Increasing the Detail and Sensitivity of Far-Infrared/Submillimeter Observations in Astrophysics: Kinetic Inductance Detector Development and Molecular Gas Dynamics in Galaxy Merger NGC 6240

by

Adalyn Elizabeth Blenman Fyhrie

B.A., University of California, Berkeley, 2013

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Astrophysical and Planetary Sciences

2019

This thesis entitled:
Increasing the Detail and Sensitivity of Far-Infrared/Submillimeter Observations in Astrophysics: Kinetic Inductance Detector Development and Molecular Gas Dynamics in Galaxy Merger NGC 6240
written by Adalyn Elizabeth Blenman Fyhrie
has been approved for the Department of Astrophysical and Planetary Sciences

Prof. Jason Glenn

Prof. Nils Halverson

Dr. Joel Ullom

Dr. Kevin France

Dr. Benjamin Oppenheimer

Date ________________

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Fyhrie, Adalyn Elizabeth Blenman (Ph.D., Astrophysical and Planetary Sciences)

Increasing the Detail and Sensitivity of Far-Infrared/Submillimeter Observations in Astrophysics: Kinetic Inductance Detector Development and Molecular Gas Dynamics in Galaxy Merger NGC 6240

Thesis directed by Prof. Jason Glenn

To push the boundaries of astrophysics we need to be able to look at the universe in increasing detail. This thesis work advances this goal in two ways: with a technology development project and by using existing technology to investigate an archetypal merging galaxy in extreme detail.

Technology development: Far-infrared observations are ultimately limited by the background radiation emitted by our galaxy and solar system. Large far-infrared observatories that are limited by this background will be critical for understanding galactic and star formation histories over time. To enable this goal, arrays of far-infrared detectors with high enough sensitivities to be astrophysical background limited need to be developed. Working with a team from the Jet Propulsion Laboratory and the National Institute for Standards and Technology, I have investigated two methods for increasing the sensitivities of kinetic inductance detectors (KIDs) in the far-infrared. First, we have fabricated low-volume aluminum and aluminum/titanium nitride bilayer devices to decrease active detector metal volume thereby increasing responsivity. To optimize future iterations on these designs, I have also measured signal lifetimes as a function of aluminum thickness. Second, we investigated phonon recycling devices (simulations and fabrications) that trap phonons generated by recombining Cooper pairs in the active area to elongate quasi-particle recombination time, thereby boosting responsivity.

Observational component: Galactic mergers are immensely complex, often resulting in tidal tails, extremely turbulent gas, active galactic nuclei (AGN), superwinds, and bursts in star formation to only name a few impacts. Using the Atacama Large Millimeter Array (ALMA), I observe the molecular gas in the well-studied galaxy merger NGC 6240 in more detail than ever before.

We analyze high-resolution observations of CO J = $3-2$ and $6-5$ of the central few kpc of NGC 6240 taken with ALMA. Using these CO line observations, we model the kinematics of the molecular gas located between the nuclei of the progenitor galaxies. Our models suggest this gas is a tidal bridge linking the two nuclei that could fall onto the nuclei prior to second pass and feed future starbursts. We also observe high velocity gas (> 300 km/s) that could be accelerated by either gravitational forces from the merger or an AGN outflow. These findings shed light onto small-scale processes that can affect galaxy evolution and the corresponding star formation, with the tidal bridge depositing molecular gas onto the nuclei while other energetic forces accelerate molecular gas further out of the nuclear region.

Dedication

For my family, blood and chosen.

since feeling is first
who pays any attention
to the syntax of things
will never wholly kiss you;
wholly to be a fool
while Spring is in the world

my blood approves,
and kisses are a better fate
than wisdom
lady i swear by all flowers. Don't cry
– the best gesture of my brain is less than
your eyelids' flutter which says
we are for each other:then
laugh,leaning back into my arms
for life's not a paragraph
And death i think is no parenthesis
–E.E. Cummings

Acknowledgements

While it is my name on the title page, this thesis was a group effort. Thank you to everyone who contributed both academically and emotionally, there are too many to call out each by name. You are no less important to me if your name is not here.

The Fyhrie family and its extra appendages: nanananananananaPatmom!, Dave's tours by Dave, KaaaaAAAaatie, Mort, Lugnut, Austintatious ... I look forward to many more silly years and Hairy Clambones. Thanks for keeping me insane.

