Analysis on Coplanar Waveguide Transmisison

Yutong Du
Physics Honors Thesis
University of California, Berkeley

May 20, 2025

Thesis Adviser Signature: ______________________ Date: ______________________

Contents

Chapter 1

Introduction

Here, we introduce the origins of the theory of dark matter. For this section, most of the history is sourced from [1].

1.1 History

To begin, we should ask ourselves the question: what is dark matter? The very origins of dark matter stem all the way back to Lord Kelvin, who came up with a way to calculate the velocity distribution of stars in the Milky Way galaxy. In this calculation, Kelvin noted that the number of stars that could reasonably exist in our Milky Way is around 10^9, but since the observed number of stars is far less than that, he concluded that many of the stars in our galaxy must be "dark bodies". This then prompted Henri Poincare to coin the term "dark matter". At the time, it was believed by notable astronomers that dark matter was composed mainly of faint stars and other "meteoric matter" – matter composed primarily of bodies like asteroids that don't radiate light, and are too faint to detect with instruments at the time. Moreover, people generally believed that though dark matter exists, it was less prevalent than ordinary (baryonic) matter.

In the 20th century, this hypothesis was about to change. In 1933, Fritz Zwicky published the first modern interpretation of dark matter: while studying the velocity distribution of galaxies in the Coma cluster, he found a velocity dispersion (effectively a measure of the "spread" of the distribution) of 1000 km/s, far greater than the expected 80 km/s that would consistent assuming the cluster only contained visible matter. This led to the now widely-accepted conclusion: that dark matter exists in "much greater quantity than luminous matter" [2]. Further, Zwicky also concluded the same paper that based on the observed speeds and the amount of visible matter, galaxy clusters like the Coma system should "fly apart", implying that the amount of visible matter present is not nearly enough to support angular velocities as observed.

At the time, people had reason to doubt Zwicky's conclusion – the conclusion was extremely counter-intuitive, as it suggested that a large majority of a galaxy's mass (and consequently the universe's) is comprised mostly of matter that is effectively invisible to us. People obviously thought this was unlikely, so they looked for alternative explanations for why this may have occurred. One interesting theory was proposed by Erik Holmberg, who suggested that the galaxies used in Zwicky's calculation were not permanent cluster members but instead galaxies on hyperbolic orbits that just happened to come across one another at the time of observation. This hypothesis was later ruled out by Martin Swarzchild, who redid the calculation whilst being careful about which galaxies to include, and still arrived at an abnormally high velocity dispersion.

This is just one example of the many ways astronomers sought to discredit the dark matter theory, but to no avail as every opposing hypothesis was eventually discredited. Eventually, by the late 1960s and 1970s, astronomers had no choice but to accept that dark matter may in fact be a real phenomenon, and began investigating what it may be composed of. Early on, Zwicky hypothesized that it could be made up of "cool and cold stars, macroscopic and microscopic solid bodies, and gases". However, this turned out to be false as Herbert Rood discovered in 1974 that the "missing mass" had to come from the space between galaxies, meaning that it could not be coming from faint stars or rocky bodies. Interstellar gas was also proposed since it fits the criterion of existing between galaxies, however subsequent X-ray analyses showed that the amount of intergalactic gas contributed only 2% of what is necessary to fully account for the required mass.

Slowly but surely, more and more "ordinary" forms of matter were ruled out as sources of dark matter. Theories which pointed to planets, dwarf stars, and other forms of less luminous matter were ultimately ruled out as subsequent observations to assess their prevalence in our own galaxy failed to produce the quantities necessary to fully account for the missing mass. Furthermore, by the late 1990s and early 2000s, calculations using the cosmic microwave background showed that the majority of the universe's mass is non-baryonic, effectively forcing dark matter to be in the form of some subatomic particle.

And that concludes our overview of the history of dark matter to the present day. Most scientists agree that dark matter is probably made of some currently undiscovered subatomic particle. Generally, people refer to these particles as Weakly Interacting Massive Particles (WIMPs); they are the leading theory for two primary reasons: first, the quantity of WIMPs that would have been produced in the early universe seems to be consistent with the estimated amount of dark matter in the universe today (around 83%). Second, most theories beyond the standard model – which, importantly, were not devised in the interest of explaining dark matter – includes the introduction of some sort of new particle, which automatically becomes a candidate for a WIMP.

1.2 Alternatives to Dark Matter

This section is not very crucial to the focus of this thesis, but it is worthwhile to note that even though *most* astronomers agree that dark matter is a real phenomenon, there are still proponents who strongly advocate for alternative theories to dark matter. One such known example is called MOdified Newtonian Dynamics (MOND), and has gained traction in recent years as the leading alternative to the theory of dark matter. Pioneered by Moredehai Milgrom in 1983, MOND stands out as a theory from the rest as it is one of the only alternative theories that manages to explain the rotational anoamalies in galaxies as well as make some correct predictions. For instance, MOND was able predict the dynamics of low-mass spiral galaxies [3], and also account for the Tully-Fisher formula which relates the rotation curve of a galaxy to its intrinsic luminosity. Beyond this though, MOND has been largely unsuccessful in conforming to observations, which is why it still remains a "fringe" theory that is not widely accepted.

Chapter 2

Dark Matter Detection

2.1 Overview

In the previous chapter, we've established that the overall scientific community believes that dark matter probably consists of some particle. Here, we will discuss some of the current experiments and initiatives launched by the scientific community to determine the identity of dark matter.

As mentioned in the previous section, most scientists believe that dark matter exists in the form of WIMPs. Because WIMPs are a stand-in for a very large class of particles, they also occupy a very large range of masses, from as small as $1\,\mathrm{GeV/c^2}$ all the way to $10^5\,\mathrm{GeV/c^2}$. This poses somewhat of a problem when it comes to designing an experiment to detect such particles, as there is no single detector that's going to be able to cover this entire range of energies. Instead, what researchers do is focus on a particular range of masses, and try to detect particles within that range. To this end, there are three primary ways that experiments go about detecting WIMPs: either we generate them at particle accelerators (like the LHC), we detect them directly from astronomical sources like the sun and other galaxies high in dark matter, or we detect their scattering interactions in low-background experiments. This last method is the one that is of interest in this thesis, and it is the one that we will focus on for the rest of the discussion.

2.2 Collision based DM Detection

Collision based detection refers to the process of designing an experiment where we deliberately prepare an experimental sample, with the objective of essentially waiting until a dark matter particle interacts with the sample. This will cause a brief but detectable signal, which is then picked up using various instruments. The actual physics behind the detectors is rather complicated, so we will only dive deeper into two particular instruments (see Section 2.3 and Section 2.4). For now though, we will focus on giving a general overview of how these collision based detection methods help us confirm and find the identity of dark matter. In this section, much of the content is derived from [4].

To begin, we use the following formula to calculate the expected rate that a WIMP scatters off a nucleus with mass m_N:

$$\frac{\mathrm{d}R}{\mathrm{d}E_{\mathrm{nr}}} = \frac{\rho_0 M}{m_n m_\chi} \int_{v_{\mathrm{min}}}^{v_{\mathrm{esc}}} v f(v) \frac{\mathrm{d}\sigma}{\mathrm{d}E_{\mathrm{nr}}}\, dv \tag{2.1}$$

Here, ρ_0 is the local dark matter density, M is the target mass of the detector; m_χ is the mass of the WIMP; σ is the scattering cross section and $f(v)$ is the velocity distribution of the WIMPs; E_{nr} is the "nuclear recoil" energy –

essentially the energy imparted to the target nucleus by the WIMP. Essentially this is just the same formula as the standard form for expectation:

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} xP(x)\; dx$$

where $P(x)$ represents the probability distribution we are interested in. v_{min} is calculated using the following formula:

$$v_{\text{min}} = \sqrt{\frac{E_{\text{nr}} m_N}{2} \frac{1}{\mu^2}} \tag{2.2}$$

Here μ represents the reduced mass of the WIMP and nucleus. The physics of how Eq. (2.1) and Eq. (2.2) are derived is too complex for this discussion; the objective of including this formula is to make the point that v_{min} is nonzero, so there is a minimum threshold energy that WIMPs must have for them to be energetic enough to be picked up by our detector. The more interesting thing about Eq. (2.1) is that there is a maximum velocity v_{esc} that we are concerned with. The reason for this is because while our detector will be able to detect WIMPs at velocities higher than this, their velocity being above the escape velocity of the Milky Way means that they are not dark matter particles that are bound by the Milky Way, and thus we are not interested in those detections.

With the equations in place, we turn to one of the central equations we will concern ourselves with:

$$N = T \int_{E_{\text{low}}}^{E_{\text{high}}} \epsilon(E_{\text{nr}}) \frac{\mathrm{d}R}{\mathrm{d}E_{\text{nr}}}\; dE_{\text{nr}}$$

ϵ is the efficiency of the detector at the energy. This equation is actually rather simple, as it is just the product of T and the rate of detection, which is given by the integral. It is weighted by the detector efficiency, since physical detectors are imperfect. This equation is extremely important to understanding the detector, since N is the quantity we ultimately measure in our experiment. What this equation really tells us is that in order to understand N, we need to have a very good understanding of our experimental setup, since $\epsilon(E_{\text{nr}})$ is crucial to determining N. Obviously, $\epsilon(E_{\text{nr}})$ will differ from detector-to-detector, which is why there are so many groups tackling this problem with unique approaches – different detectors operate in different E_{nr} regimes, so while the full formula for N is what was written, the actual value of N also differs from this due to these experimental limitations.

One other complication that arises from calculating N is that the quantity $\frac{\mathrm{d}R}{\mathrm{d}E_{\text{nr}}}$ is not constant over time. This is primarily due to the fact that the Earth orbits around the sun; and since the sun orbits around the Milky Way, this means that for half the year the Earth is moving with the direction of the solar system through the Milky Way and against it for the other half. Ultimately this means that the relative velocity of WIMPs will vary throughout the year, in the fashion of a modulated signal:

$$S(t) = B + S_0 + S_m \cos\left[\omega(t - t_0)\right] \tag{2.3}$$

where $S(t)$ denotes the number of signal events at any time of the year t. This modulation is rather small (on the order of 5%), but it is nevertheless non-negligible. There are also more intricate factors like finer resolution of the dark matter distribution which should also be taken to account when we calculate $S(t)$ over the long-term (1+ years), but we won't enter that discussion as it detracts from the main focus of this thesis.

The objective of this section is to highlight the process by which dark matter can be detected using a collision based detector: first, we design an experiment to test for a WIMP with a very specific energy – this will correspond to a predicted particle of some sort based on a theory beyond the standard model. Then, we run the detector and record the signal $S(t)$ it produces over the course of a year, and compare its signal to its theoretical $S(t)$. If $S(t) \approx B$, then we know that the only thing the detector picks up is background, hence the proposed particle either does not exist or does not contribute to dark matter as we expect. On the other hand, if we do indeed see a modulated signal of the form in Eq. (2.3), then we know that what we've detected, at the very minimum, one particle that likely contributes to dark matter.[1]

Now that we've gone over the basics of how we can use a detector-based experiment to examine for the evidence of dark matter, we now focus on one particular approach, which uses cryogenic fridges to develop extremely sensitive devices that detect when a collision on the detector has occurred. First, we will focus on one of the primary technologies used today, the Transition Edge Sensor (TES).

2.3 Current Technology: TES Detectors

In this section, we will go over the basic workings of the Transition Edge Sensor (TES) and how they are generally fabricated into extremely sensitive detection devices. As a cryogenic device, TESs rely on the effect of superconductivity to achieve their high sensitivity, and as this is the case, it will be fully impossible to flesh out how TESs work on a fundamental level here. For further information on how they work, we refer to [5] for a more in-depth description.

As mentioned, TESs leverage the phenomenon of superconductivity to make these collision detections. To begin explaining the mechanism, it will help to first understand how superconductivity works: at a very high level, some conducting materials have the unique property that when they are cooled below a critical temperature $T < T_c$, their electrical resistance drops rapidly and approaches zero. In short, the reason this occurs is because below T_c, some electrons transition into a superfluid state, allowing them to move freely in the conductor with no resistance, which manifests itself macroscopically as a conductor with a resistivity of zero. The specifics of why this superfluid state occurs and how zero resistivity is achieved is too complex for this thesis, for now all that needs to be known is that the theory predicts that below T_c the electrons are said to "pair up" into *Cooper pairs*, which allows them to behave differently than ordinary free electrons.

Now, consider a system where a metal is cooled to T_c. Because the electrical resistance is a continuous and differentiable function, the resistance does not drop immediately to zero, but instead rapidly descends to zero, dropping as much as several ohms in within a range of a few millikelvins. This steep drop is specifically what TESs leverage, as described in Fig. 2.1. The TES usually sits on top of a substrate material (e.g. silicon). When a particle collides with the substrate, it emits phonons in the substrate, which are then absorbed by the superconducting portion of the TES. This causes the metal to heat up slightly, and because its temperature is around T_c, this raises its resistance dramatically.[2] If we are to connect the superconducting metal to a current, then we would see a detectable drop in the current through the circuit as the superconductor resistance increases, which constitutes our signal. After the event, the superconductor eventually radiates the heat away and returns back to the critical temperature T_c.

Now that we understand how TESs work on a basic level, it is important to discuss how they are realized in a detector. Specifically, how do we *select* the detection of a particular particle? That is, if we are interested in detecting WIMPs of a particular energy scale, how do we generate TESs that scan that specific range of energies? To answer this question, we turn to the *proximity effect*.

2.3.1 The Proximity Effect

The proximity effect refers to the technique through which experimentalists create superconducting TESs with a selected T_c. As mentioned, the reason this is desirable is because each particle deposits a different amount of energy

---

[1] Obviously, we are glossing over factors like device sensitivity but the overall logic is still sound.

[2] In reality, we cool the TES down to well below T_c then run a DC current to heat the TES up to T_c, because controlling the DC current is easier. In principle though, the effect is the same.

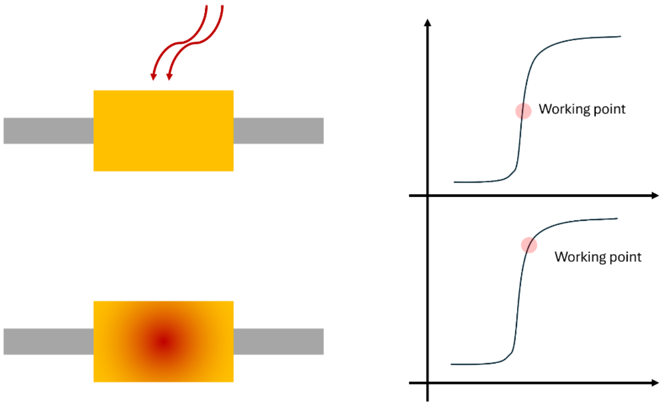

Figure 2.1: Simplified diagram of TES operation. When a particle collides with the material and emits a phonon, it is absorbed by the TES, which converts all the energy into heat. This raises the temperature of the TES, causing its resistance to increase sharply, which produces a detectable signal.

into the TES, so if we are to design a detector to search a particular range of energies, then we need to engineer our TES so that its thermal response is maximized (or at the very least detectable) when a particle of that energy collides with the substrate the material. The way this is achieved is through the *proximity effect*.

Without going too deep, the idea behind this mechanism is to combine two superconducting metals together by layering one on top of the other. If the layers are made thin enough, then the wavefunction that governs the electrons in the superconducting material will be able to "penetrate" the normal conductor. The overall effect is that we create a superconductor with a T_c that is between the T_c of the initial metals themselves, with its precise value controlled by the ratio of the two metasl used.

The exact equations for how this is done is not very important to the central focus of this thesis. For more information about TESs in general, see [5] for more information. For now, what matters here is understanding how TESs work, and that it is possible to engineer these TESs to work under a wide range of conditions by engineering T_c to the exact temperature we desire to most easily detect particles of a particular energy.

2.3.2 Readout

So far, TESs seem to be the perfect kind of detector – its theoretical principles and the necessary engineering techniques are already well developed, and they work for detectors spanning a very wide range of energies. So why is there even a need for other detection devices? Here, we turn to on one of the primary "issues" with a TES detector: the way we read out a signal from it.

As mentioned before, a "signal" from the TES is generated via a change in the resistance of the superconducting element due to heat, which reads out as a change in current. In the case where we use the TES as a particle detector, the power P_{ext} delivered to the TES can usually be modeled as a delta function $P_{\text{ext}} = E_0\delta(t - t_0)$, and in this regime the current through the TES can be approximately given by:

$$I(t) = -\frac{1}{V}\frac{E_0}{\tau_{\text{eff}}}e^{-\frac{t}{\tau_{\text{eff}}}}$$

Again, the specifics of this equation don't really matter too much – what matters in this equation is noticing that the amplitude is proportional to E_0, and since the energy deposited by any particle colliding with the detector is extremely small the change in $I(t)$ needs to be amplified in order for it to be detectable by our experimental

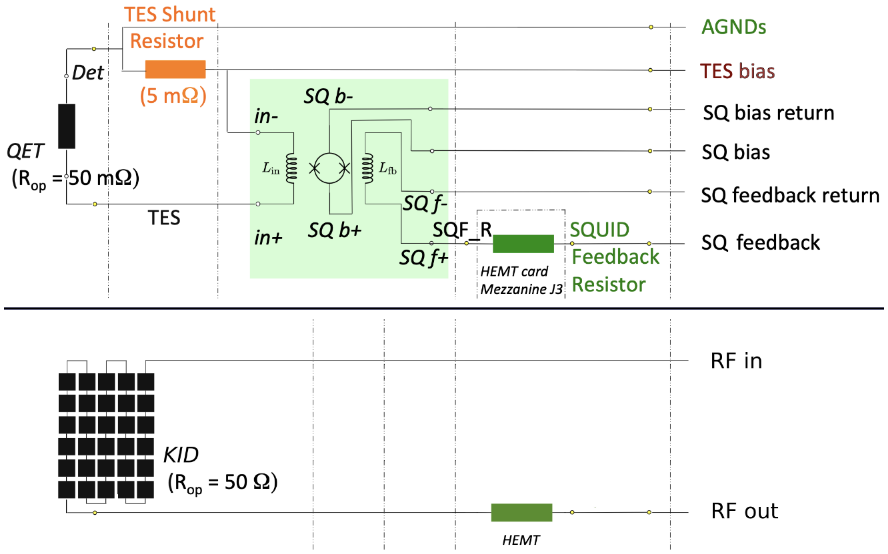

Figure 2.2: Circuit block diagram comparison of a TES-based detector with SQUID amplification (top) compared to that of a TES-based detector (bottom). The SQUID amplification circuitry is boxed in green. Of note is the complex circuitry required to calibrate TES readout, whereas no such amplification is needed in principle for a KID-based detector. This diagram is taken from [6].

instruments. The most popular way this is done is through the use of Superconducting QUantum Interference Devices (SQUIDs). SQUIDs are electrical circuits that help amplify the signal of a TES, and in general an array of SQUIDs are generally required to amplify the signal produced by the TESs in order to get a measurable readout. According to [5], arrays of 100 SQUIDs are commonplace, which give an output signal of several millivolts.

The main problem with TESs as they stand is the fact that SQUIDs, which are more or less required, are relatively complex circuit elements, especially when they are connected together in an array. Furthermore, there tends to be many different sources of noise one needs to consider when designing a SQUID array, which makes them a rather complex circuit element that requires very special attention. Furthermore, there is the added complication that due to the complexity of SQUID arrays, it makes them much harder devices to diagnose and fix when a malfunction occurs. It is these reasons precisely that motivates a new kind of detector that doesn't require such complex circuitry – the Kinetic Inductance Detector (KID).

2.4 A "New" Device: KID Detectors

A Kinetic Inductance Detector (KID) is another kind of detector proposed as an alternative to TESs for dark matter detection. Unlike TESs, the circuitry involved for a KID-based detector is far simpler than that of a TES-based one (see Fig. 2.2), making KIDs a competitive choice as its simplicity makes fixing problems with the detector much easier compared to a TES-based detector.

Fig. 2.2 shows a diagrammatic comparison between the circuitry involved in a TES-based detector with SQUID amplification to that of a KID-based one. As shown in the figure, the SQUID amplification step, boxed in green in the top diagram, contributes to most of the complexity involved in a TES-based detector. By comparison, a KID-based detector in principle only requires an RF feedline and an HEMT amplifier to read out a signal. While both devices require amplification, the HEMT used has been widely commercialized in in the radio industry, and also requires far

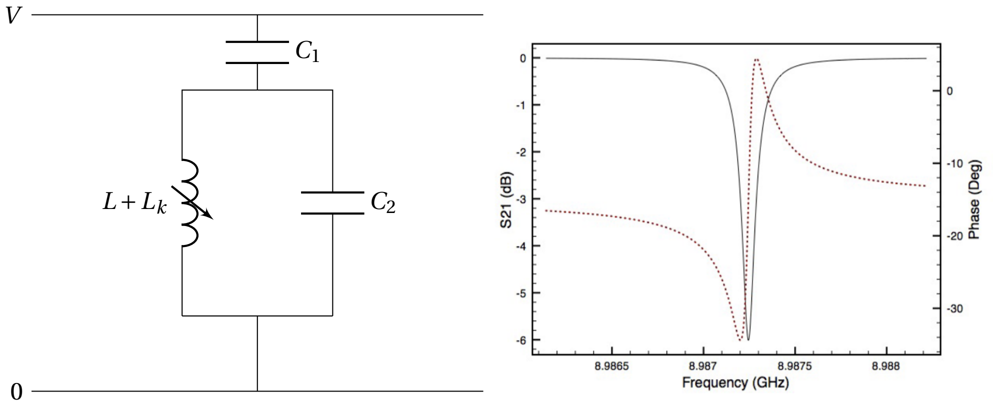

(a) **(b)**

Figure 2.3: (a) circuit schematic of a KID, with the feedline connected at the top and bottom. (b) Example frequency response of a KID. Notice the suppression in the signal at all frequencies besides a narrow band. Figure (b) is taken from [7].

less circuitry than SQUIDs. These two factors combine to make a detector that is much simpler circuitry wise, which makes fixing malfunctions easier while also being cost-effective.

2.4.1 Mechanism

Here, we will go over a high-level overview of how a KID-based particle detector works, much in the same way we did for TESs earlier. Without going into too much detail, one way to understands KIDs is to model them as a LC filter as in Fig. 2.3a, with the top line set to an AC voltage $\tilde{V}$ and the bottom line acting as a ground. For now, we can imagine the input voltage to be a simple sine wave:

$$\tilde{V} = V_0 e^{i(kx-\omega_0 t)}$$

As a filter, the frequency response of the circuit depends heavily on the frequency ω, as shown in Fig. 2.3b. In essence, this means that the amplitude of the output wave depends on that of the input frequency, in a way where only a narrow band of frequencies are suppressed. In particular, the resonant frequency of this system is given by:

$$\omega = \frac{1}{\sqrt{LC}}$$

which is the standard formula for the resonant frequency in an LC filter. Now, the key to the mechanism is that the inductor in this setup is variable with an inductance of $L+L_k$. The value of L remains fixed, while L_k is a quantity that depends on the superconducting characteristics of the KID. In particular, when a phonon interacts with a KID, it alters the number of cooper pairs present, changing the value of L_k, changing its resonant frequency. So, the mechanism of operation can be summarized as follows: when a particle interacts with the detector, it releases phonons that travel through the substrate. Then, as the phonon interacts with the KID, it alters the value of L_k, which changes its resonant frequency $\omega+\delta\omega$ by a small amount. Then, because we never change the input frequency set at ω_0, the shift in resonant frequency results in a change in the output signal which manifests as a pulse, which we record as a detection.

In a typical detector, we place many such KIDs in series with each other, all connected to the central transmission line, typically called the "feedline". Because we have intimate control over the inductance $L+L_k$ of each KID, this

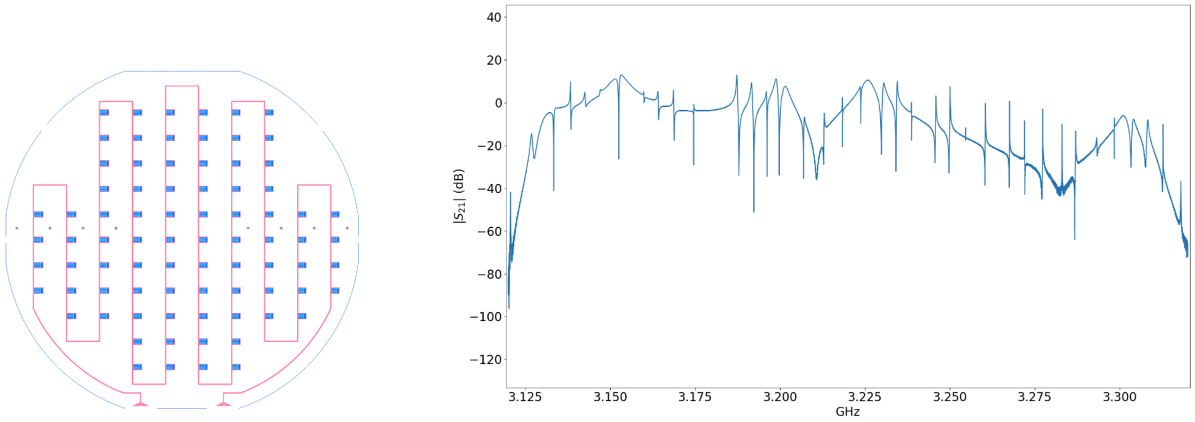

(a)

(b)

Figure 2.4: (a) Physical layout of a KID-based detector. Each blue rectangle in the diagram is an individual KID, each with a different resonant frequency ω_0. (b) Experimental frequency response for this KID device. Notice the distortions present in the frequency response as the overall shape is not flat, and the attenuation of each KID is nonuniform. Both of these diagrams were sourced from [6].

means we can set each KID to have a slightly different resonant frequency, allowing us to scan over a wide range of frequencies. We then send in a mixed signal:

$$\tilde{V} = V_0 \sum_i e^{i(kx-\omega_i t)}$$

where each ω_i corresponds to the resonant frequency of each KID in the circuit. If ω_i is chosen correctly, then a phonon detection would only activate one of the KIDs in our array, giving us information about the energy of the phonon. This can then be extrapolated back to the energy of the particle fairly easily.

2.5 Focus of Study

One aspect of making a successful KID-based detector is the ability to engineer a proper feedline that minimizes the interference introduced to the KIDs by its presence. To illustrate this point, Fig. 2.4a shows a physical realization of what the KIDs look like on a silicon wafer, and Fig. 2.4b shows its frequency response. The red line in Fig. 2.4a is the feedline through which we send in our signal of mixed sine waves, and each blue rectangle represents an individual KID. The resonant frequency of each KID is apparent in Fig. 2.4b, marked by the numerous sharp spikes present in the frequency response.

In an ideal world, the frequency response should resemble a flat horizontal line, with spikes at even intervals from the KIDs. However, this is clearly not the case in Fig. 2.4b – not only is the general shape of the curve not horizontal, but the degree to which each KID attenuates the signal is nonuniform, which makes pinpointing detections much more difficult. All in all, these effects contribute to making KIDs an overall weaker detector. It is believed that one of the primary reasons that gives rise to this effect is the interference caused by introducing a bend in the feedline, which radiates the supposiedly KID-reading signal to excite the box's resonant cavity modes, which in turn couple to the KID resonances and heavily affect the detector's designed performance.

The focus of my thesis is to verify, through simulation, the validity of this claim. Originally, the focus of this thesis was to investigate the use of a slow-wave structure proposed in [8], but I was not able to fully investigate that effect in time for this thesis. Instead, we focus on the feedline itself, seeking to investigate which particular factors in the feedline's design affect its transmission the most.

Chapter 3

Feedline Design

In Section 2.5, we outlined the focus of this thesis, which is to investigate the factors which affect the transmission we get from the feedline. In particular, we will focus on the geometric properties of the feedline, and investigate how different features in the geometry affect its transmission. We will do that in Chapter 4, but before that it will be useful to go over some basic theory of the feedline structure, which we will do in this chapter.

3.1 Theory

For this section, we will largely follow the derivation provided in [9], specifically sections 3.1.1 and 3.1.2. For a more thorough explanation, we refer to the remaining chapters in [9]. The basic structure of the feedline is called a coplanar waveguide (CPW). A diagram of its basic geometry is shown in Fig. 3.1: it is characterized by having three metal strips placed on top of a substrate layer of permittivity ϵ_r, with the two outer strips being grounded and the center being the signal strip. In the general case of a wave propagating through a waveguide, one may write the **E** and **B** fields generally as:

$$\mathbf{X}(x, y, z) = [\mathbf{x}_t(x, y) + x_z(x, y)\hat{\mathbf{z}}]e^{-j\beta z}$$

Here, **X** is a general vector field that represents both the electric field **E** and auxiliary field **H**. In other words, we can decompose the vector field into its transverse and longitudinal ($\hat{\mathbf{z}}$) components; this is fine because both

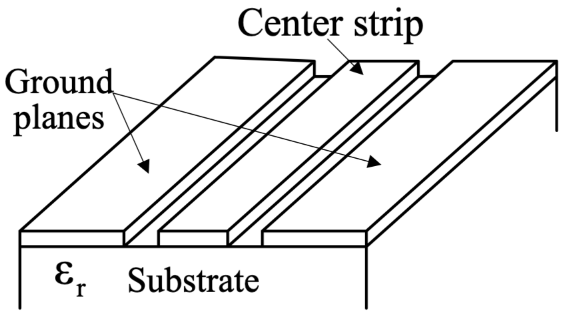

Figure 3.1: General geometry of a coplanar waveguide. Notice that the planes on either end are set to ground, and the center strip is set to a voltage $V(t,z)$. All three strips lie on the substrate ϵ_r. This diagram is taken from [9].

components are linearly independent. To find how waves propagate in a CPW, the standard approach is the same as we use for any waveguide: we just solve Maxwell's equations to see which solutions the CPW configuration admits. In the case where the metal strips are perfect conductors and we submerge the CPW in a homogeneous medium (i.e. we the entire space is filled with a substrate of permittivity ϵ_r), then the solutions to Maxwell's equations indicate the presence of a pure TEM mode. This mode can essentially be understood as a perfect standing wave, where the longitudinal (i.e. $\hat{\mathbf{z}}$ direction) component of $\mathbf{E}$ and $\mathbf{H}$ are zero and the transverse current density $\mathbf{J}$ also goes to zero.

However, the case in any realistic CPW is not so simple: as indicated by Fig. 3.1, the space in which the CPW exists does not have uniform ϵ_r, and secondly, the metal strips have finite conductivity, and are thus not perfect conductors. These two factors combined makes the pure TEM mode impossible with our superconducting CPW. However, as outlined in [9], a superconducting CPW gets *close enough* to a pure TEM mode, in that the non-TEM component of the wave remains small. For this reason, we say that the CPW admits a quasi-TEM mode – in essence, we are able to use the theoretical results that come with the pure TEM without worrying too much because the non-TEM effects are negligible.

3.1.1 Impedance

One characteristic of the CPW we must pay close attention to is its impedance. This is particularly so because in an ideal world, we want a constant impedance throughout the feedline, so as to minimize distortions to the signal that we send into the detector, and be able to impedance match to our start and end ports. Thus, it pays to spend some time discussing the factors that govern the impedance of our CPW.

To do so, we first identify that the CPW has a characteristic impedance and capacitance, denoted L and C. Its characteristic impedance is calculated as $Z_0 = \sqrt{L/C}$, hence the only two quantities that we need to calculate are L and C themselves. As mentioned at the end of Section 3.1, because we can approximate the field solution as a TEM mode (i.e. standing wave), this means that the transverse electric field $\mathbf{e}_t$ and auxiliary field $\mathbf{h}_t$ can be analytically solved by considering it as a two-dimensional statics problem.

First, we tackle the case of the electric field $\mathbf{e}_t$ over all space. To begin, because we are dealing with a superconductor, it is sufficient to make the argument that we are dealing with an approximately perfect conductor – this means that we can assume $\mathbf{e}_t$ inside the superconductor to be zero, and $\mathbf{e}_t$ outside the superconductor to be approximately equal to the case of a perfect conductor. Because there are no charges outside the superconductor, it follows that the electric potential outside the superconductor satisfies Laplace's equation:

$$\nabla^2 \Phi = 0$$

which we can then back out the electric field using $\mathbf{e}_t = \nabla\Phi$. From here, we can calculate the capacitance C, which we detail how to do in Section 3.1.2. For the magnetic field, we can make a similar assumption, approximating the superconductor as a perfect conductor. Doing this, we conclude that the vector potential $A_z(x, y)$ is constant within the superconductor, and outside the superconductor it also satisfies Laplace's equation:

$$\nabla^2 A_z = 0$$

Now that we have the setup, how do we actually go about solving such equations? To do this, one common technique involves using a conformal mapping, which we will detail in the section below.

3.1.2 Conformal Mapping

Conformal mapping is a popular and powerful technique to solving such two-dimensional static problems. In particular, consider the two-dimensional Laplace's equation we ultimately wish to solve:

$$\frac{\partial^2 \Phi}{\partial u^2} + \frac{\partial^2 \Phi}{\partial v^2} = 0$$

The solution to any complex differential equation is a function of the form $z = f(w)$, which can be understood as a function which maps one coordinate (u, v) to a new coordinate (x, y). Complex analysis now gives us a very powerful tool: if f is analytic and is also one-to-one, then it follows that Laplace's equation is invariant under f, so we can pick any other domain (x, y) to solve over so long as it is related to the original (u, v) by an appropriate f. Essentially, what this property allows us to do is solve Laplace's equation in an easier domain (x, y), which then also solves the equation in the original (u, v) domain via f.

One particular f that is useful to us is called the Schwarz-Christoffel mapping. Without going into too much detail, the function f is expressed as:

$$z = f(w) = f(w_0) + c \int_{w_0}^{w} \prod_{j=1}^{n-1} (w' - w_j)^{\alpha_j - 1} \, dw' \tag{3.1}$$

This mapping essentially turns a closed polygon and "straightens" it out, mapping it to the upper half complex plane. Our CPW can be easily modeled as a closed polygon, so this is indeed an appropriate mapping to choose. Then, as described, we solve Laplace's equation in the upper half plane, and use f to map that solution back onto our CPW.

For our particular CPW, it's thin enough that it can be approximated by a geometry with zero thickness, which is beneficial to us because it makes the mathematics significantly easier. This process is summarized in Fig. 3.2, showing how the conformal mapping of the CPW changes its geometry into one that is more desirable on the right, from which it is easier to calculate the capacitance. The quantity K and K' denote the width of the lower and upper CPWs after the transformation, which are calculated using elliptical integrals:

$$K = K(k) = \int_0^1 \frac{1}{\sqrt{(1-x^2)(1-k^2x^2)}} \, dx$$
$$K' = K(k') = \int_1^{1/k} \frac{1}{\sqrt{(x^2-1)(1-k^2x^2)}} \, dx$$

Here $k = a/b$ is a ratio determined by the gap between the center and ground lines, and $k' = \sqrt{1-k^2}$. The exact means by which we derive this formula is not important at the moment, but what is important is that these two quantities, $K(k)$ and $K(k')$ directly relate to the capacitance and inductance:

$$C = \frac{1+\epsilon_r}{2}\epsilon_0 \frac{4K(k)}{K(k')} \quad L = \mu_0 \frac{K(k')}{4K(k)}$$

The key insight in this equation is C and L, and thus $Z_0 = \sqrt{L/C}$, depends *solely* on two quantities: the gap width and the relative permittivity of the substrate given by ϵ_r. So, in theory, what this means is that as long as we don't change the ratio a/b and keep the substrate the same throughout the system, the impedance of the CPW should remain constant.

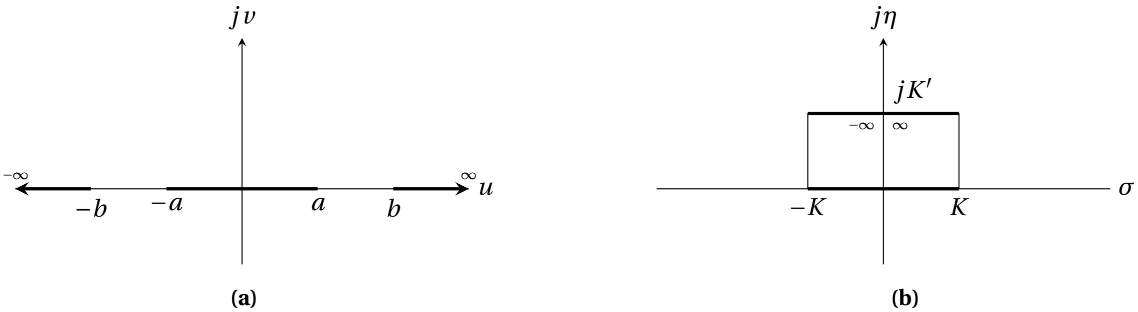

Figure 3.2: Conformal mapping of coplanar waveguide with zero thickness using the SC-mapping. (a) The original diagram for the CPW. (b) The shape of the CPW after the conformal mapping. The "ground" strips are assumed to be infinite; notice how they "bend" into the shape of a rectangle. This figure is taken from [9], reproduced here by for clarity.

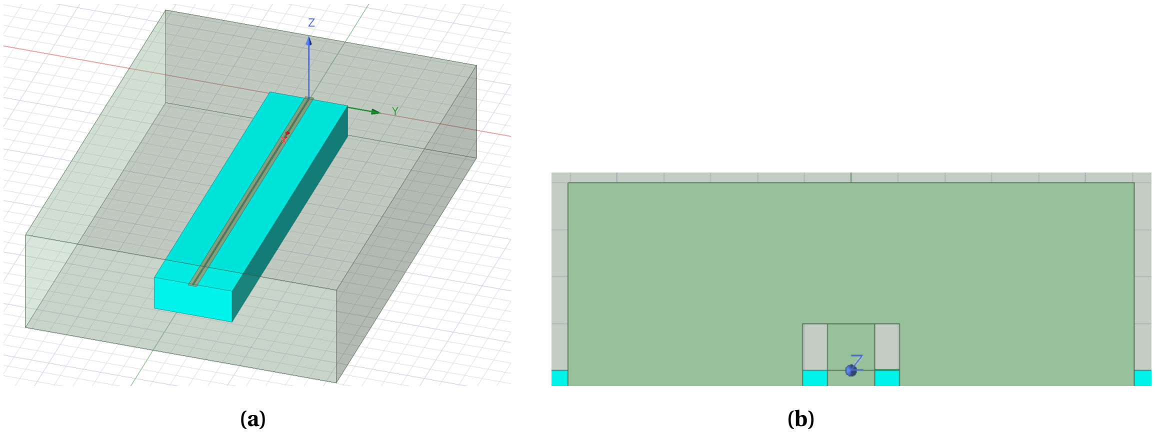

Figure 3.3: (a) Screenshot showing one of the simulation models we performed. The silicon substrate is colored in blue, with the CPW sitting on its surface. The entire apparatus is enclosed in a conducting box. (b) Close-up screenshot of one of the lumped port excitations in the simulation. The lumped port is the small rectangle in the middle, which is surrounded by the center line just below it and the ground lines connected to it from above. Notice that the ground lines are connected together, this is just to ensure that they are both at the same voltage.

3.2 Simulation Setup

Here, we discuss the setup of the CPW simulations that were performed. The simulations were completed in Ansys HFSS, a high-precision finite element analysis (FEA) software that specializes in RF simulations. Fig. 3.3a shows a typical example of a model that was simulated. In each simulation, we use a CPW with a central line of width of 20 microns, and ground lines of width 100 microns. The gap between was calibrated to 10.75 microns. The substrate is silicon, with a relative permittivity of 11.9. Given these specifications, we used a calculator provided by this link to determine the impedance Z_0, which comes out to about 50 ohms. We also enclose the entire waveguide inside of a conducting box – this is also for experimental reasons, as current prototypes for KIDs are built with a similar box around them.

As mentioned, the focus with this thesis is that the impedance of our waveguide remains constant from end to end. Further, we explored at the end of Section 3.1.2 that this condition can be guaranteed, at least in theory, when we don't change the ratio $k = a/b$ throughout the length of the waveguide. The specific reason we chose 50 Ohms is to be consistent with existing experimental setups, but in principle this number can be altered and it wouldn't matter very much for the focus of this investigation.

One detail that should be mentioned is that in Section 3.1.2, we assumed the ground lines to be infinitely long, as shown in Fig. 3.2a, yet this is obviously impossible with physical simulations. However, recall that the only thing we need to ensure is that the gap remains constant, so the width of the ground strips don't matter very much, as long as we make them long enough. In our case, the width of the ground strip is roughly ten times times larger than that of the central line, which is large enough that making the ground strip larger would not make a large difference to the overall impedance.

The simulation is excited by two lumped ports with impedance set to 50 ohms, labeled ports 1 and 2, placed at either end of the waveguide. These ports are responsible for sending the input signal into the waveguide, which we set to frequencies between 1GHz and 10GHz. This specific frequency range was selected since these ports are placed so that they don't touch the substrate material but are in contact with the center and ground planes, via the geometry shown in Fig. 3.3b. As shown, the lumped port hangs off the side of the substrate, which is important because the port is already calibrated to an impedance of 50 ohms; the presence of the substrate could alter this value, leading to impedance mismatch.

3.2.1 Voltage Standing Wave Ratio (VSWR)

The Voltage Standing Wave Ratio (VSWR) is the particular quantity that we calculate in our simulations which characterizes the transmission quality of our CPW. It is a common measurement quantity in RF simulations, as it is a measurement giving us information about how well the impedance matching in our system is. At its core, the VSWR is calculated as:

$$\text{VSWR} = \frac{|V_{\text{max}}|}{|V_{\text{min}}|} = \frac{1+|\Gamma|}{1-|\Gamma|}$$

Here, Γ is called the *reflection coefficient*, expressed as:

$$\Gamma = \frac{V_r}{V_f}$$

Due to the nature of the absolute value in the calculation for VSWR, its range of values is between 1 and infinity. When the value of Γ is zero, this leads to a VSWR value of 1, which therefore corresponds to no reflection (alternatively, perfect transmission). On the other hand, the larger the VSWR value is, more of the wave is said to be reflected as $|\Gamma|$ is larger. In practice, Γ is a complex number that also encodes the phase of the reflected wave, but for our purposes we only care about the magnitude.

Chapter 4

Results

This section will detail the major simulations that were made, and the conclusions we can draw from them in regards to their impact on the overall transmission of the system. To begin, we started with the simplest model: just a straight waveguide according to the specs outlined in Section 3.2. We simulated waveguides of two different lengths, one that is on the order of the characteristic input wavelength and one that is much larger. As it turns out, the scale of the waveguide makes a significant impact when it comes to the transmission parameter we measure.

With the straight waveguide simulated, we then proceed to add one bend to the structure, and investigate its effects on the overall transmission of the waveguide. In particular, we first start by introducing a sharp bend, then we introduced a rounded bend and compared the results of the transmission against each other.

4.1 Straight Waveguide

To begin, we look at the straight waveguide. Here, we simulated two waveguides of two different lengths, one of length 10mm, and the other 20cm. We picked these particular lengths because we wanted to investigate how the transmission of the waveguide changes when the length of the waveguide is on the order of, and also much larger than the wavelength of the incoming wave.

A plot of the VSWR for the short and long simulation are shown in Fig. 4.1. As shown in Fig. 4.1a, the VSWR plot shows one particular frequency at which the VSWR is minimized, at around 6 GHz. This makes sense given the geometry of the CPW, since using the typical formula $\lambda = c/f$ means that a 6 GHZ signal corresponds to a wavelength of 40mm. Given that our straight waveguide is a multiple of the wavelength, it means that this particular wave propagates at a "resonant" frequency, hence the low reflection coefficient. By comparison, none of the other frequencies are resonant, meaning they will all reflect a wave of some amplitude, explaining their higher VSWR values. This fact is also further shown when we slightly vary the length of the waveguide, as shown in Fig. 4.2. In this case, we've increased the length of the short waveguide to 13mm from 10, and as a result in its VSWR plot we now see two minima, indicating that there are two resonant frequencies, at roughly 4.5 and 9 GHz.

The same resonance logic can be applied to the longer waveguide: we just have more resonance frequencies because the waveguide is longer, but the principle still applies. However, there are some interesting features in the long waveguide that we cannot yet explain: firstly, the region between 5 and 7 GHz seems to have no resonance frequencies, and the oscillatory nature of the VSWR begins to break down at higher frequencies. Interestingly enough, this phenomenon is present in Fig. 4.1b, but not present in Fig. 4.2b, possibly suggesting that this is simply

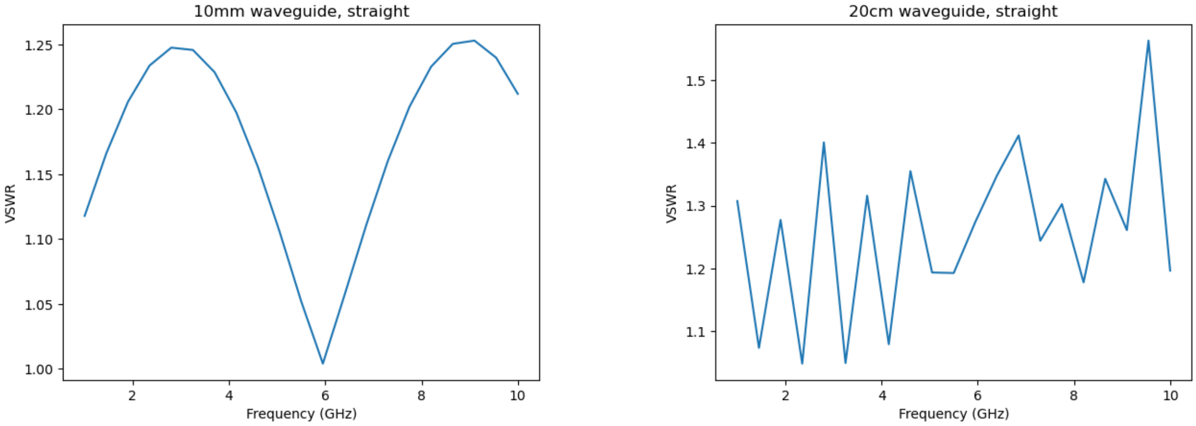

(a) **(b)**

Figure 4.1: Plots of the VSWR for the (a) short and (b) long waveguides. Notice that in the short case, there is one resonant frequency around 6 GHz, and in the case of the long waveguide there are many more resonant frequencies.

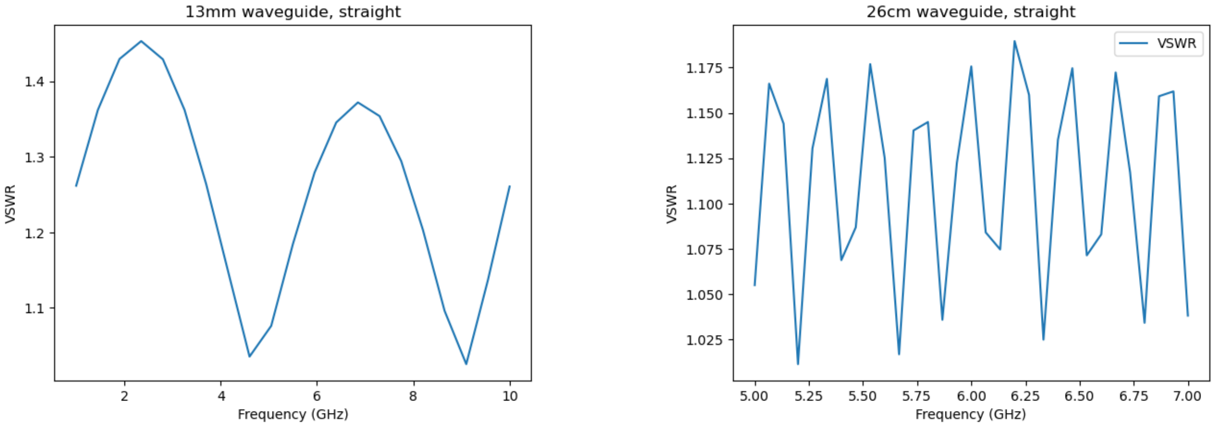

(a) **(b)**

Figure 4.2: Plots of the VSWR for the (a) short and (b) long waveguides, which are slightly longer than those used in Fig. 4.1a and Fig. 4.1b. Notice the difference in the number of minima in each case, strongly suggesting that there is a resonance effect in the waveguide.

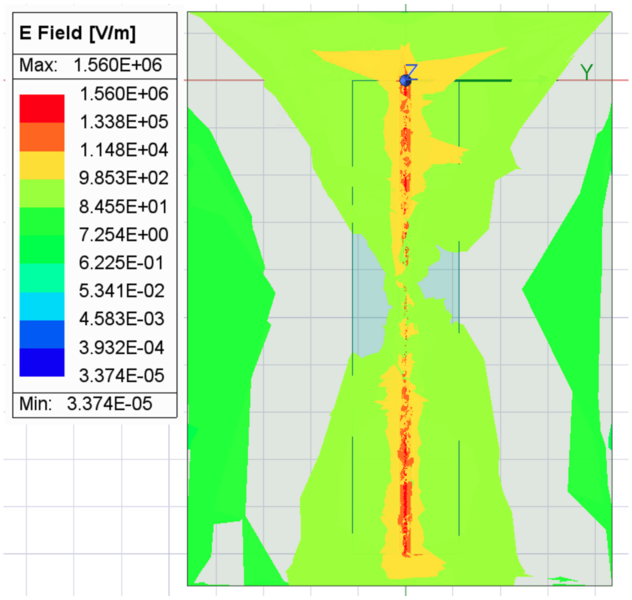

Figure 4.3: Plot showing the electric field solution for a wave of a particular frequency through the short waveguide. Of note is the oscillatory nature of the electric field in the horizontal direction, showing interaction between the CPW and the surrounding box and indicating the presence of box modes.

a numerical artifact and not a physical phenomenon. The latter observation of degrading VSWR with increasing frequency is a common phenomenon seen in many of our simulations, but we do not have a concrete answer for why this occurs.

Another observation we made was that due to the fact that our surrounding box is conducting, the incoming wave seemed to also exhibit standing wave modes with the surrounding box, which we call "box modes". These box modes are particularly problematic, because the reflection of the wave off the surrounding box can interact again with the wave traveling through the waveguide, causing interference in the signal. Evidence of these box modes are shown in Fig. 4.3, where we can clearly see the nonzero electric field that oscillates between the central CPW and the surrounding box on the sides. The mitigation of these box modes are an entirely different area of investigation, which we will not dive into here. This could explain the anomalies we find in the VSWR plot for the long waveguide, but we are unsure at the moment.

Finally, I want to end this section by discussing the importance of simulating these simple waveguides. Although they are non-representative of the waveguides we use in the actual detectors, these simulations help us grasp a basic understanding of how the transmission in a waveguide behaves numerically. This is especially true in our case where many mathematical approximations need to be made to arrive at an analytic solution, so numerical solutions may produce different results than our theoretical results. Furthermore, these simulations also help us establish a baseline for the expected behavior when we add complications to the waveguide, as we will see in the following section.

4.2 Adding Bends

Here, we will add one complexity to our simulation by introducing a bend into the waveguide. The reason we chose to focus on this particular feature compared to others is because when ultimately designing a KID detector, we

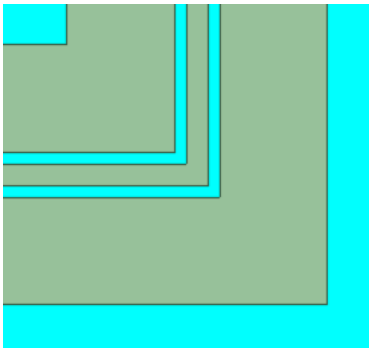

Figure 4.4: Sharp bend introduced to the waveguide.

would ideally want to fit as many KIDs onto our substrate as possible, and the most straightforward way to do that is to have a "meandering" feedline that snakes across the material, like in Fig. 2.4a. As a result, it is important to investigate the effect of bends on our CPW. In particular, we investigate two kinds of bends: first, we introduce a sharp 90° bend, and also a rounded bend. Further, as the ultimate goal of this project is to investigate the viability of the slow-wave structure from [8], it is therefore important to first understand the effects of a bend first.

4.2.1 Sharp Bends

Fig. 4.5

The simplest kind of bend we can introduce is a 90° bend, shown in Fig. 4.4. In order to make the model as close to the straight waveguide as possible, we introduce the bend exactly halfway through the waveguide, and ensure that the *total* length of the waveguide remains the same compared to the straight waveguide. For example, in the case of the 10mm waveguide, this means that the wave travels 5mm, takes the bend, and travels the remaining 5mm.

The simulated results are shown in Fig. 4.5. Of note in these simulations is the fact that in the case of the short waveguide, the presence of the bend doesn't really seem to affect the VSWR plot very much, at least in terms of the resonance. However, the standing wave ratio for the non-resonant frequencies is much worse than Fig. 4.1a, reaching a maximum VSWR of roughly 1.40 compared to 1.25. One possible explanation for this is that in the case of the bent waveguide, the coupling effect to the surrounding box is stronger (as there is less distance for the wave to travel from one port to the other, geometrically speaking), and therefore this could result in more interference especially in the non-resonant frequencies.

The same conclusion cannot be said for Fig. 4.5b, where we see a significant difference between it and Fig. 4.1b. In particular, the resonances seem to not exist for low frequencies, and seems to be more frequent for higher frequencies. The latter observation could be explained by the fact that because the distance to the bend is half that of the full waveguide, this leads to twice the number of resonances. The former observation, however, we do not currently have an explanation for.

Despite these differences however, the magnitude of the VSWR for the long waveguide does not seem to be affected by the bend very much. Our primary hypothesis for this phenomenon goes back to Section 3.1.2. In particular, using the same process outlined in Section 3.1.2, we can create a conformal map for the dielectric constant of the bent waveguide and "straighten" it out, which leads to the mapping shown in Fig. 4.6. As shown in the figure, we can see that introducing a bend into our waveguide is equivalent to altering the dielectric constant in the waveguide – this argument can also be made intuitively, if you imagine "bending" the waveguide and making it straight, then the outer corner would be compressed while the inner corner will be stretched which leads to a variation in the

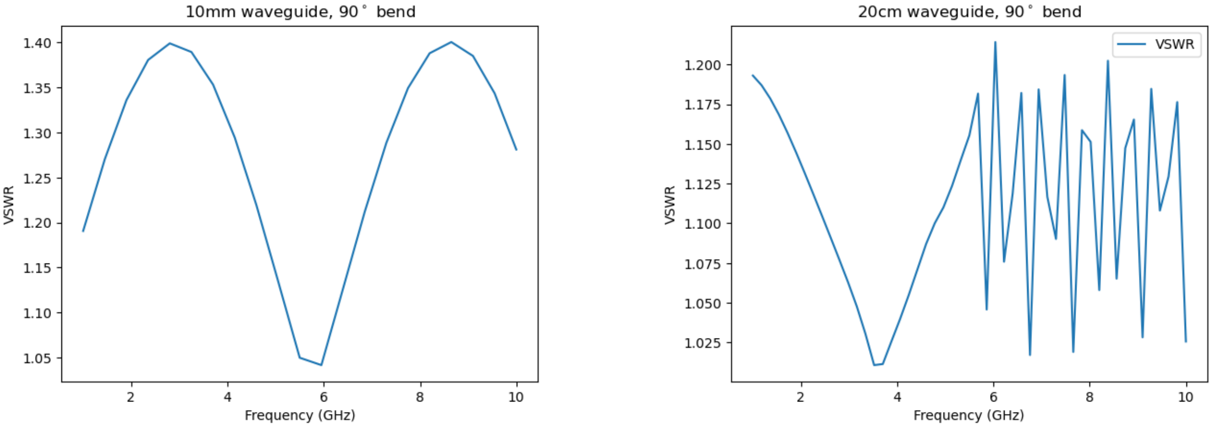

(a) **(b)**

Figure 4.5: Plots of the VSWR for the (a) short and (b) long waveguides with a 90° bend included. Notice that in the case of the short waveguide, the VSWR plot looks nearly identical to that of the straight waveguide, whereas that of the long waveguide does not.

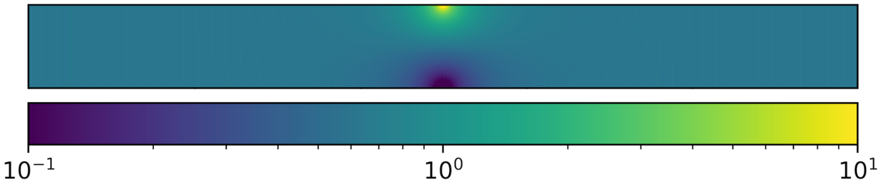

Figure 4.6: Dielectric constant of the bent waveguide after a conformal mapping. Notice that the introduction of a bend manifests as a variation in the dielectric constant of the waveguide, which alters its impedance.

dielectric constant.

Now, in principle, this should change the impedance of the waveguide. However, because the wavelengths that we deal with are far larger than the size of this feature, we believe that the wave simply "skips" over such a feature, which is why we don't see a significant change in the VSWR plot when we introduce the bends.

4.2.2 Rounded Bends

Now, we move on to the third simulation: changing the type of bend from a sharp 90° bend to a rounded one. The motivation for this is twofold: firstly, in the case of a sharp bend, the location of the bend has a different impedance than the rest of the waveguide, because the gap ratio slightly distorts due to the geometry of the bend. This phenomenon is eliminated with a rounded waveguide, as we are able to maintain the size of the gap between the center and ground at all points. In theory, we initially believed that this would result in better transmission.

The bend structure we used is shown in Fig. 4.7. We designed the bend so that the gap between the center and ground lines are maintained, which in theory means that the impedance should not change over the course of this bend. The corresponding VSWR plots are shown in Fig. 4.8, and we can identify some significant differences between this and those shown in Fig. 4.5. Firstly, in the short waveguide, the VSWR for non-resonant frequencies is much larger than that in the case of a sharp bend, indicating that the rounded bend contributes to more reflection than the sharp bend. This is rather surprising, considering that we believed the rounded bend to be unilaterally better than the sharp bend. We do have a hypothesis for this, which we will elaborate on later.

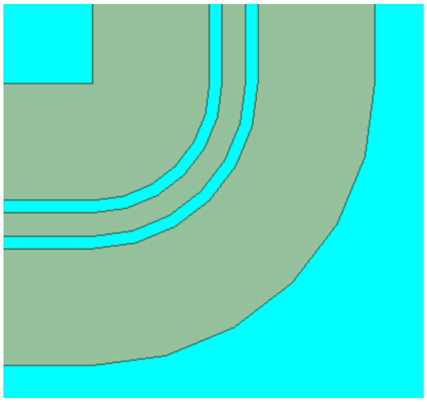

Figure 4.7: Rounded bend introduced to the waveguide.

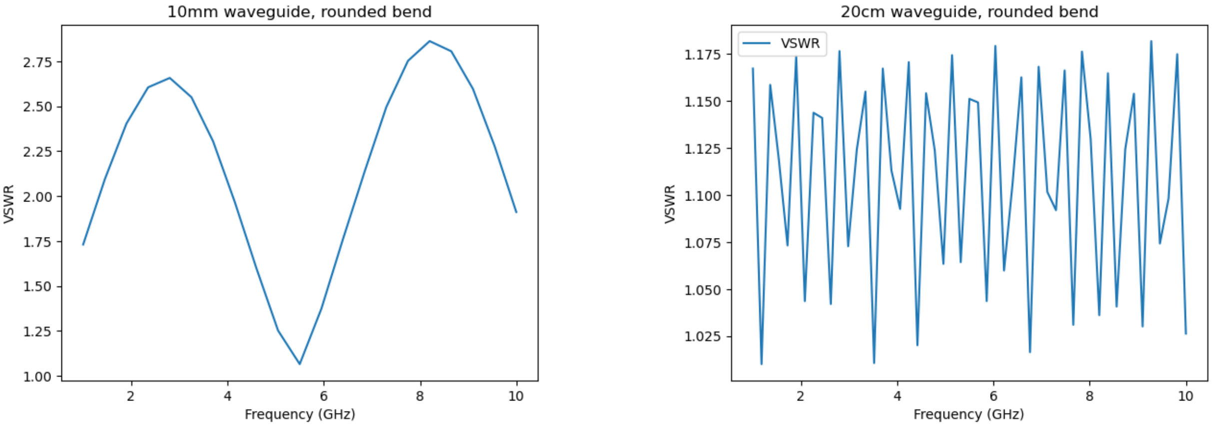

(a) **(b)**

Figure 4.8: VSWR plots for the (a) short waveguide and (b) long waveguide, with a rounded bend. Notice that in the case of the short waveguide, the resonant frequency is untouched, but the VSWR for non-resonant frequencies is much worse.

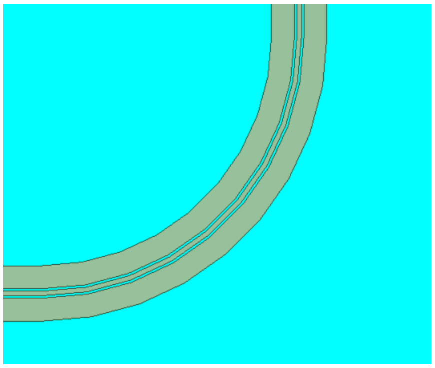

Figure 4.9: Example of the large rounded bend we introduced.

The rounded bend has also affected the long waveguide in interesting ways – firstly, the resonances at low frequencies are present, unlike the plot shown in Fig. 4.5b. However, this seems to be about the only difference that is noticeable by eye. In terms of magnitude, the VSWR seems to match pretty well with that of the Fig. 4.5b, which seems to indicate that in this case, the difference between a sharp and round bend does not really affect the overall transmission. We suspect this could be due to the size of the feature: on a large scale, small differences like rounding out the bend should not do much to the overall transmission because the feature size is small, hence why we don't see much variation in the long waveguide but we do see changes in the short waveguide.

4.2.3 Large Round Bends

The fourth and final simulation we conducted was a variation on the rounded bend. To fully investigate the effect of the rounded bend, we decided to increase the radius of the bend. Theoretically, this has the effect of "smoothening" out the bend even further compared to the earlier iterations, as the bend rounds a larger corner and thus bends slower. Again, we preserve quantities like the gap between the center and ground, so that the impedance of the CPW remains constant throughout.

A figure showing the modification to the bend is shown in Fig. 4.9. Notice that here, the turning radius of the bend is much larger than that in the previous section, and as a result the waveguide bends much more smoothly. Our motivation behind studying this bend was to investigate the effect of "smoothening out" the bending process, to see if this made the transmission of the waveguide better, as this would mitigate some of the reflection that was present in the previous waveguides.

The VSWR plots for these two simulations are shown in Fig. 4.10. Interestingly, these plots look significantly different, and in the case of the long waveguide, significantly worse in terms of transmission compared to the simulations with the smaller bends. Our leading hypothesis in why this occurs goes back to our discussion on conformal mappings in Fig. 4.5. In particular, at the end of that section we detailed how we believe that as long as the distortion to the dielectric constant is much smaller than the wavelength of the signal, then the overall transmission is unlikely to be affected by the variations as the wave would "skip over" these differences. For that same reason, when we introduce a large bend whose size is no longer negligible, it appears that this has a drastic effect on the transmission of the waveguide.

This leads us to the primary result that we've been able to conclude thus far: in terms of creating a waveguide that minimizes the interference observed by impedance perturbation, it is more effective to choose a bend that is small rather than one that better maintains the ratio a/b and yet has an extent comparable to the wavelength. At first, this conclusion may seem counter-intuitive since we are inclined to believe that gradually changing the bend should

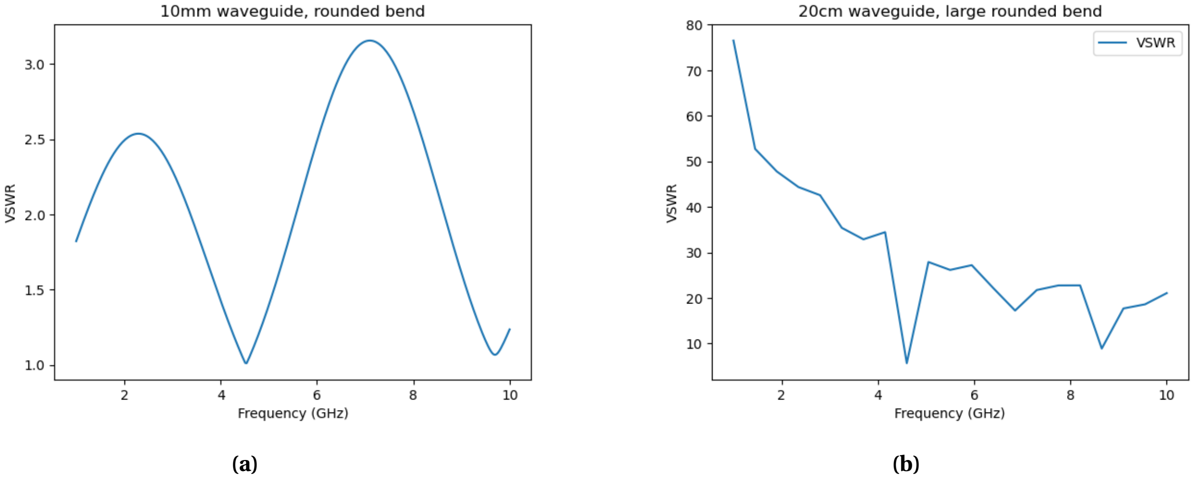

(a) **(b)**

Figure 4.10: VSWR plots for the (a) short waveguide and (b) long waveguide, with a larger rounded bend. Notice that here, we have drastic changes in both VSWR plots, compared to the previous plots in Fig. 4.5 and Fig. 4.8.

allow the wave to bend nicely through the waveguide and minimize reflection, we find that this hypothesis is clearly false, as our simulations and Fig. 4.10 clearly shows.

Chapter 5

Discussion

This chapter is dedicated to a brief discussion about the results we obtained in Chapter 4. In particular, the conclusion that we arrived that large rounded bends result in a worse transmission than that of a smaller bend is extremely interesting, and opens the door to questions about what other large-scale features in the design of a CPW we should be avoiding to achieve an ideal transmission. Furthermore, there are also other questions yet to be answered, such as the effect on the CPW when we start to add more than one bend. In addition to fully simulating the slow-wave structure presented in [8], these are ultimately questions that we didn't have time to get to, which we will definitely investigate in the future following the conclusion of this honors thesis.

There is also the question of how these modifications we investigate behave in the real world. As the results we've concluded here are only based on idealized simulation results, it would be interesting, and also worthwhile, to fabricate a real device and prove that the effects we observe here are actually real. Results obtained from physical device testing will be invaluable to furthering our research into the design of the feedline, as it will be able to show which modifications among the ones we simulated have the greatest effect experimentally.

I want to conclude with some thoughts about my progress over the past two semesters. Overall, despite having failed in our initial goal of investigating the slow-wave structure, I am nevertheless extremely proud of what I've been able to accomplish over the past year. Throughout the past year, being involved in this project has really taught me *how* to tackle problems in the field of research. In particular, it's given me lots of practice in identifying problems, forming hypotheses, then having the ability to devise a plan to test that hypothesis. Mastery of these skills are essential to become a good researcher, and I am thankful to have had this opportunity.

I want to particularly thank Dr. Yen-Yung Chang and Prof. Matt Pyle for agreeing to sponsor me and supervise me through this project, especially given the short time frame I gave you when I decided to move forward with an honors thesis. Without your support this thesis would not have been possible, and for that I am infinitely grateful.

Bibliography

[1] Gianfranco Bertone and Dan Hooper. "A History of Dark Matter". In: *Rev. Mod. Phys.* 90.4 (Oct. 15, 2018), p. 045002. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.90.045002. arXiv: 1605.04909 `[astro-ph]`. URL: http://arxiv.org/abs/1605.04909 (visited on 04/22/2025).

[2] Heinz Andernach and Fritz Zwicky. *English and Spanish Translation of Zwicky's (1933) The Redshift of Extragalactic Nebulae*. Version 1. 2017. DOI: 10.48550/ARXIV.1711.01693. URL: https://arxiv.org/abs/1711.01693 (visited on 04/29/2025). Pre-published.

[3] Mordehai Milgrom and Robert H. Sanders. "MOND Rotation Curves of Very Low Mass Spiral Galaxies". In: *ApJ* 658.1 (Mar. 20, 2007), pp. L17–L20. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/513695. arXiv: astro-ph/0611494. URL: http://arxiv.org/abs/astro-ph/0611494 (visited on 04/26/2025).

[4] Marc Schumann. "Direct Detection of WIMP Dark Matter: Concepts and Status". In: *J. Phys. G: Nucl. Part. Phys.* 46.10 (Oct. 1, 2019), p. 103003. ISSN: 0954-3899, 1361-6471. DOI: 10.1088/1361-6471/ab2ea5. arXiv: 1903.03026 `[astro-ph]`. URL: http://arxiv.org/abs/1903.03026 (visited on 04/27/2025).

[5] Mario De Lucia et al. "Transition Edge Sensors: Physics and Applications". In: *Instruments* 8.4 (Oct. 31, 2024), p. 47. ISSN: 2410-390X. DOI: 10.3390/instruments8040047. arXiv: 2411.01968 `[physics]`. URL: http://arxiv.org/abs/2411.01968 (visited on 04/29/2025).

[6] Yen-Yung Chang. "SuperCDMS HVeV Run 2 Low-mass Dark Matter Search, Highly Multiplexed Phonon-mediated Particle Detector with Kinetic Inductance Detector, and the Blackbody Radiation in Cryogenic Experiments". In: ().

[7] S. Doyle et al. "Lumped Element Kinetic Inductance Detectors". In: *J Low Temp Phys* 151.1-2 (Apr. 2008), pp. 530–536. ISSN: 0022-2291, 1573-7357. DOI: 10.1007/s10909-007-9685-2. URL: http://link.springer.com/10.1007/s10909-007-9685-2 (visited on 05/15/2025).

[8] Hosaeng Kim and R. Franklin-Drayton. "Wire-Bond Free Technique for Right-Angle Coplanar Waveguide Bend Structures". In: *IEEE Trans. Microwave Theory Techn.* 57.2 (Feb. 2009), pp. 442–448. ISSN: 0018-9480, 1557-9670. DOI: 10.1109/TMTT.2008.2011207. URL: http://ieeexplore.ieee.org/document/4752856/ (visited on 05/15/2025).

[9] Jiansong Gao. "The Physics of Superconducting Microwave Resonators". In: ().