. M. Vadiraj, A. Ranadive, N. Nehra, M. P. Patankar, J. Aumentado, A. A. Clerk, and R. Vijay. Broadband parametric amplification with impedance engineering: Beyond the gain-bandwidth product. *Applied Physics Letters*, 107:262601, 2015. doi:[10.1063/1.4939148](https://doi.org/10.1063/1.4939148).
- [47] N. E. Frattini, V. V. Sivak, A. Lingenfelter, S. Shankar, and M. H. Devoret. Optimizing the nonlinearity and dissipation of a SNAIL parametric amplifier for dynamic range. *ArXiv Preprint*, 2018. URL <https://arxiv.org/abs/1806.06093>.
- [48] S. Simbierowicz, V. Vesterinen, L. Grönberg, J. Lehtinen, M. Prunnila, and J. Hassel. Flux-driven Josephson parametric amplifier for sub-GHz frequencies fabricated with side-wall passivated spacer junction technology. *ArXiv Preprint*, 2018. URL <https://arxiv.org/abs/1805.07307>.
- [49] M. Malnou, D. A. Palken, L. R. Vale, G. C. Hilton, and K. W. Lehnert. Optimal operation of a Josephson parametric amplifier for vacuum squeezing. *Physical Review Applied*, 9:044023, 2018. doi:[10.1103/PhysRevApplied.9.044023](https://doi.org/10.1103/PhysRevApplied.9.044023).
- [50] G. Liu, T.-C. Chien, X. Cao, O. Lanes, E. Alpern, P. Dekker, and M. Hatridge. Josephson parametric converter saturation and higher order effects. *Applied Physics Letters*, 111:202603, 2017. doi:[10.1063/1.5003032](https://doi.org/10.1063/1.5003032).
- [51] N. E. Frattini, U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret. 3-wave mixing Josephson dipole element. *Applied Physics Letters*, 110:222603, 2017. doi:[10.1063/1.4984142](https://doi.org/10.1063/1.4984142).
- [52] S. Al Kenany, M. A. Anil, K. M. Backes, B. M. Brubaker, S. B. Cahn, G. Carosi, Y. V. Gurevich, W. F. Kindel, S. K. Lamoreaux, K. W. Lehnert, S. M. Lewis, M. Malnou, D. A. Palken, N. M. Rapidis, J. R. Root, M. Simanovskaia, T. M. Shokair, I. Urdinaran, K. A. van Bibber, and L. Zhong. Design and operational experience of a microwave cavity axion detector for the 20-100 μeV range. *Nuclear Instruments and Methods in Physics Research A*, 854:11, 2017. doi:[10.1016/j.nima.2017.02.012](https://doi.org/10.1016/j.nima.2017.02.012).
- [53] S. Pogorzalek, K. G. Fedorov, L. Zhong, J. Goetz, F. Wulschner, M. Fischer, P. Eder, E. Xie, K. Inomata, T. Yamamoto, Y. Nakamura, A. Marx, F. Deppe, and R. Gross. Hysteretic flux response and nondegenerate gain of flux-driven Josephson parametric amplifiers. *Physical Review Applied*, 8:024012, 2017. doi:[10.1103/PhysRevApplied.8.024012](https://doi.org/10.1103/PhysRevApplied.8.024012).
- [54] B. Abdo, Chavez-Garcia J. M., M. Brink, G. Keefe, and J. M. Chow. Time-multiplexed amplification in a hybrid-less and coil-less Josephson parametric converter. *Applied Physics Letters*, 110:082601, 2017. doi:[10.1063/1.4976962](https://doi.org/10.1063/1.4976962).
- [55] W. F. Kindel, M. D. Schroer, and K. W. Lehnert. Generation and efficient measurement of single photons from fixed-frequency superconducting qubits. *Physical Review A*, 93:033817, 2016. doi:[10.1103/PhysRevA.93.033817](https://doi.org/10.1103/PhysRevA.93.033817).

- [56] A.A. Clerk, M.H. Devoret, S.M. Girvin, F. Marquardt, and R.J. Schoelkopf. Introduction to quantum noise, measurement, and amplification. *Reviews of Modern Physics*, 82(2):1155, 2010. doi:[10.1103/RevModPhys.82.1155](https://doi.org/10.1103/RevModPhys.82.1155).
- [57] H. Nyquist. Thermal agitation of electric charge in conductors. *Phys. Rev.*, 32(1):110, 1928. doi:[10.1103/PhysRev.32.110](https://doi.org/10.1103/PhysRev.32.110).
- [58] L. Zhong, E. P. Menzel, R. Di Candia, P. Eder, M. Ihmig, A. Baust, M. Haerberlein, E. Hoffmann, K. Inomata, T. Yamamoto, Y. Nakamura, E. Solano, F. Deppe, A. Marx, and R. Gross. Squeezing with a flux-driven Josephson parametric amplifier. *New Journal of Physics*, 15:125013, 2013. doi:[10.1088/1367-2630/15/12/125013](https://doi.org/10.1088/1367-2630/15/12/125013).
- [59] S. Franssila. *Introduction to Microfabrication*. John Wiley & Sons, Ltd., 2010. ISBN 978-0-470-74983-8.
- [60] Van Leeuwenhoek Laboratory. Kavli nanolab equipment page, July 2018. URL <https://www.tudelft.nl/en/faculty-of-applied-sciences/about-faculty/departments/quantum-nanoscience/kavli-nanolab-delft/equipment/>.
- [61] G.J. Dolan. Offset masks for lift-off photoprocessing. *Applied Physics Letters*, 31(5):337, 1977. doi:[10.1063/1.89690](https://doi.org/10.1063/1.89690).
- [62] J. Dalle. Quantum circuits for mechanics. mathesis, Delft University of Technology, 2018.
- [63] A. van den Broek. Josephson junctions made by atomic layer deposition. mathesis, Delft University of Technology, 2018.
- [64] LNF-LNC4_8C. Low Noise Factory, rev: march 2017 edition, March 2017. URL http://www.lownoiseactory.com/files/1114/8943/8809/LNF-LNC4_8C.pdf.
- [65] AFS3-04000800-07-10P-4. L3 Narda-MITEQ, 100 Davids Drive, Hauppauge, NY 11788, 2018. URL <https://nardamiteq.com/viewmodel.php?model=AFS3-04000800-07-10P-4>.
- [66] P.D. Nation, M. P. Blencowe, and E. Buks. Quantum analysis of a nonlinear microwave cavity-embedded dc SQUID displacement detector. *Physical Review B*, 78:104516, 2008. doi:[10.1103/PhysRevB.78.104516](https://doi.org/10.1103/PhysRevB.78.104516).
- [67] V. Singh, B. H. Schneider, S. J. Bosman, E. P. J. Merks, and G. A. Steele. Molybdenum-rhenium alloy based high-q superconducting microwave resonators. *Applied Physics Letters*, 105:222601, 2014. doi:[10.1063/1.4903042](https://doi.org/10.1063/1.4903042).
- [68] D. D. Coon and M. D. Fiske. Josephson ac and step structure in the supercurrent tunneling characteristic. *Physical Review*, 138(3A):744, 1965. doi:[10.1103/PhysRev.138.A744](https://doi.org/10.1103/PhysRev.138.A744).
- [69] C. Camerlingo, M. Russo, and R. Vaglio. Zero-field Fiske steps in small Josephson junctions. *Journal of Applied Physics*, 53:7609, 1982. doi:[10.1063/1.330134](https://doi.org/10.1063/1.330134).
- [70] W. C. Danchi, F. Habbal, and M. Tinkham. Photon-assisted tunneling and AC Josephson effect at 246 and 604 GHz in small-area superconducting tunnel junctions. *IEEE Transactions on Magnetics*, 19(3):498, 1983. doi:[10.1109/TMAG.1983.1062376](https://doi.org/10.1109/TMAG.1983.1062376).
- [71] Steele Lab. Stlab, 2018. URL <https://github.com/yausern/stlab>.
- [72] *Keysight Technologies E8257D PSG Microwave Analog Signal Generator Data Sheet*. Keysight Technologies Inc., 2017.