My chosen family: when laughter failed, whiskey sufficed. I lost a lot of bets over these past 6.5 years in CO (never bet against Becky's tongue or Dan's "photographic memory"), and gained a lot of friends. From bar crawls to rock walls, we have each other's backs (and fronts. and beers.) Jordan! thanks for letting me cry in our shared office without question and teaching me how to use power tools. Thanks to all the ladies/etc. of ladies' night (MikeNicoleLauraBrianaKatieEmily-HilaryBeckyBethanyTasha) and Disney/Ke$ha power hour, Thursdays were a respite of emotional support and hearty laughter. To those from Berkeley (DhanyaPattyEricaGloriaPaz) who stayed with me through that slog and this, you remind me that friendship can span both decades and distance. To all the horses who carried me over the finish line, you can't read but you're important too. Bethany! Putting a diaper on a cat together means we are bonded for life. I love you. Nick! I'm friengaged to your wife so you're stuck with me too. You are one of the reasons I came to CO, and now a reason for me to stay. Becky! I am so lucky to have one of my heroes also be one of my closest friends. You are incredible. Gluten! Thanks for letting me terrorize the neighborhood by screaming into a microphone for a number of years. Daniel! You brought me sweet foods and coffee breaks exactly when I needed them. Austin! You literally pick me up when I am down and I'm so excited to be your forever friend.

Lovelovelove, Addi

Contents

Tables

Table

Figures

Figure

Chapter 1

Introduction

1.1 Far-Infrared and Submillimeter Science: Star Formation and Galaxy Evolution

1.1.1 Far-Infrared and Submillimeter Science

The far-infrared (FIR, $\sim$30-1000 μm Farrah et al. (2019)) is a critical band for understanding the star formation history of the universe. Stars are formed while enshrouded in dust, which is opaque to optical and UV light but transparent to the FIR. Over the entire history of the Universe, approximately 50% of energy in the form of starlight has been absorbed and reprocessed by dust into the FIR (Puget et al., 1996). Additionally, high-redshift (up to z=7.5, around the epoch of reionization) dusty galaxies have been recently discovered; a systematic study of these galaxies in the FIR could unveil star formation's link to the reionization of the universe (Watson et al., 2015). The FIR contains continuum emission and absorption from dust grains at around 15 to 100 K that are heated by emission at shorter wavelengths, making it crucial for understanding energy re-processing and dust properties of galaxies.

The FIR and submillimeter contain a vast range of molecular and atomic lines, from the frequently observed CO to the less commonly observed space booze (CH_3CH_2OH; ethanol). These molecular lines often trace cool molecular gas, a critical component of the Universe because it is the fuel for star formation. Understanding the Universe's star formation history therefore requires an understanding of the molecular gas in galaxies over time.

The most abundant molecule in molecular gas is H_2, but it is notoriously difficult to observe despite its prevalence. H_2, being a linear homonuclear molecule, does not have a permanent dipole moment. The oscillator strength, the probability that a molecule will transition between energy states via absorption or emission of a photon, is proportional to the dipole moment squared. This means H_2's spectral lines can be quite dim. Additionally, the energy levels of H_2 have quite a large average spacing and require hot environments ($>$ 500 K) to excite them (Lacy et al., 2017). In order to observe molecular gas, especially cold molecular gas, astronomers must therefore observe other abundant molecular species. CO is used the most often for this purpose, with an abundance second only to H_2 and bright lines easily excited in cool clouds ($>$ 5.5 K). The amount of H_2 is inferred from that of CO, assuming a relative abundance of $CO/H_2 \sim 10^{-4}$.

When determining the observability of a molecular line, the critical density of a molecular line is important to consider. The critical density quantifies the density at which the radiative de-excitation is occurring at the same rate as collisional de-excitation. Many astronomers use the critical density as a representative density above which a line will appear bright (be considered "observable"). This is a good rule of thumb, but there are caveats to consider. The critical density is often 1-2 orders of magnitude higher than the effective density, the density at which an integrated line brightness of 1 K km/s is observed (considered an observable line) (Shirley, 2015). This indicates that lines can become bright well below their critical densities, and as such the presence or absence of a line does not inherently determine the density of a gas. Regardless, the critical density is a good rough indicator of what gas density a molecular line traces. The low-level CO J lines have modest critical densities of 10^3-10^4 cm^{-3}, indicating they can trace relatively thin gas (Combes, 2018). Molecules such as HCN, HCO+, and CS have critical densities more than two orders of magnitude higher than these CO lines, and as such they are used to trace dense gas (Combes, 2018).

The star formation rate (SFR) of a galaxy depends on the physical properties of the molecular gas of a galaxy, especially its mass, density, and temperature. Bursts in star formation indicate changes to the molecular gas reservoirs that power them. Galaxy mergers and interactions have been shown to be associated with enhanced star formation, indicating alterations to their molecular gas reservoirs triggered by these processes. These changes could take the form of either enrichment of the molecular gas reservoir or an increase in efficiency of turning extant molecular gas into stars (Pan et al., 2018).

While mergers are associated with bursts in star formation, the contribution of mergers to the overall stellar content of the Universe has not been well quantified. The peak of star formation occurs around a redshift z of 2 with about half of all stellar mass formed by z of 1 (see Conselice (2014) and Williams et al. (2011) for a review). The evolution of galaxy merger rates is not as well quantified: some studies suggest galaxy merger rates are relatively constant across $z \sim$ 1-2 (Stott et al., 2013), while others find the galaxy merger rate increases from z = 0-1 with a peak at 1.5 followed by a plateau or slight decrease, reminiscent of the star formation history (Fvre et al., 2013). Observations suggest that major mergers are unlikely to be responsible for the bulk of the stellar mass growth (e.g. Rodighiero et al. (2011), Stott et al. (2013), Lofthouse et al. (2016)). The claimed percentage of stellar mass major mergers are responsible varies, with some claiming up to 30% of stellar mass in galaxies of mass $>10^{10.5}$ $M_{\odot}$ is derived from major mergers since $z{\sim}3$ (Fvre et al., 2013). Some simulations find that major and minor mergers together contribute around 35% of star formation at $z{\sim}3$ and 20% at $z{\sim}1$. Regardless of the specifics, mergers certainly play an important role in the star formation history of the Universe.

Some of the most dramatic SFR's in merging systems are seen in (ultra)-luminous infrared galaxies, or (U)LIRGs. These galaxies are defined by their infrared luminosities, L_{IR}= $10^{11}-10^{12}$ $L_{\odot}$ for LIRGs and 10^{12} - 10^{13} $L_{\odot}$ for ULIRGs. The IR luminosity originates from dust heated by star formation and, in some cases, active galactic nuclei (AGN) (Lonsdale et al., 2006). In most cases this activity appears to be triggered by galaxy mergers and interactions, though there are exceptions. As compared to other interacting galaxies, (U)LIRGs have both higher star formation rates and star formation efficiencies (SFE; defined as star formation per molecular mass). This difference is even more dramatic for high-z (U)LIRGs, with SFE's around 1.5 orders of magnitude higher than local interacting galaxies (Pan et al., 2018). Because molecular gas and star formation are intimately linked, understanding the molecular gas properties of galaxy mergers, especially those classified as (U)LIRGs, is paramount in understanding star formation processes over time.

Luminous infrared galaxies at high redshifts ($z \gtrsim 2$) are commonly called submillimeter galaxies (SMGs). It is common to view SMGs as the high-redshift counterparts to (U)LIRGs, however, there are significant differences between luminous galaxies locally ((U)LIRGs) and at high redshifts (SMGs). One recently quantified difference is that (U)LIRGs are more compact by nearly an order of magnitude and consequently have higher infrared surface brightness (luminosity / size) than SMGs (Rujopakarn et al., 2010). Currently, it is not clear from observations whether (U)LIRGs constitute the same IR-bright galaxies that existed in the early Universe. However, with the peak of star formation occurring around $z \sim 2$, local LIRGs still serve as a laboratory for important star forming merging galaxies at least out to the peak of star formation.

In addition to their elevated star formation rates, galaxy mergers are a critical group of galaxies to study due to their place in the hierarchical structure of galaxy buildup over time. Current cosmological models propose that massive galaxies in the local Universe were assembled via successive mergers of lower-mass galaxies. Local galaxy mergers provide snapshots in time of this assembly process. Detailed observations of many galaxy mergers at different stages are needed to allow astronomers to build up a complete picture of galaxy evolution.

1.1.2 This Work: Galaxy Merger NGC 6240 as a Local Astrophysical Laboratory

Galaxy merger NGC 6240 is an example of a LIRG in the local universe (z = 0.02448) in the midst of a gas-rich major merger event (Fried and Schulz, 1983) that is triggering high star formation rates (Genzel et al. (1998); Tecza et al. (2000)) and AGN activity in its two nuclei (Vignati et al., 1999). NGC 6240's far-infrared (FIR) luminosity $L_{FIR} \approx 10^{11.8}$ $L_{\odot}$ (Sanders et al. (1988); Thronson et al. (1990); Sanders and Mirabel (1996)) is just below that needed to classify it as a ULIRG, though it is expected to cross this threshold when a second starburst is triggered during final coalescence (Engel et al., 2010). As such, NGC 6240 presents an excellent opportunity to study an example of a merger-driven ULIRG just before it passes this threshold. Its proximity facilitates high-resolution, detailed observations of the processes occurring during this merger. Lessons learned from the detailed study of NGC 6240 can be extended to higher redshift merging galaxies that cannot be studied at such high resolution.

In Chapter 6 we analyze high-resolution observations of CO J = $3-2$ and $6-5$ of the nuclear region of NGC 6240 taken at the Atacama Large Millimeter Array (ALMA). Using these CO line observations, we model the kinematics of the molecular gas located between the nuclei of the progenitor galaxies. Our models suggest this gas is a tidal bridge linking the two nuclei that could fall onto the nuclei prior to second pass and feed future starbursts. We also find possible support for the presence of an AGN outflow or gravitational slingshot accelerating molecular gas away from the nuclear region. These findings shed light onto small-scale processes that can affect galaxy evolution and the corresponding star formation, with the tidal bridge depositing molecular gas onto the nuclei while other energetic processes push molecular gas further out of the nuclear region.

1.2 History of Far-IR and Sub-mm Observatories

The first observatory to use far-infrared light to study galaxy evolution was the *Infrared Astronomical Satellite* (IRAS) launched in 1983. This satellite made the first maps of the sky at 12, 25, 60, and 100 μm. The Cosmic Background Explorer (COBE) satellite launched in 1989 made the first measurements of the cosmic infrared background (CIB), interpreted as high-energy radiation from stars in the early Universe that was absorbed and reprocessed by dust. This interpretation was supported by observations from the *Infrared Space Observatory* (ISO), launched in 1995 and observing in both the mid- and far-infrared from 2.5 to 240 μm. Its spectrometer also pioneered the use of mid-infrared spectra to study the energy sources and physical conditions of dusty galaxies (Genzel and Cesarsky, 2000). Launched in 2003, the *Spitzer Space Telescope* was ISO's successor, with similarly general-purpose infrared observing capabilities but with a larger primary mirror. *Spitzer* found faint, massive galaxies at $z > 6$, resolved dusty galaxies out of the cosmic infrared background at 20μm $<\lambda<$ 200 μm, and measured star formation out to $z >$ 3 (Soifer et al., 2008). The AKARI satellite (Murakami et al., 2007) launched in 2006 and operated until 2011 built on the observations of IRAS and completed a more in-depth all-sky survey, and also took detailed mid-infrared observations of nearby objects. At a similar time, ground-based observatories like the Submillimetre CommonUser Bolometer Array (SCUBA) camera on the 15-m James Clerk Maxwell Telescope (JCMT) on Mauna Kea observed the sky in 450 and 850 μm, discovering hundreds of high-redshift galaxies that are extremely luminous in the FIR/sub-mm (Blain et al. (2002), Casey et al. (2014)). These galaxies have some of the highest star formations at redshifts around $z \sim 2$ (Lutz, 2014).

Herschel Space Observatory (Pilbratt, G. L. et al., 2010) is the most recent general-purpose space-based infrared observatory, operating from 2009 to 2013 and observing from 60 to 670 μm. Its science objectives spanned from the study of comets and Kuiper belt objects to studying the reservoirs of gas in nearby galaxies. It mapped nearly 1300 deg^2 of the sky, resulting in the detection of more than one million galaxies. These efforts are matched by ground-based observatories like the South Pole Telescope observing at 1.4-2 mm that have discovered hundreds of high-redshift lensed dusty galaxies. These leaps in discovery helped characterize the statistical properties of early star-forming galaxies and the star formation history of the universe.

1.2.1 The Atacama Large Millimeter/Submillimeter Array (ALMA)

The Atacama Large Millimeter/Submillimeter Array (ALMA) is currently a powerful workhorse for the astronomical community. It is also the telescope used for observations presented in this thesis, and as such warrants discussion here.

ALMA observes from 35 to 950 GHz (300 μm to 8.5 mm) using 66 dish antennae operated interferometrically, 54 of which are 12 m in diameter and 12 of which are 7 m in diameter. Its observing frequencies are broken up into 10 observational bands that fit into the atmospheric transmission windows, shown in Figure 1.1. The primary absorber in this wavelength range in Earth's atmosphere is water vapor, meaning ground-based IR and FIR observatories are located at high elevation, extremely dry sites. As such, ALMA is located in one of the driest places on Earth – the Chajnantor Plateau in the Atacama desert in Chile at an elevation of 5000 m.

ALMA's scope earns it the title of the most complex ground-based astronomical observatory.

Figure 1.1: Atmospheric transmission at ALMA for precipitable water vapor (pwv) of 0.25, 0.5, 1, and 2 mm created with the ALMA atmospheric modeler https://almascience.eso.org/about-alma/atmosphere-model. At this observing site, the pwv is typically 1.0 mm and is below 0.5 mm 25% of the time during the driest observing months from May to September.

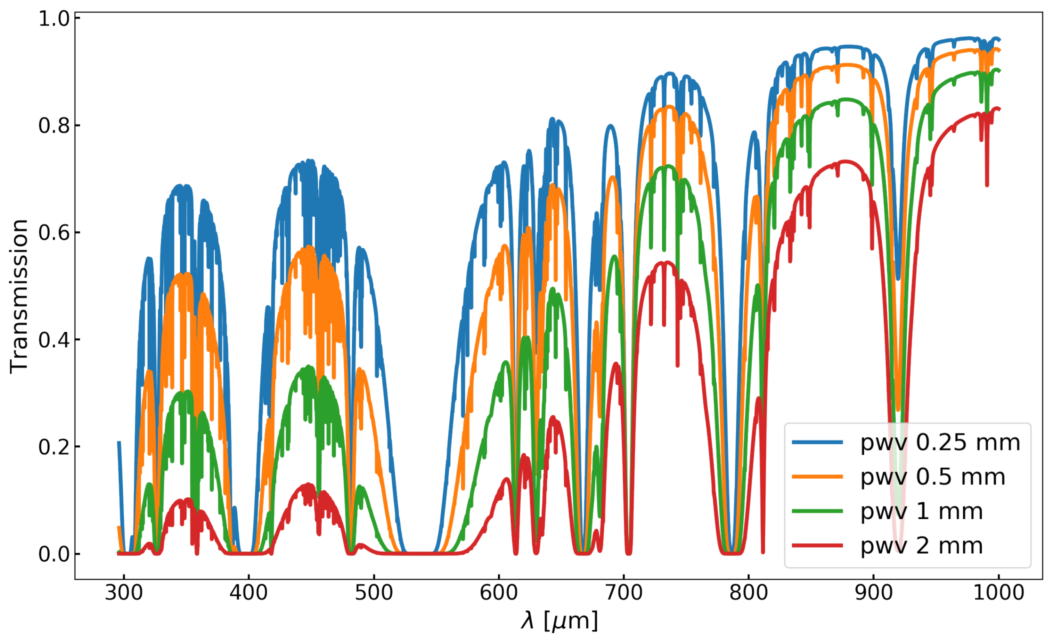

The dishes are mobile, able to be transported on the backs of trucks into different configurations anywhere from 150 m to 16 km across. Wider configurations of the 12 m dishes provide better angular resolution and compact configurations facilitate deep observations of dim, large objects. The 7 m dishes, along with four 12 m dishes, comprise the Atacama Compact Array (ACA) used for imaging of large-scale structures. Its angular resolution in the most compact configuration is $0.6''$ at 675 GHz and $4.8''$ at 110 GHz. In the extended configuration, its resolution is 6 mas at 675 GHz and 37 mas at 110 GHz.

ALMA's flexibility, sensitivity, and high angular resolution have facilitated many ground-breaking discoveries. Its images of planets forming in protoplanetary disks and discovery of organic molecules in these disks have made headlines. It observed the nuclear region of a comet in our Solar system, confirming that it is rich in organic molecules. It participated in the Event Horizon Telescope project, assisting in producing the first image of a black hole's event horizon. Its observations spectroscopically and spatially resolved two objects at z of 7.15, revealing them to be the earliest example of a galaxy merger observed to date (Hashimoto et al., 2019). The multitude of discoveries facilitated by ALMA cannot be listed exhaustively here. Nevertheless, its power can be understood by considering these examples that range from the astronomically tiny to the gigantic, from nearby to extraordinarily distant.

1.3 Astronomical Detectors from the sub-mm to IR: State of the Field

The observatories discussed above could not exist without detectors to digitize the light they collect. Below, we present the current state-of-the-art (far)-infrared to sub-mm detectors that make these and future observations possible.

A common measure of a detector's sensitivity is the noise equivalent power (NEP), details of which are presented further along in this work. For this overview, one only needs to know that lower NEP's indicate more sensitive and/or less noisy detectors, with future cutting edge observatories requiring NEP's around 3×10^{-20} W/$\sqrt{Hz}$ to be limited by photon noise from astrophysical backgrounds (Section 1.4).

1.3.1 Transition Edge Sensors

A transition edge sensor (TES) is a superconducting bolometer operated in the transition temperature between superconducting and normal-state. At this temperature, the slope of resistance with temperature is quite sharp. Therefore, when the material absorbs a photon and heats up, the change in resistance results in a clearly detectable signal. TES bolometer arrays have been used on many prominent CMB experiments (e.g. Ade et al. (2015), Henderson et al. (2016), Hubmayr et al. (2016), among many others), as well as for calorimeters in wavelengths from the optical to the γ-ray. TES arrays are also used in the infrared. SCUBA2, a 10,000 pixel camera observing at 450 and 850 μm on the James Clerk Maxwell Telescope has been observing since 2001. HAWC+ and HIRMES are TES-based cameras observing on SOFIA. The SPICA mission, an upcoming space-based observatory with a 2.5m diameter mirror observing from 12 to 230 μm, has baselined TES bolometers for its far-IR spectrometer SAFARI. This mission will be the TES's first demonstration on a space-based observatory (Farrah et al., 2019).

These missions have advanced TES arrays to a level of high technological maturity and extreme sensitivity. Derived NEP's as low as 7×10^{-20} W/$\sqrt{Hz}$ have been achieved in the Far-IR Kenyon et al. (2009), and NEP's of 1×10^{-19} W/$\sqrt{Hz}$ have been demonstrated at the IR by SRON (Suzuki et al., 2016).

The main downside of TES bolometer arrays is their complex fabrication and readout schemes. In earliest iterations of the bolometer arrays, each sensor required its own superconducting quantum interference device (SQUID) amplifier and its own readout line. Recent advances in multiplexing chips have mitigated this issue (e.g. de Haan et al. (2019)), allowing for up to 132 TES's to be read out on a single coaxial line. However, on-chip multiplexing systems are still not existent and the largest multiplexed arrays still require thousands of wirebonds for their functionality. These complex readout systems are undesirable for large space-based cameras due to weight, space, and complexity concerns, and are not going to be easily scalable above $\sim10^4$ pixels.

1.3.2 Kinetic Inductance Detectors

Kinetic Inductance Detectors (KIDs) are superconducting resonators whose inductance changes when they absorb photons resulting in a measurable shift in the resonance frequency. A single room temperature coaxial cable and amplifier can read out thousands of KID responses at once due to KIDs' natural ability to pack their resonant frequencies tightly together in frequency space ('frequency multiplexing'). KIDs are a relatively recent technology with their first demonstration in 2003 (Day et al., 2003).

Far-infrared KID arrays have been demonstrated on a number of ground-based telescopes. On the Caltech Submillimeter Observatory, MAKO and MUSIC completed successful observations in the sub-mm in the mid 2010's. Other ground-based far-infrared KID cameras include A-MKID at APEX and NIKA/NIKA2 at IRAM. KID cameras have also been demonstrated in the optical/near-infrared, including DARKNESS on the Palomar Observatory that will conduct direct imaging of exoplanetary systems. Balloon-borne experiments using KID cameras have also been commissioned, for example BLAST-TNG launching in December 2019.

The development of KID technologies has become a worldwide endeavor, with many promising NEP's resulting from these efforts: Janssen et al. 2014 Janssen et al. (2014) achieved an NEP of 1.4×10^{-18} W/$\sqrt{Hz}$ at 350 μm for optical loading of 0.1 fW, and P. Diener et al. Diener et al. (2012) reached an array-averaged dark NEP of 5.4×10^{-19} W/$\sqrt{Hz}$, with some pixels as low as 4.4×10^{-20} W/$\sqrt{Hz}$. More recently, J. