A System for Cancellation of Two-Level System Noise in Kinetic Inductance Devices

Thesis by

Alex Meiburg

Submitted in partial requirements

for the requirements of the degree of Bachelor of Science

May 18, 2018

Abstract

Kinetic Inductance Detectors (KIDs) are showing promise in a variety of low-light applications photometry applications, notably in observing B-mode polarization of the cosmic microwave background. These devices are read out by modulating the inductance of an LC resonator through light, and observing the shift in resonant frequency. Among several contributing sources of noise is Two-Level System noise (TLS noise) that causes low-loss drift in the frequency. Under certain assumptions of the source of the noise, we propose a new dual-resonator design that would allow the TLS noise to be observed independently of the signal, and thus cancelled out. This design comes at a roughly factor-of-2 cost in component size and sensitivity. We designed a manufactured a niobium-on-silicon chip, but encountered issues in that we were unable to observe enough TLS noise to conclusively say that the cancellation works.

Acknowledgments

I would like to thank Professor Jonas Zmuidzinas, my adviser, for his leadership on this thesis. I am totally indebted to Roger O'Brient and Bryan Steinbach for many hours of meetings, design revision, and even car rides to JPL and back. Roger's insistence to get parts of this written up piece by piece saved me in terms of time management. I would also like to thank Albert Wandui for much aid in completing my experiments and designing my chips, and Lorenzo Minutolo for his excellent software – and patience as I learned how to use it.

And of course, I would like to thank my mom and dad: my dad who got me excited about physics and research from a young age, and my mom who made sure that I would have the ability to get this thesis written. You have both inspired me in so many ways.

Contents

Chapter 1

Background

1.1 Cosmic Microwave Background

Under the standard cosmological model, the early universe had much higher density and temperature than today's. When the average energy of photons was still much higher than the ionization energy of hydrogen (13.6eV), photons were subject to abundant Thomson scattering, a type of interaction in which the hydrogen would ionize and exchange momentum and energy with the photon. This implied that the universe was optically opaque to the majority of electromagnetic radiation. As the universe expanded, the density of gas decreased, and photons lost energy due to expanding spacetime elongating their wavelengths. Ultimately, they decoupled, and a large fraction of the light from this "time of last scattering" has continued on a free path ever since. Although this temperature was approximately 3000K, spacetime continued to expand, so that these photons were gradually redshifted down to 3K. This 3K radiation is a very uniform feature of the sky, known as the Cosmic Microwave Background. It has been precisely measured to a temperature of 2.72K, when corrected for the blueshift due to the movement of Earth through our solar system and galaxy.

The expansion of spacetime was not entirely uniform, however. Inflationary gravitational waves are believed to have existed during the time of the last scattering and shortly thereafter, reflecting the turbulence of inflation far earlier in the universe's history. This gravitational waves would have modified the nearby photons, selectively modifying their energy along different axes. Although true thermal radiation (as the CMB was) is unpolarized, this

would introduce patterns of polarization, correlated across the sky.

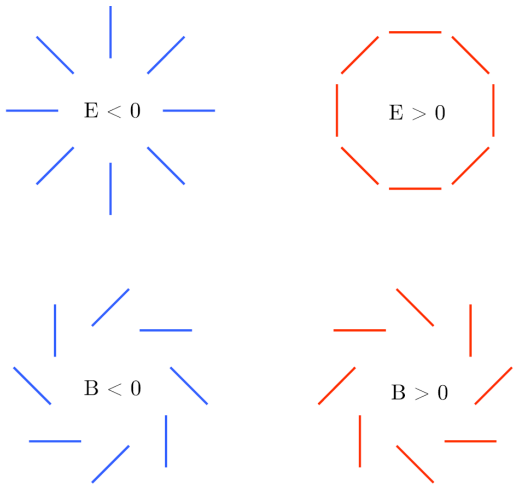

Figure 1.1: Characteristic Pattern of E-mode and B-mode Polarization (Dodelson et al., 2009)

Other perturbations in the cosmic medium, such as local density fluctuations, also create E-mode polarizations, with B-mode polarizations quadratically suppressed. Gravitational waves, however, produce both types of polarization in equal quantity: a B-mode is simply an E-mode, rotated by $\pi/4$.

There are a number of efforts underway to do precision measurements of the B-mode polarization of the cosmic microwave background. By placing constraints on the development of the early universe, parameters of cosmological and standard model can be inferred to new levels of precision, including inflation rates and neutrino masses.

1.2 BICEP and the Keck Array

The BICEP (Background Imaging of Cosmic Extragalactic Polarization) and Keck Array are two such surveys. Located at the Scott Amundsen station at the South Pole, they have unique advantages of observation, and challenges in terms of construction. The extremely low humidity removes one of the largest barriers to land-based observation. Extremely stable weather due to constant temperatures and flat landscaping leads to lower atmospheric

turbulence. Finally, high altitude removes a large fraction of atmosphere, period.

Figure 1.2: The Dark Sector Laboratory at Amundsen-Scott South Pole Station. At left is the South Pole Telescope. At right is the BICEP2 telescope. (WM Commons, CC-BY-SA-3.0)

Figure 1.3: Keck Array at the Martin A. Pomerantz Observatory (foreground) and Dark Sector Laboratory (background) at Amundsen-Scott South Pole Station. (WM Commons, CC-BY-SA 3.0)

Existing Keck array and BICEP3 detectors are based on Transition Edge Sensor (TES) bolometers. These devices function by slightly increasing temperature when light shines on them, edging a material out of superconducting temperature regime into normal; the transition is detected as the appearance of resistance to electrical current. A relatively new type of detector, known as a kinetic inductance device, operates based on the principle of instead directly transforming photons into quasiparticles in a superconductor – their functioning will be explained more thorough in the next section. It is believed that these devices are capable of outperforming TES devices, and the next generation of BICEP telescope is expected to built on them. They are still a relatively new (10 years since inception) technology though, and their design is far from a problem with an established solution.

Chapter 2

Kinetic Inductance Devices

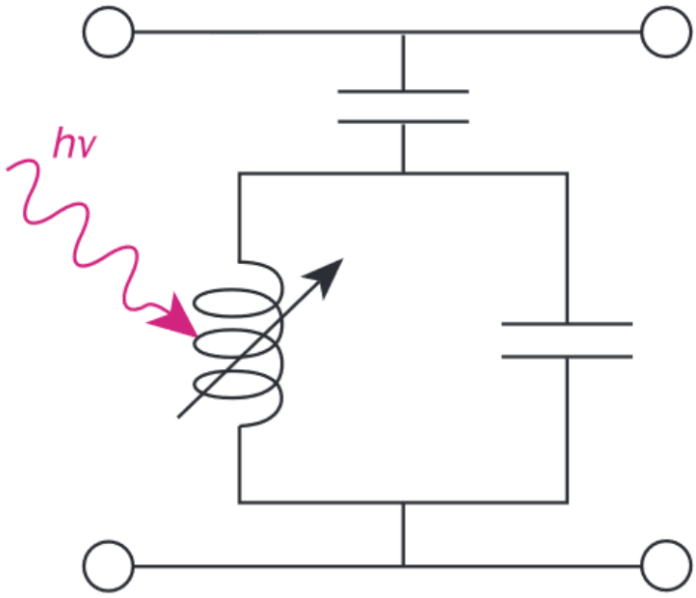

Figure 2.1: Schematic of an abstract KID circuit element, consisting of two capacitors and an inductor across a transmission line. A photon impinges on the inductor, interacting to create quasiparticle excitations. The top capacitor is responsible for coupling the resonator to the transmission line by a controlled amount, while the capacitor on the right creates the resonance. (After Mazin et al 2012)

The KID is focused on one RLC lumped element, placed across a transmission line. The circuit material is necessarily superconducting, so losses due to resistance are low (but nonzero): rather, the imperfect coupling through the coupling capacitor (top of Figure 2.1) limits the quality of the resonator. The inductive term is not created from a tradition magnetic element, but rather from the *kinetic inductance* of the wire material. Kinetic inductance is a high-frequency correction that occurs in all conductors (not just superconductors) when the relaxation time of the electron motion is noted to be nonzero. For phasor voltages and currents, the Ohmic law

$$\tilde{V} = \tilde{I}R$$

is now necessarily modified to include a complex resistance. In the the Drude model of electrical conduction, the complex resistance becomes

$$R = \frac{m(1 + i\omega t)}{ne^2\tau}$$

where τ is the relaxation time, m is mass of the charge-carrier, ω is the source frequency, n is the charge carrier number density. In normal materials are below-optical frequencies, $\omega\tau \ll 1$, and so imaginary corrections are negligible: the material is Ohmic. For superconductors however τ grows up towards scales of 1 second, so that already at megahertz frequencies there are resistances dominated by their imaginary parts. In a superconductor, the exact relaxation time can be computed from more conventional data using the fact that energy needed for one Cooper pair of flow must equal the energy stored in kinetic induction. This leads to the relation that

$$L_K = \frac{m\ell}{2n_s e^2 A}$$

where L_K is the kinetic inductance, m is the mass of the charge carrier, n_s is the Cooper pair density, ℓ is the wire length, e is electron charge, and A is the wire cross sectional area. The factor of two arises because Cooper pairs have twice the mass of the charge carriers, being formed of two carriers. As all of the other parameters in this relation will stay essentially fixed over a given circuit, any changes in n_s will be responsible for a change in the inductance. This change in inductance will then be observable as a different resonant frequency. By playing tones on the transmission line near the resonant frequency of the resonator, and observing the resulting

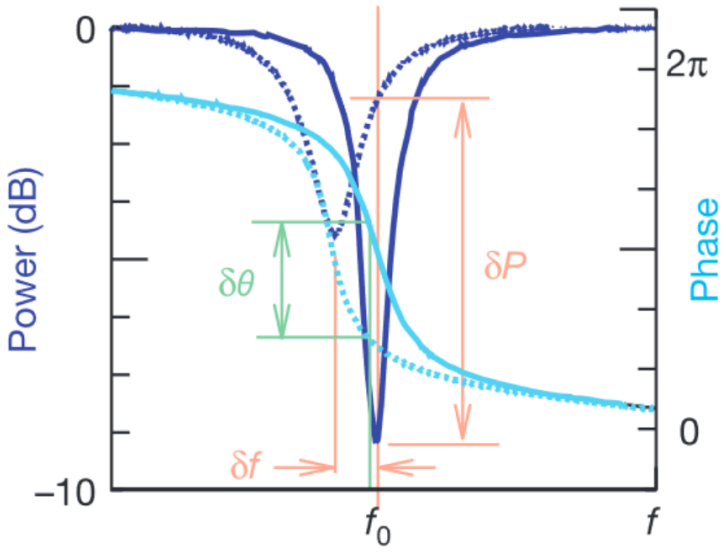

Figure 2.2: Shift of resonant response pattern in response to reduction in n_s. The resonant frequency shifts by an amount of δ_f. The phase shift observed at (undisturbed) resonance, f_0 is δ_θ. The power transmission ratio observed at f_0 is δ_P. (After Mazin et al 2012)

absorption and phase shift, modulations in resonant frequency can easily be observed down to one part in a hundred thousand. This corresponds to an equal fractional shift in the quasiparticle density, so that impinging photons create a dramatic effect. Long and thin wires are used to modify ℓ and A to bring the inductance within a range not too plagued by other factors (such as geometric inductance of the loop).

At this point, the capacitor can be chosen so as to bring the resonant frequency within a sufficiently low to have easily available readout hardware. Because the changes in resonant frequency are so slight and each only require

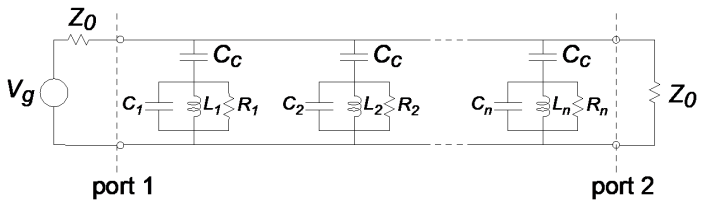

Figure 2.3: Circuit diagram of a multiplexed KID transmission line.

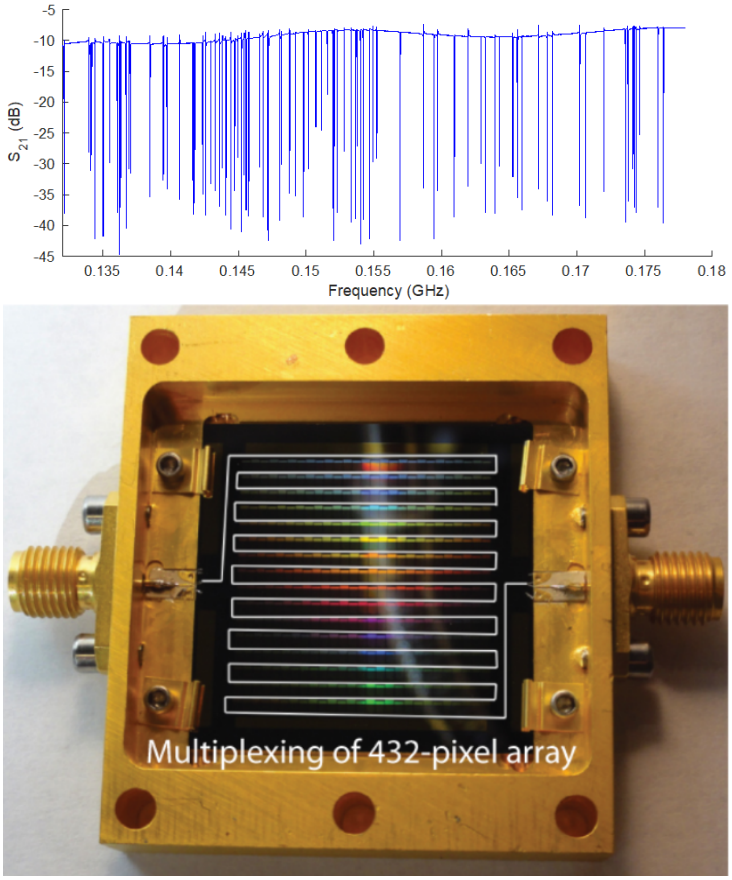

Figure 2.4: Top: Power transmission spectrum on a line of 432 KID resonators in series. Each high-Q resonator appears as an extremely narrow trough in transmission, so that they can be individually tracked as they move in response to incoming light. Bottom: The physical chip. Approximately 1 in^2.

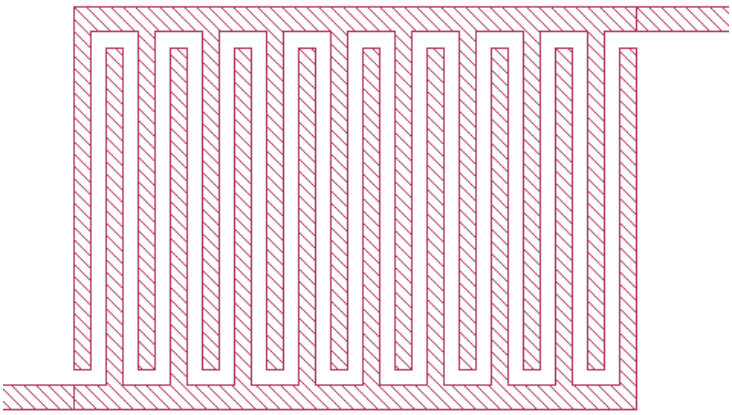

Figure 2.5: Semiconductor artwork of a simple interdigitated capacitor.

a transmission line, these devices are easily multiplexed: a thousand of them are placed on the same transmission, sharing virtually all fixed cost. Each device is engineered to have a slightly different resonant frequency which will then shift independently, and a collection of tones is then played on the transmission line and demodulated out afterwards. See Figure 2.3 for a circuit and Figure 2.4 for an example spectrum. To allow all these devices to be manufactured in one semiconductor fabrication iteration, without attaching additional components, the capacitors are constructed from *interdigitated* wires. In this setup, many wires from two electrodes are interlaced, allowing them to act many parallel "plate" capacitors in parallel (although these plates are very thin in one direction!) For a sense of scale, our chips used on the order of 100 fingers, each approximately 1mm long, and 1μm thick and in their gaps, and these produced capacitances on the order of of picoFarad. See Figure 2.5 for an example of interdigitated capacitor semiconductor artwork.

2.1 ADR Cooling

To function effectively as a photodetector, the KID must be cooled not only to the superconducting transition, but indeed to well below the 2.7K of the CMB, for otherwise the ambient heat of the surrounding material will overwhelm the sky. Thus in a typical application (and our tests), a KID is cooled

to 100mK, within a factor of a couple. Understanding how this cooling is achieved is relevant to understanding the noise concerns we wished to study. These low temperatures are achieved through a three-stage refrigerator. The first stage consists of a simple liquid nitrogen pump, bringing the whole mass to 77K. The second stage pumps on liquid helium, down to the range of 3K. At this point the third stage can be run, which is an adiabatic demagnetization refrigerator (ADR). The principle of an ADR is the *magnetocaloric effect*, in which magnetization of a metal removes some degrees of freedom, warming it up – and then removing the magnet later introduces new freedom again, allowing it to cool down its surroundings. Mounted on the third stage is a large molten salt pill, with the capacity to magnetically polarize. While it is thermally coupled to the second stage, a magnetic field is turned on (via a large inductive coil) over the course a minute. The magnet falls into the new potential for its dipoles, gaining thermal excitation but giving up entropy in the dipoles. Fifteen minutes elapse, and the excitation from the third stage has flowed into the second stage, and the system has re-equilibrated at 3K. A clamp is opened to thermally de-couple the third stage from the second, and the third stage is suspended only via very thin, thermally-noncoductive rods. Now the magnetic field is gradually removed, and the dipoles are freed to suck thermal energy from the environment. This brings the third stage down to 50mK.

2.2 KID Noise

KIDs have their strong point in extremely low-light environments. They require long exposure times (ideally under 30fps) in order to sample a detector (or all, simultaneously) for sufficient time to determine the frequency shift. The frequency of light they are sensitive is to some degree constrained by the material: the impinging photons will excite superconducting Cooper pairs to a normal state, and so the energy of the photon must be at least equal to the bandgap energy of the material in the superconducting phase. According to the BCS theory of superconductivity, the energy gap is

$$\Delta E = 1.764 k_B T_c$$

where k_B is the Boltzmann constant and T_c is the critical temperature of the material. For materials like niobium or titanium nitride which KIDs are often made of, which have crtical temperatures of 9.26K and 5.6K respectively,

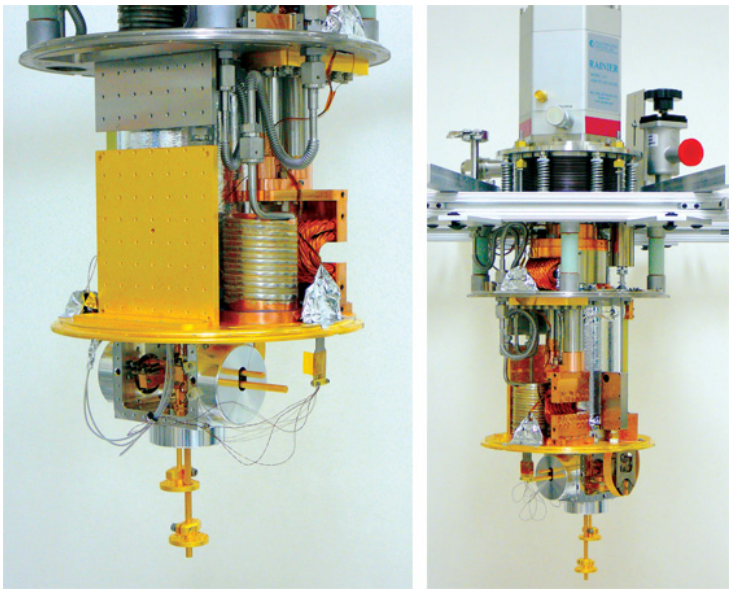

Figure 2.6: Entire HDP 103 Cryostat, with first stage starting near the top and the device to be mounted at the bottom.

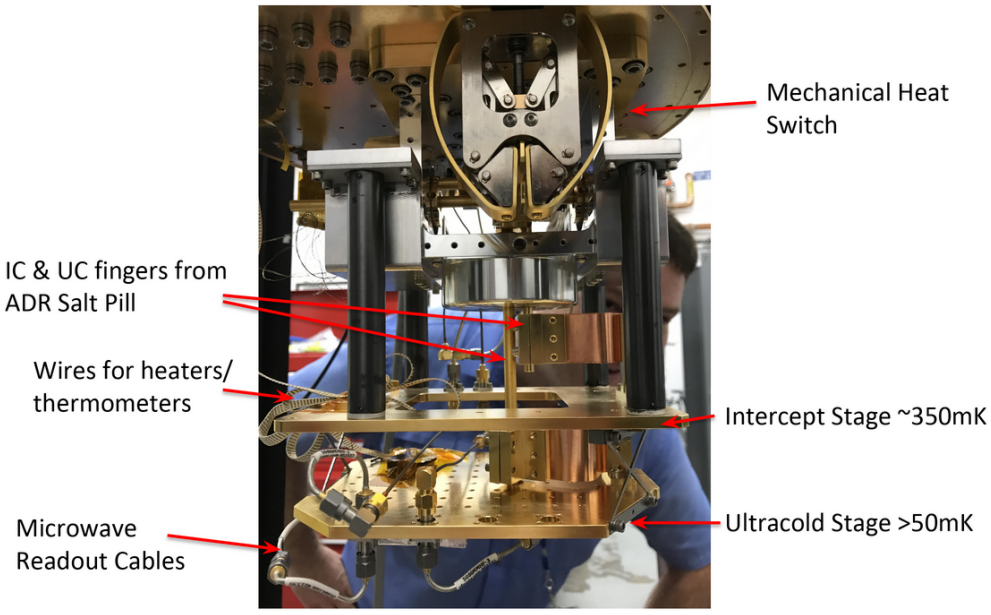

Figure 2.7: Labelled HDP 103 Cryostat, at the interface between second and third stages.

these are bandgaps of 1.4meV and 0.849meV. Thus these absorb radiation best on the order of 16K or 9.9K, respectively.

Beyond evaluating their sensitivity, there is the matter of considering what types of noise are observed in the detectors. The noise we would ideally be limited by is the *shot noise* of the incoming photons, that our signal arrives in particle packets and not as a continuous stream of energy. If our device were limited by shot noise, we would in a sense be getting the best data coming through the atmosphere.

In reality, there are a variety of other sources of noise that affect our readout. At the basic level, the device is *not* at 0K, and so produce its own random thermal excitatiosns. The number of quasiparticles is not fixed and will randomly fluctuate up and down, with fluctuations increasing at higher temperatures and decreasing with a larger band gap. This is termed "generation-recombination" noise, or GR noise. We chose niobium for our tests, which has a relatively large band gap specifically to suppress this type of noise: its critical temperature is 9.26K, one of the highest of any elemental superconductors.

The cryostat's helium pulse tube produces a lot of (audible) noise. Although best efforts are put in to de-couple its vibrations from the final stage, including vacuum inside the cryostat and minimal supports, some phonons are transmitted through the cryostat. These phonons are also capable of exciting quasiparticles, and so this shows up as signal as well. The pulse tube has a rhythm of approximately 2Hz. This frequency shows up very noticeably in the Fourier transform of the readout.

There is noise in the readout line. In order to not "drown" the detector in energy when the test tone is sent down the transmission line, the tone must be very small in power. If such a tone is transmitted from outside however, it will be drowned out entirely by thermal noise in the transmission line itself. In order to overcome this, the tone is transmitted at higher power, and each stage of fridge the input signal is attenuated to bring it to an energy level accordant with that temperature. By the time it reaches 50mK, the signal is at an appropriately small level. This is then amplified back up to (roughly) the same magnitude of power as it goes back up to warmer stages. These amplifiers necessarily create noise as they amplify thermal noise back up. This all means we will get warmer than normal noise temperature at the end of the transmission line. This type of noise does not directly modify the apparent resonant frequency, but will make it more difficult to read out exactly what the transmission coefficient is, ultimately hurting our readout

speed in a slightly different way.

There is usually some capacity for the circuit to couple to ambient electromagnetic fields. Radio waves or the pulse tube, for instance, can produce oscillating fields thart can induce small currents in the resonator. This poses an issue because kinetic inductance as a phenomenon is slightly nonlinear. Superconductors have a critical current beyond which a superconducting phase cannot be maintained, and as this limit is released, the kinetic energy of the Cooper pairs increase slightly superquadratically. Thus, there is a small increase in inductance when the average current in the resonator is higher. This effect is actually exploited in so-called Nonlinear KIDs where deliberate excitations are now used to "tune" the kinetic inductance, and thus the resonator frequency, to different value at the time of experiment (Kher, 2017). In any application however, stray excitations cause noise in the resonant frequency. This is mostly tackled through extensive magnetic shielding around the device.

The last type of noise is the one with the least complete model of its operation, and the one which we focus on tackling. The surface of the superconducting metal is typically coated in a thin oxide layer, which will be amorphous and not participating in the superconductivity. Its random, amorphous structure means that – by the large of numbers – some atoms will be in "loose" positions, and likely able to shift between nearby states with only slightly different energies and a low energy barrier. This is a "Two-Level System", or TLS. If such a TLS then couples to the nearby superconductor, it is possible for the TLS to, as it switches, modulate the effective parameters of the superconductive material (such as carrier density or critical temperature). An ensemble of maybe TLSes then can produce noticeable drift in frequency. These dominate at relatively low frequencies, below 1Hz, meaning that there is not a strict ordering between TLS and other sources of noise in terms of amplitude. In this thesis we propose a new design of KID that we hope can remove the vast majority of TLS noise.

2.3 Dynamics: TLS Noise and Resonators

The dependence of TLS and KIDs on temperature and power is interesting and subtle. Even when superconducting, the resonator does experience some resistance (according to the so-called "Two Fluid" model, which has a superconducting and normal current side-by-side), and so will heat up from energy.

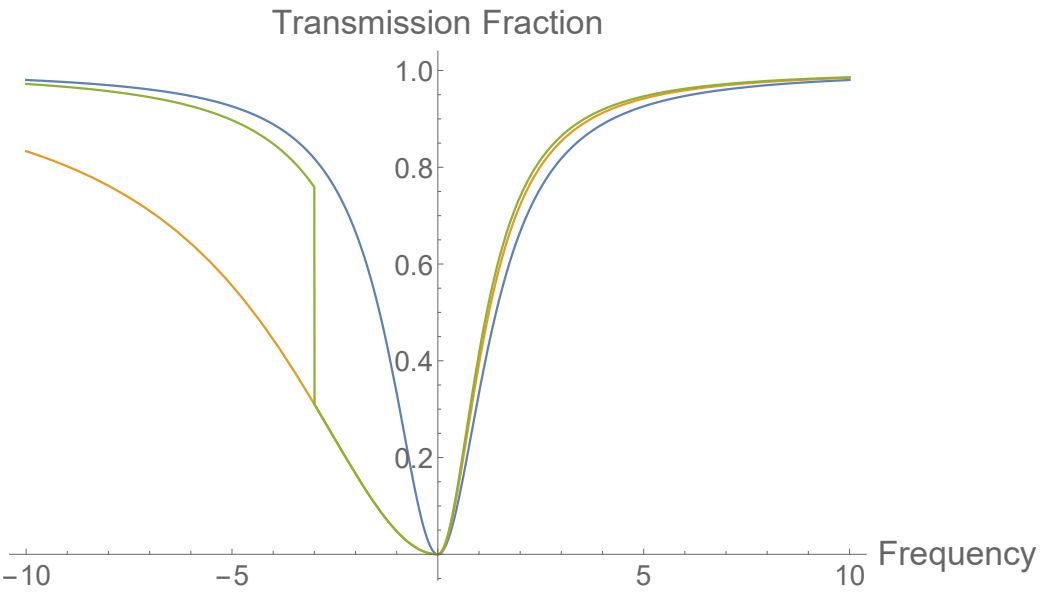

Figure 2.8: Demonstration of the KID hysteresis in an idealized resonator. Blue line is true response of the resonator, as would be observed with an infinitesimal beam power. Orange line is the observed spectrum if the beam is swept downwards in frequency. Green line is observed with a sweep upwards in frequency.

The energy absorbed is proportional to the excitation of the resonator, thus, only significant near resonance. This heating will in turn shift the resonant frequency, which can further modulate the energy absorbed. This means that sampling the resonator introduces a hysteretic term. See for instance Figure 2.8. During a sweep in frequency of the beam tone, the frequency of the resonator will change. When the beam tone is somewhat below the nominal resonance, there is a bistable configuration: in one state, the energy absorption is high, so the resonant frequency is dropped to match the tone, so the absorption is high. In the other state, the energy absorption is low, so the resonant frequency is near its nominal value and does not match the tone, so the energy absorption is low. As the frequency varies, this bistable state can merge or reach instability, giving the capacity for hysteretic cycles.

This hysteresis and nonlinearity means that there is substantial influence of how the frequency domain is searched for resonances. The five primary methods that could be used are

- Rising-frequency chirp: This produces the green line in Figure 2.8, with a discontinuity. The exact point of the discontinuity may be usable as a detectable feature.

- Falling-frequency chirp: This produces the orange line, and pushes the low end of the beak out. This gives the broadest range of frequencies over which to do a curve fit.

- White noise: This has the same uniform power spectrum has both chirps, but with random phases on each tone so that there is roughly equal excitation at all frequencies at all points in time. This may produce a slightly less clean response curve, as the resonance will be jittering back and forth during the tone, but suffers less from hysteresis and stability issues.

- Beam tones: Placing a test tone roughly half the way up the right side of the resonance curve can suffice to observe signals, as shifts in resonant frequency will move the beam tone response up and down in transmission, and change its phase. This is stable and only requires demodulation (as opposed to a full FFT), but provides less redundancy in the curve fit of the resonance.

- "Fast chirp": This approach consists of a brief, high-power blast of white noise down the line, followed by a period of silence. The resonators gradually ring down, and the produced signal can be decomposed to recover their resonant frequencies.

At extremely low temperatures, TLS excitations will not occur, and so there is no noise whatsoever. However, the ambient power of current in the wire will provide plenty of available energy (even if it's not fully thermalized energy), so that TLS will appear quite strongly. On the other hand, at higher temperatures, the TLS excitations can saturate to being completely random: as a result, the time variation in the number is low, and the TLS noise is diminished. As the TLS noise modulates the resonant frequency, it can also affect the energy absorption, in turn modulating the amount of TLS noise present. This can create a complex interaction where the resonant frequency and magnitude of TLS noise create even more types of stability; and both are affected by temperature in different ways.

In terms of the precise origin of the TLS noise on the circuit, there are mixed results. Recent studies have suggested that it primarily manifests

on the capacitors of the circuit, where it modulates the complex impedance (either capacitively or inductively, it is unclear). The evidence is unclear, because a set of unrelated experiments appeared to demonstrate that the magnitude of TLS depends only on the width of the wire in the inductor (but not the length).

Chapter 3

Dual Resonator Design

In order to overcome the TLS noise, we would like to consider the existence of a circuit that allows us to "factor out" the noise in the capacitor. If we model the TLS noise as stemming entirely from the capacitor, then we naturally have a problem that our one observable is directly influenced by the noise the same way that it is by the signal. We need to ask, then, if there are alternative circuit topologies that could help us become resilient to this noise in some way.

3.1 Design principles

In its current form, the resonator has an idealized circuit like the image on the left:

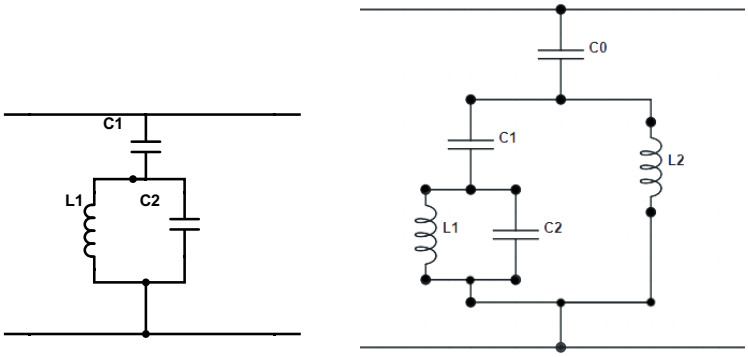

where the capacitor C1 is very small compared to C2, and there for the purpose coupling of the tank circuit to the transmission line. The resonant frequency is then $f_0 = 1/\sqrt{L_1 C_2}$. This is then just as dependent on the capacitance as the inductance. Since our signal is contained in the changing inductance, any changes in the capacitance will cause noise in the calculated inductance. As we add more components to the circuit in a non-trivial topology, we see new resonances emerge. The circuit on the right has 2 resonant frequencies, with 2 inductors and 2 capacitors (the minimum). Dimensionally, we know that the resonant frequencies must follow the same scaling relation that they will double if both inductances or both capacitances are reduced by a factor of 4. There are two other degrees of freedom in this circuit though, the dimensionless ratio of the inductances and the dimensionless ratio of the capacitances:

$$\ell_R = \frac{L_1}{L_2}, \quad c_R = \frac{C_1}{C_2}$$

The resonant frequencies both depend on these ratios as well as the overall scale.

If we assume that all the radiation lands on one of the inductors – say, L2 – then there are three unknowns: the inductance (which tells us the incoming radiation rate) and the capacitance of each of the two capacitors. If we choose our ratios correctly, though, we can make it so that a change in either capacitor "looks" the same, to first order. Put precisely, if we write the resonant frequencies ω_1, ω_2, we desire

$$\frac{\frac{d\omega_1}{dC_1}}{\frac{d\omega_2}{dC_1}} = \frac{\frac{d\omega_1}{dC_2}}{\frac{d\omega_2}{dC_2}}$$

With a CAS, we try to find variables to satisfy this property. Writing the impedance of the element gives us zeros at the factors of

$$1 - C_1 L_1 \omega^2 - C_2 L_1 \omega^2 - C_1 L_2 \omega^2 + C_1 C_2 L_1 L_2 \omega^4$$

After solving for the roots, computing the derivatives needed, and turning dimensionless, the constraint that we end up needing is

$$c_R = \frac{\ell_R}{1 + \ell_R}$$

As an example of some parameters, and how the capacitance noise could be avoided, we can take $\ell_R = 1, c_R = 1/2$. That is, the two inductors are equal in size, and C_1 is half the size of C_2. The two resonant frequencies are then $\sqrt{2 \pm \sqrt{2}}/\sqrt{C_2 L_2}$, or numerically

$$\omega_1 = \frac{0.765}{\sqrt{C_2 L_2}}, \quad \omega_2 = \frac{1.85}{\sqrt{C_2 L_2}}$$

If we compute the fractional derivatives $\frac{1}{\omega}\frac{d\omega}{dC_1}$, we get

$$\frac{1}{\omega_1}\frac{d\omega_1}{dC_1} = \frac{1}{\omega_2}\frac{d\omega_2}{dC_1} = -\frac{1}{4C_1}$$

$$\frac{1}{\omega_1}\frac{d\omega_1}{dC_2} = \frac{1}{\omega_2}\frac{d\omega_2}{dC_2} = -\frac{1}{4C_2}$$

The fractional derivatives with regards to a change in inductance are

$$\frac{1}{\omega_1}\frac{d\omega_1}{dL_2} = -\frac{0.427}{L_2}, \quad \frac{1}{\omega_2}\frac{d\omega_2}{dL_2} = -\frac{0.0732}{L_2}$$

Or in general, to first order,

$$\begin{bmatrix} \Delta\omega_1/\omega_1 \\ \Delta\omega_2/\omega_2 \end{bmatrix} = \begin{bmatrix} -1/4 & -1/4 \\ -0.427 & -0.0732 \end{bmatrix} \begin{bmatrix} \Delta C/C \\ \Delta L_2/L_2 \end{bmatrix}$$

where C can be either capacitor or both, it doesn't matter. Inverting this matrix will give us ΔL_2, from which we can recover the amount of light on the detector, hopefully without the TLS noise.

It is worth noting that if it turns out that TLS noise is a feature modulating inductance, instead of capacitance, then the above circuit topology will still work, because of the reciprocity of capacitors and inductors. Instead, the necessary constraint will be $\ell_R = c_R/(1 + c_R)$, and this would allow us to measure TLS noise on the inductors. However, this wouldn't offer any benefits for light detection, since motion of the inductance would be both our signal and noise and thus couldn't be separated. Building such a resonator would be an interesting test for the source of TLS noise, however.

From a practial standpoint of layout on the chip, the above suggest parameters would roughly double necessary space for circuitry, as an inductor of equal size and capacitor of half the size are being added. However, in some detector designs where the TLS noise is enough of a problem that it creates

the limiting factor, the capacitors would normally have to be manufactured larger in order to compensate; if the TLS noise can be cancelled out, this could allow the capacitors to be substantially smaller, perhaps even being almost as space efficient as before.

There is separately the issue of how to distribute the frequencies. With the examplle paramters given, the higher resonance is always at $1.85/0.765 = 2.41$ of the lower one. This is easy to distinguish certainly, but a high cost of in necessary bandwidth. In general, for a given value of ℓ_R, the ratio of frquencies is given by

$$\frac{\omega_2}{\omega_1} = \sqrt{1 + 2\ell_R + 2\sqrt{\ell_R + \ell_R^2}}$$

which rises with ℓ_R, and goes to 1 in the limit of very small ℓ_R. That is, if we are willing to make one inductor much larger than the other, then the frequencies can be placed closer to one another. For instance, with $\ell_R = 1/5$ and $c_R = 1/6$, the frequencies are only separated by a ratio of 1.365. The scale of that frequency is set by $1/\sqrt{C_1 L_0}$, so if we wanted this to have generally similar parameters to a given standard-topology resonator, we would be adding a second capacitor be 1/6th the size of the standard one, and a second inductor 5 times larger than the standard one. Since it's most space-economical to adjust the capacitors' and inductors' scale so that they take up similar quantities of space, the capacitors could be scaled up and the inductors scaled down by a factor of $\sqrt{5}$, so that utlimately this would take roughly 3 times the surface area of a standard detector – and again, requiring a bandwidth sufficient to see two frequencies 1.365 times apart. The dependencies become

$$\begin{bmatrix} \Delta\omega_1/\omega_1 \\ \Delta\omega_2/\omega_2 \end{bmatrix} = \begin{bmatrix} -1/4 & -1/4 \\ -1.7603 & -0.73969 \end{bmatrix} \begin{bmatrix} \Delta C/C \\ \Delta L_2/L_2 \end{bmatrix}$$

This was, again, all assuming that all light hits the L_2, and L_1 is dark. The computations are slightly different if this is not the case, as long as the fraction of light hitting each inductor is constant. For instance, if L_1 receives all the light, then the second row in the above matrix will be different. If the fractions of light are unsure at the time of construction, then a calibration measurement by observing behavior at any two levels of brightness (and integrating over a long time to average out TLS noise) should one to identify the common mode. Note that the two inductors will not affect the resonators

with proportional amounts (i.e. $\frac{\frac{d\omega_1}{dL_1}}{\frac{d\omega_2}{dL_1}} \neq \frac{\frac{d\omega_1}{dL_2}}{\frac{d\omega_2}{dL_2}}$), so that measuring the TLS noise will require knowing how light is hitting the two inductors. Although having them be proportional, this cannot be achieved while $c_R = \ell_R/(1 + \ell_R$ in this topology. It is not necessarily impossible in higher-order constructions, though, e.g. 3-inductor / 3-capacitor / 3-resonance.

3.2 Symmetric resonators

We are focusing on topologies of circuit, what we call the "symmetric double-resonator" and the "asymmetric double-resonator". The former is easier to reason about intuitively because we can consider the symmetric and anti-symmetric modes; the latter has fewer components and may end up being more favorable in terms of design constraints. For now we discuss the mathematical aspects of the symmetric design, and what component quantities would yield a feasible detection device with similar design parameters (size/sensitivity/resonator quality) as the traditional, single-resonator devices.

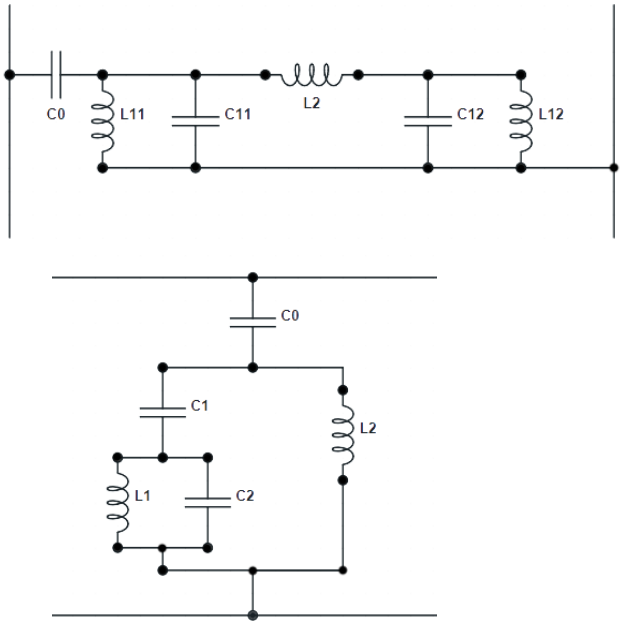

By computing the impedances of these devices on the line, and finding zeros of those impedances, we can identify their resonant frequencies. A computer algebra system aids this process. The designs have resonant frequencies, respectively, at the roots of

$$P_S(\omega) = L_{11} + L_{12} + L_2 - (C_{11}L_{11}L_{12} + C_{12}L_{11}L_{12} + C_{12}L_{12}L_2 + C_{11}L_{11}L_2)\omega^2 + C_{11}C_{12}L_{11}L_{12}L_2\omega^4$$

$$P_A(\omega) = 1 - C_1L_1\omega^2 - C_2L_1\omega^2 - C_1L_2\omega^2 + C_1C_2L_1L_2\omega^4$$

These frequencies are nominally affected by the coupling capacitor strength C_0, but it is typically several times larger in capacticance than the other C_i, and for the algebraic discussion here we can with good approximation here neglect it by taking the limit as $C_0 \to 0$. In simulations the term will be included, though.

We can model these detectors as system that take changes in the circuit parameters and produce observable resonant frequencies. For the purposes of this design, we are assuming that TLS noise stems entirely from capacitative noise on the capacitors; the signal (incoming light) will be focused onto inductor L_2 in either design. Thus we can think of a topology as defining a pair of functions function,

$$\omega_{S1}(C_{11}, C_{12}, L_2) \qquad \text{and} \qquad \omega_{S2}(C_{11}, C_{12}, L_2)$$

For detection however, we don't care about the global structure of the function as much as the local motion, the derivatives. As all our shifts are small fractions of the component's nominal state, we can write

$$\Delta\omega_{Si} = \frac{\partial\omega_{Si}}{\partial C_{11}} \cdot \Delta C_{11} + \frac{\partial\omega_{Si}}{\partial C_{12}} \cdot \Delta C_{12} + \frac{\partial\omega_{Si}}{\partial L_2} \cdot \Delta L_2$$

or equivalently

$$\begin{bmatrix}\Delta\omega_1 \\ \Delta\omega_2\end{bmatrix} = \begin{bmatrix}\partial\omega_1/\partial C_{11} & \partial\omega_1/\partial C_{12} & \partial\omega_1/\partial L_2 \\ \partial\omega_2/\partial C_{11} & \partial\omega_2/\partial C_{12} & \partial\omega_2/\partial L_2\end{bmatrix} \begin{bmatrix}\Delta C_{11} \\ \Delta C_{22} \\ \Delta L_2\end{bmatrix}$$

and from this system we would like to be able to solve ΔL_2 from the $\Delta\omega_i$. In the system above, we will refer to 3x2 matrix as the $\mathbf{A}$ matrix; it describes reponsitivity observables to hidden variables. We can solve for ΔL_2 only if the third column of $\mathbf{A}$ is linearly independent from the first two, which in turn implies that the first two must be columns must be proportional. So

our first design condition, to allow us to ignore capacitative TLS noise, is given by

$$\frac{\partial\omega_1/\partial C_{11}}{\partial\omega_1/\partial C_{12}} = \frac{\partial\omega_2/\partial C_{11}}{\partial\omega_2/\partial C_{12}} \tag{3.1}$$

For the symmetric design this gives the constraint (using computer algebra systems) that

$$(C_{11} + C_0)L_{11}(L_{12} + L_2) = C_{12}L_{12}(L_{11} + L_2) \tag{3.2}$$

For the asymmetric design the constraint reduces to $\frac{C_2 - C_1}{C_0 + C_1} = \frac{L_2}{L_1}$. The motivation between the symmetric design is to give it symmetry by setting $C_{11} = C_{12}$ and $L_{11} = L_{12}$; it is clear that this satisfies the constraint above when $C_0 = 0$. However, the presence C_0 breaks the symmetry slightly, such that in reality $C_{11} = C_{12} - C_0$. Under these substitutions

$$L_{11} = L_{12}, \quad C_{11} = C_{12} - C_0$$

all terms of C_0 cancel exactly from our expressions for resonant frequencies, since we have 'restored' symmetry by absorbing an appropriate amount of capacitance to the other side of the resonator. Now the resonant frequencies are given by

$$\omega_1 = \sqrt{\frac{1}{L_{12}C_{12}}}, \qquad \omega_2 = \sqrt{\frac{1}{L_{12}C_{12}} + \frac{2}{L_2C_{12}}} \tag{3.3}$$

This form is reassuring, as the single resonator has the same $1/\sqrt{LC}$ form. Note that ω_1 (the symmetric mode) actually ends up being entirely independent of our detection inductor L_2. We can think intuitively of our readout as taking the first frequency ω_1, knowing that all shifts in ω_1 are due to capacitative noise; using to compute the capacitative noise; and subtracting that term on ω_2 to get a value for ΔL_2.

Just as we cannot recover ΔL_2 at all if $\boldsymbol{A}$ has a degeneracy with its third column, we will have a *difficult* time recovering ΔL_2 if $\boldsymbol{A}$ has a near-degeneracy. Looking at this quantitatively, we would like to know the sensitivity and noise level on ΔL_2 given our observables ω_1, ω_2. There will be other noise terms $\vec{n}$ that contribute to the variation in frequencies $\vec{\omega}$, and the components of $\vec{n}$ have some covariance matrix $\mathbf{N}$. Then our system is modelled as

$$\vec{\Delta\omega} = \boldsymbol{A}\vec{s} + \vec{n}$$

where $\vec{s}$ comprises $[\Delta C_{11}, \Delta C_{12}, \Delta L_2]$. This is a linear system, specifically a Gauss-Markov model, and we can then build an optimal least-squares estimator as

$$\hat{s} = (\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{A}^T \mathbf{N}^{-1} \vec{\Delta\omega} \tag{3.4}$$

Since we don't need to actually recover the capactiance changes ΔC_{11} and ΔC_{12} separately, we can just drop all the terms and rows related to ΔC_{12} so that now $\hat{s}$ is only length 2 and $\boldsymbol{A}$ is 2x2. We can understand this as treating all noise as stemming from just one of the capacitors, and our estimator still holds.

Now the second row of $K = (\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{A}^T \mathbf{N}^{-1}$ describes how large of a change in inductance we need to observe a given change in frequency. We can build a similar model for the traditional single resonator case, and use these to compare sensitivities. By taking the reciprocal of the elements of the last row of K, we get our effective, degeneracy-adjusted sensitivities, as Hz of shift per Henries of inductance shift. Additionally, the element in the second row and second column of $V = (\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A})^{-1}$ tells us the variance in our estimated inductance. This will be proportional to the expected variance of the resonant frequencies due to other sources of noise, which is described in N.

Now that we can compute numerical values for how sensitive and noisy our detector will be, we can select a set of parameters and compare to a single-resonator detector. Since we plan to be testing with a Software Defined Radio device with only 100MHz of bandwidth, for this initial test we will need to select parameters that place both resonant frequencies well within 100MHz of each other. Intuitively, the detecting inductor L_2 is also what breaks the symmetry: by making its inductance very large, the mode in which current flows through L_2 starts to see that inductor increasingly as a ground. The symmetric mode, in which current doesn't flow through L_2, always sees it as a ground. Thus increasing L_2 causes the resonant frequencies to draw closer to one another. Increasing the capacitances causes all frequencies to scale down, by simple dimensional analysis, so increasing them will also force the resonant frequencies closer together. Lowering the $L_{11} = L_{12}$ inductance draws the resonant frequencies somewhat lower, but less significantly. Unfortunately, all of these approaches also decrease degeneracy-adjust sensitivity by comparable amounts for region of parameter space we checked. For applications where size on the chip is more of a limitation than the bandwidth, we would recommend to make the L_2 inductance comparably small in order

to increase sensitivity and reduce size of chip.

The following table contains the parameter values we selected, the resulting frequencies, the degeneracty-adjusted sensitivity $1/K_2$, and the variance in estimated inductance relative to noise in resonant frequency V_{22}; for a comparable dual-resonator and single-resonator model.

	C_0	C_1	L_1	L_2	ω_1	ω_2	$1/K_2$ in MHz/H	$\sqrt{V_{22}}$ in H/Hz
Dual	0.1pF	40pF	3nH	18nH	459MHz	530MHz	(58.4, -66.3)	2.28×10^{-8}
Single	0.1pF	12pF	7nH	-	549MHz	-	-275	3.67×10^{-9}

This suggests that we will have roughly 3 times the chip area (as the capacitors and inductors are roughly 3 times the size), while lowering our sensitivity by approximately 4dB and raising our sensitivity to other noise by approximately 8dB.

Chapter 4

Methods

4.1 Design simulation

With the desired component values determined, appropriate dimensions and layout of such components needs to be computed. All our inductance and capacitance is essentially "geometric" in the sense that it can be produced by appropriate layout of 2D metalization, as opposed other attached components.

The Sonnet simulation package provides faculties for laying out metalization and computing its frequency response patterns. Our capacitors were unfortunately going to be too large to simulate directly, which meant appropriately extrapolating based on smaller regions of the capacitor pattern. Before simulation could begin, we would need to select a material: at this point, we chose niobium. Niobium was desirable because it has a relatively large band gap, which would suppress generation-recombination noise.

4.1.1 Capacitor simulations

All our capacitors had a finger width of 2μm, inter-finger spacing of 2μm, a terminal width of 3μm, and an end gap of 2μm. The gaps were more conservative (larger) than the gaps in the inductor, because any short in the capacitor would lead to total failure of that entire double-resonator. With those parameters determined, the dimensions to adjust to get an appropriate capacitance were the number of finger pairs (FP), and μm finger overlap (FO). These each adjusted dimensions of the capacitor along one axis. Our test data from Sonnet was as follows:

FP	FO (μm)	C (pF)
20	42.0	0.115
30	63.0	0.238
40	84.0	0.395
60	126.0	0.864
90	189.0	1.925

Where the test dimensions were chosen to give a fixed FO/FP ratio of 2.1μm, so that we only had to fit along one dimension. The data fit well to a quadratic ($\pm$ 0.004pF max error), given by

$$C = 0.0216704 + 0.0000532499x^2$$

where C is capacitance in pF and x is the Finger Overlap in μm. We used this to arrive at parameters for our desired components, where first took the nearest good FP (constrained to be an integer), and then the FO that gave the product closest to the desired (real) value. Our design parameters were then

C (pF)	FP	FO (μm)
0.1	19	42.0
0.3	19	126.0
12	229	480.0
40	417	876.0

4.1.2 Inductor simulations

A similar set of simulations and design work happened for the inductors. We kept fixed a conductor width and conductor spacing 1μm and a meander length of 60μm. This then allowed the number of legs (NL) to determine the inductance. Our test simulations were

NL	L (nH)
20	0.657
50	1.585
100	3.131
150	4.697
200	6.257

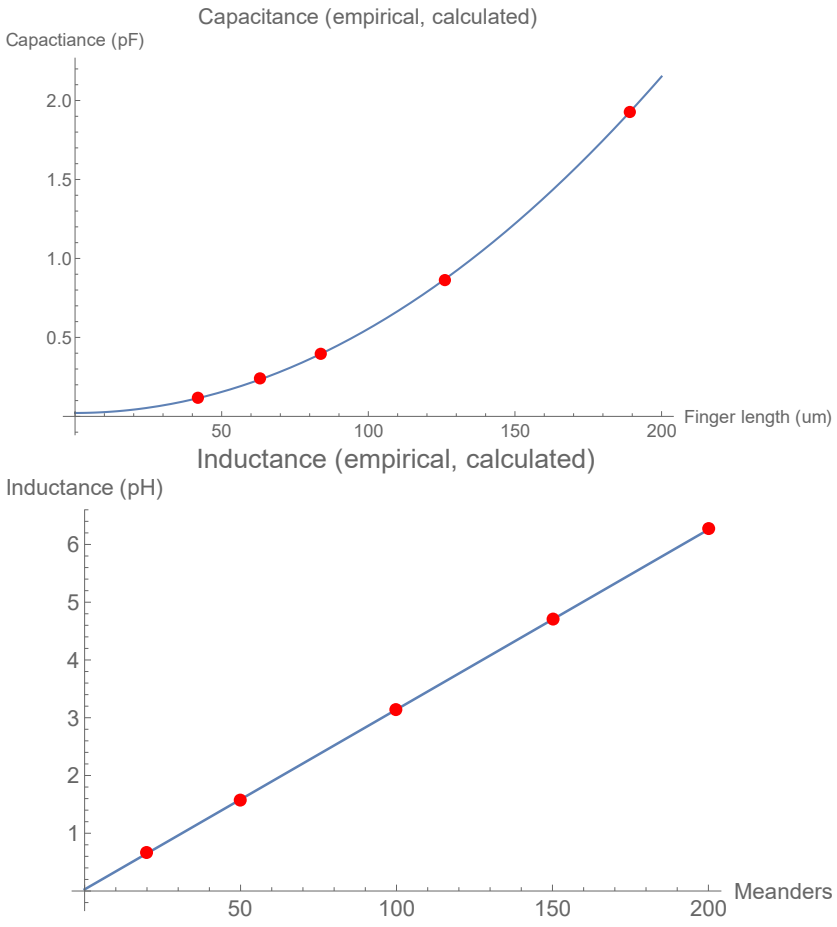

Figure 4.1: Inductor and capacitor model fits to simulation data

These were fit with a linear fit (maximum error 0.01nH). The fit was

$$L = 0.0308537 + 0.0310829x$$

where L is inductance in nH and x is the number of legs. This led to design parameters of

L (nH)	NL
3	95
3.1	99
3.2	103
7	224
18	577

4.1.3 Layout

In terms of arranging our components correctly, we had to take into account how they would couple to the ground plane behind, and how they would couple to each other. After an initial "first pass" at layout, we calculated coupling values to check for safety.

For modelling the interactions between two capacitors of the same resonator, we treated them as two rectangular plates, of width 886μm, separation 471μm, and length 3330μm, in a medium of $\epsilon = 3.4$. This predicted a coulping of 0.19pF, which while not ideal, should keep anything from still functioning correctly.

Similar figures were used to model the interaction of resonators across dual-resonator pairs. Flipping with width and length variables, and setting the separation to 700μm, led to a coupling capacitance of 0.06pF. This was sufficiently small that we felt we could call the dual resonators isolated.

The capacitors couple to the ground plane as a parallel plate. The material between them (silicon) has $\epsilon = 12$, and we model the capacitor as a full parallel plate (ignoring the 50% filling fraction) because the voltage between the fingers should remain roughly constant. The capacitor is 0.5mm away from the ground plane. This leads to a coupling capacitance of 0.635pF. As a metric of whether this will noticeably degrade the performance of the resonator, we can compute the coupling Q:

$$Q_C = \frac{C}{\pi\omega C_C^2 Z_0} = \frac{40\text{pF}}{\pi(500\text{MHz})(0.1\text{pF}^2)(50\Omega)} \approx 51000$$

and then see how this is affected by the coupling capacitance. The effective coupling capacitance C'_C, as determined by the normal C_C and parasitic C_P, is now

$$C'_C = \frac{C_C C_P}{C_C + 2C_P} = 0.046\text{pF}$$

which drops our Q more than we would like. In order to improve this situation, we will etch away the back of the chip to provide 2mm of air between the capacitor and the ground plane. This reduces ϵ from 12 to 1, and the coupling capacitance becomes safely negligible. This produces a parasitic coupling of only 0.015pF, and gives us a Q in the range of at least 30000. This was also what led to our decision of how large our coupling capacitor should actually be.

To estimate whether or not the geometric inductance of the layout will play a role, we want to calculate whether or not two resonator loops will couple to each other. We model each half-dual as a "right triangle" (in terms of area), with dimensions of 1200μm by 3000μm. Now approximate them as circles of that area, i.e. a radius of 840μm. The center-center distance of these triangles is 1400μm. Applying a formula from Jackson's Electrodynamics (page 234), we get a mutual inductance

$$M_{12} = \frac{\mu_0 \pi r^4}{4d^3} = 0.2\text{nH}$$

which is very small compared to the inductances of any of the individual resonators, so we view this as a reasonable layout.

It is possible that the geometric inductance of the loops might affect the resonant frequencies somewhat, but it will do so geometrically: the symmetry of the design means any geometric inductance will shift the whole system down, but it will continue to function just fine in terms of TLS cancellation, as long as each half of the dual-resonator is affected equally.

4.2 GPU readout

With the cryostat set up and cooled, there is the question of transmitting and processing signals efficiently. In the past, our group has used Field Programmable Gate Arrays (FPGAs) for this task, however these are expensive and relatively difficult to configure. Our group has instead started to use off-the-shelf Software Defined Radio (SDR) hardware for our processing. An

Ettus Research USRP suffices for transmitting an arbitrary waveform down the transmission line, and receiving the signal back. This is then relayed back to a computer for processing, via 10GB Ethernet. The processing requires requires large, real-time FFTs, however, and so is not easily done on the computer.

For this purpose, I wrote prototype software last summer for communicating with the USRP, receiving the data, and moving immediately to a GPU. The GPU is then programmed in Nvidia's CUDA language, where the FFT is carried out, resonant peaks are identified, and then this data is then moved back to the CPU for logging. This software has now gone through a few iterations, and we hope that it can be released as a general library for Kinetic Inductance and Quantum Capacitance readout systems.

Chapter 5

Measurements

With the mask designed and ordered, we ended up making two groups of chips, where the second was attempt to make TLS noise more visible over the first.

5.1 First chip

In the first batch of chips, three chips were produced. These three were put under a microscope to search for defects – shorts in the capacitors, primarily – and exactly one was defect free. See images in Figure 5.2 and Figure 5.3. Due to a slight miscommunication, they were stripped of oxide during manufacturing, which may have reduced TLS. The resonant frequencies were found to be, in pairs,

251.6MHz	281.7MHz
252.3MHz	282.6MHz
253.1MHz	284.3MHz

These were substantially lower than expected. We believe this can be attributed to geometric inductance in the circuit design, raising the effective resonant mass. The small loop, within one resonator, has dimensions approximately 3350μm by 100μm. Modelling the structure as a rectangular loop with a wire width of 1um and relative permeability of 1, textbook formulas would predict geometric inductance of 7nH. Compared with the intended 3nH of kinetic inductance, this is enough to make a roughly 2-factor frequency drop, and so is the order of magnitude to explain the discrepancy.

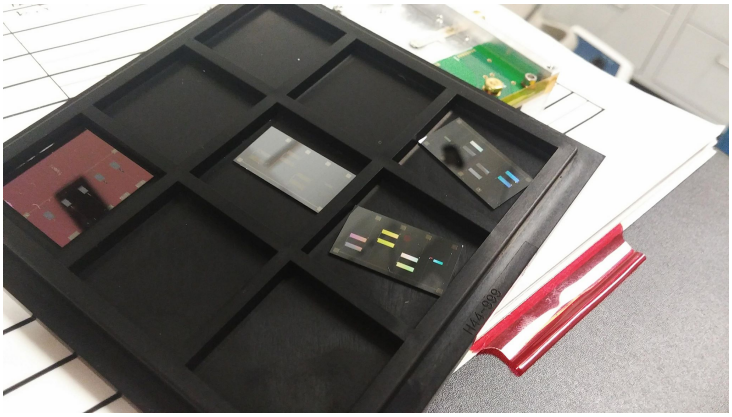

Figure 5.1: First batch of chips

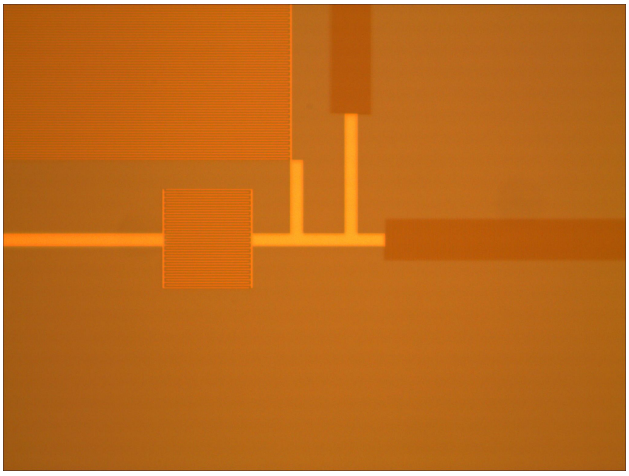

Figure 5.2: Image of a few of the connections in a detector. To the bottom is the line joining the two resonators, pulling off to the right. Vertically in the middle is the inductor, and off to the left is the capacitor. The small box in the middle is the coupling capacitor, attaching the lumped element to the transmission line.

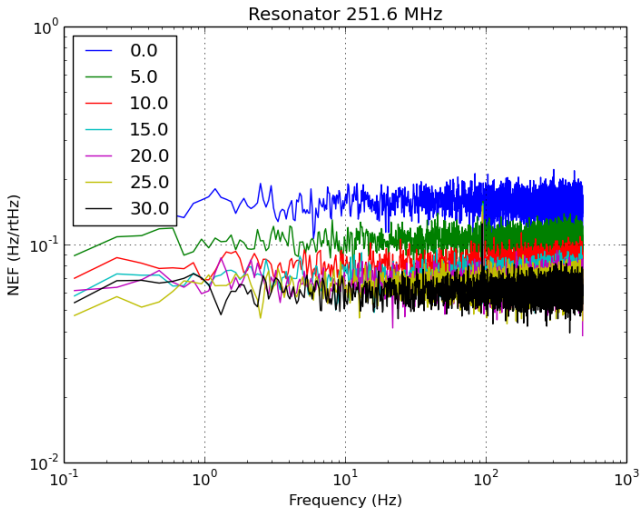

Figure 5.4: Effective noise power at various transmission powers, for the 251MHz resonance.

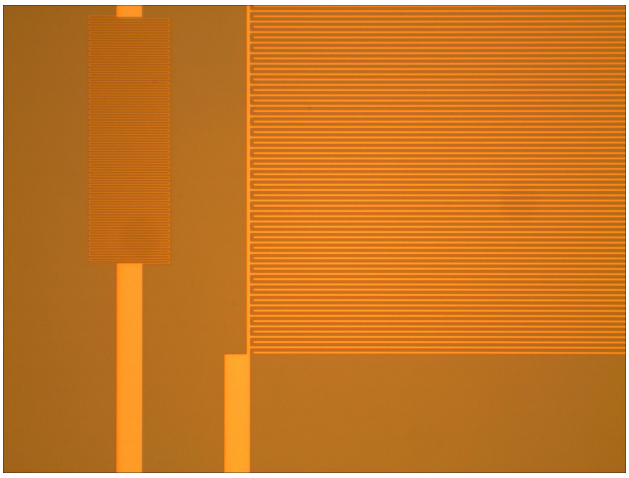

Figure 5.3: Closer zoom on another joint. The interdigitation and meander are directly visible.

Once we had characterized where the resonant frequencies were, we took reference noise measurements at each frequency. As a particular example, one spectrum can be seen at Figure 5.4; the others are collected in Figure 5.5. The transmission power dependence is as expected: at low transmission powers

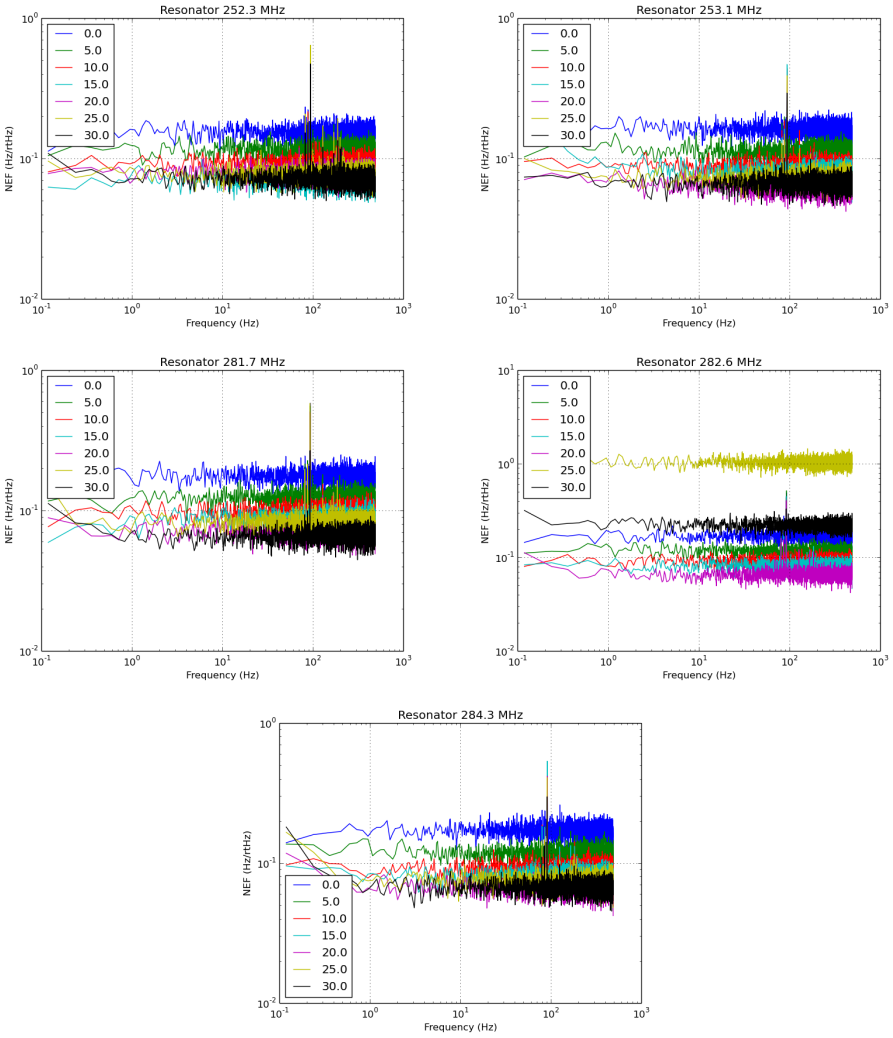

Figure 5.5: Effective noise power for other resonators

(0.0dbm down the line), the Noise Effective Frequency is high. Here, we are dominated by uncertainty coming from the readout itself, the noise in the transmission line. As we move higher towards 20.0dbm, we hit a "noise floor" where our transmission power is adequately probing the device. As can be seen in plots 2 and 5 of Figure 5.5, at sufficiently high powers the noise just barely begins to creep up again, as we are injecting enough power to excite the device. This increases temperature and GR noise.

These individual noise spectra are very flat, which indicates something concerning: the noise is almost entirely white noise. We would ideally have seen a sloping upwards in the left, as TLS noise goes as $1/f^{1/2}$ in power ($1/f^{1/4}$ in amplitude), and so should gradually begin to dominate at low frequencies. The real test for this chip comes in looking at the cross-correlation in the noise between two resonators in a pair. The plots are visible at Figure 5.7. None of these correlation coefficients is notably away from 0, and indeed all the correlation that is observed at low frequencies is within random variation. There is also the risk of contamination from gradual temperature drift over time, which would shift resonances in all resonators similarly, and dominates at low frequencies.

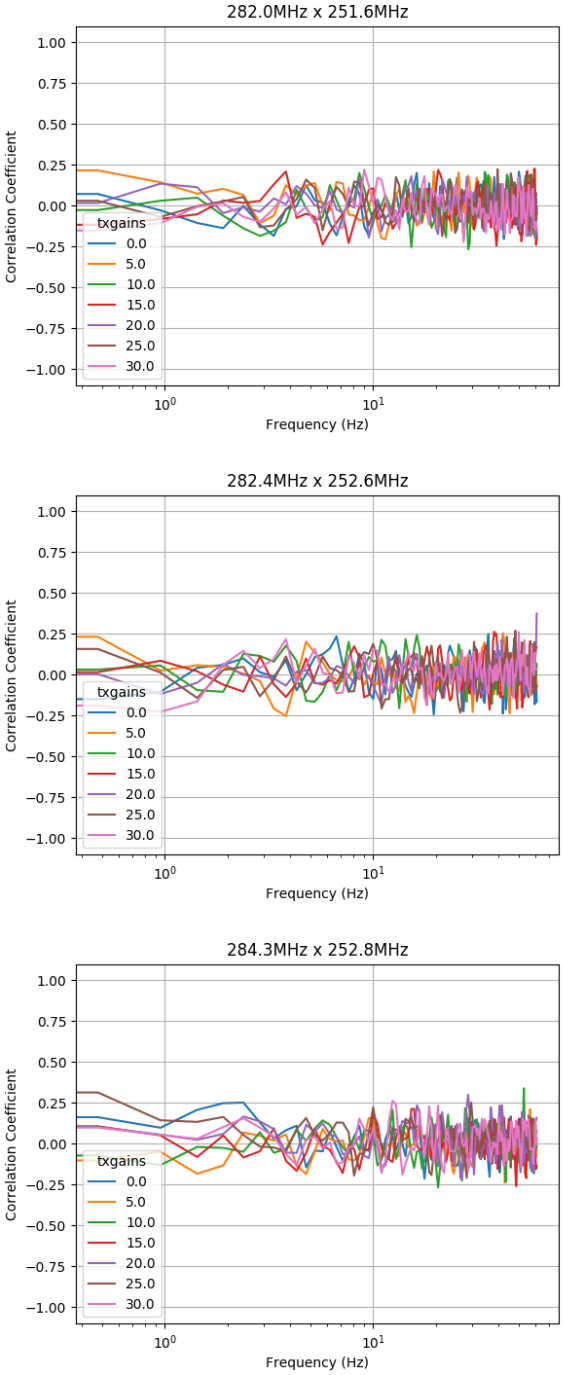

Figure 5.6: Cross correlation between coupled pairs of resonators.

Finally, we noticed that the resonators had an unusually high quality factor, and a very low temperature dependence in their quality factor:

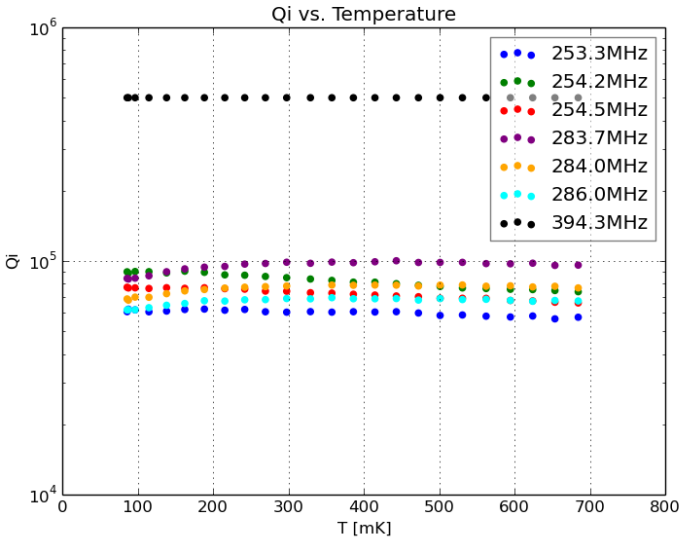

Figure 5.7: Temperature dependece of Q's in first chip

We expected much more, if TLS were to be playing a significant role. When it became apparent that these devices were not displaying enough TLS noise that we could find, our group decided to make another batch of chips, with a layer of silicon oxide placed over. The intent was that this extra oxide would couple more two-level systems to the wire in order to create more TLS noise.

5.2 Oxidized chips

The second batch of chips was made. One chip was discarded due to a large short. In order to reduce problems with "overlap" between the resonances, which may or may not be far off from intended, two of the dual resonators were cut with a scalpel. This way we could just focus on one. The resonant frequencies of the dual resonators were 319.6MHz and 358.8MHz. There was also still a single resonator, as reference, at 494.9MHz.

The noise spectrum showed that were not hitting the noise floor (see Figure 5.8). So, 20dB of attenuation was removed. This did leave some

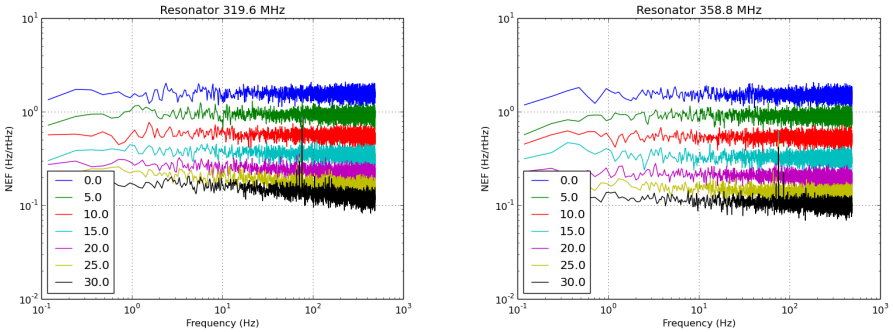

Figure 5.8: The oxidized chip: Effective noise power at various transmission powers.

curvature in the noise spectrum, see Figure 5.9, which we hoped was TLS now.

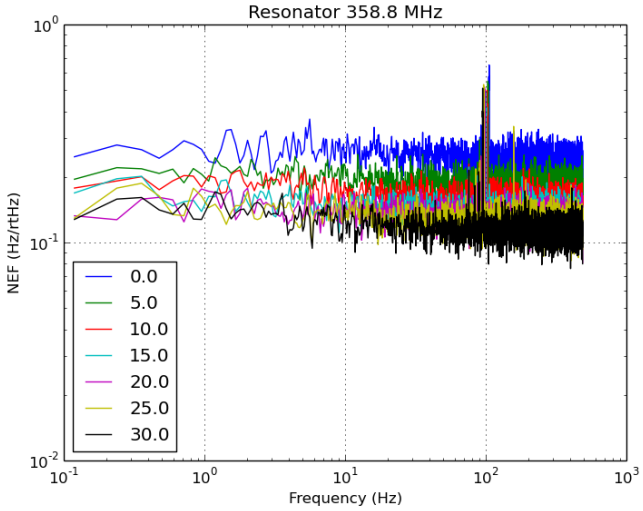

Figure 5.9: The oxidized chip: Noise floors with 20dB less attenuation.

To characterize whether or not this was TLS, we did a network analysis of the resonance at a variety of temperatures and tone powers. The dependence of the resonant frequency and Q (quality factor) would then be a good indicator of the dynamics.

5.3 Comparison with expected TLS

As (Zmuidzinas, 2012) explains, the dependence of TLS noise on temperature and resonant frequency can be modelled with a loss tangent,

$$\delta_{TLS}(\omega, T) = \delta_0 \tanh\left(\frac{\hbar\omega}{2k_B T}\right)$$

where δ_0 is "proportional to the density of TLS per unit volume and energy", and $\hbar\omega/k_B T$ becomes a dimensionless parameter describing the ratio of thermal energy to resonant energy. The resonant frequency shift is then described by

$$\frac{\delta\omega}{\omega} = \frac{F_{TLS}\delta_0}{\pi}\left[\Re e\Psi\left(\frac{1}{2} + \frac{1}{2\pi i}\frac{\hbar\omega}{kT}\right) - \ln\frac{\hbar\omega}{kT}\right]$$

Where $\Re e\Psi$ is the real part of the complex digamma function, and F_{TLS} is a filling fraction of the TLS energy (expected to be order unity).

Zmuidzinas 2012 also includes a chart of empirical data for TLS levels in various types of resonators, specifically Fig 14 on p192. refers to 200nm niobium traces on silicon, probed at 5.1GHz and 120mK. Our probes are considerably lower in frequency (more than an order of magnitude) and the width is unspecified, but if TLS is detected, trying to extend this data would provide useful confirmation.

Chapter 6

Conclusions and Future Work

We do not very well understand why our chips failed to display TLS, or at least failed to display correlated TLS that we could detect. It is additionally surprising that adding the additional SiO coat did not seem to create much more TLS, and that in fact the quality of the resonators noticeably increased when the oxide layer was applied. Although we were initially hopeful that we would see strong correlation in the TLS noise basically as soon as put the chip in, this did not happen, and it is not clear why. Our group will continue pursuing this problem through more careful characterization of the noise of each chip. Our group had not previously made any all-niobium chips on silicon – when we had niobium chips, it was typically a niobium capacitor and aluminum inductor, for instance, as niobium's high band gap makes it somewhat harder to excite with incoming light. It is possible that TLS is some way dependent on aluminum oxide and other oxides, in a way that does not manifest nearly as strongly with niobium and silicon oxide.

To respond to our particular question, whether this design of detector allows us to cancel out TLS, we can only say inconclusive, as we could not conclusively identify TLS noise to cancel out, yet. That the resonators should be coupled in terms of any noise on their capacitors is almost certainly still going to hold, but it is possible that the design ultimately fails, if the TLS noise manifests primarily on the sensing inductor.

To specifically outline what we would like to do in the future, we need to more carefully characterize the chips we have, both oxidized and not. We did see correlated low-frequency noise at times, but we are not sure to what degree this was temperature drift. To be satisfied with our understanding of these chips, we will need to look at a noise spectrum at each of a variety

of temperatures and tone powers, where we more carefully have controlled temperature. This will mostly mean running it for a while at a given tone power, allowing it to "warm up" to the tone. We are also considering doing single-tone measurements, which inject far less power into the system, and keep the temperature more stable over the course of a reading. Single-tone readings will allow us to track resonant frequency shifts. This will require having a calibration of the Q's of each resonator at these temperatures as well. If we are unable to observe TLS effects, we will investigate fabricating similar designs out of other metals for comparison.

References

1. Jonas Zmuidzinas. *Superconducting Microresonators: Physics and Applications*. Annual Review of Condensed Matter Physics. Vol. 3:169-214. (2012).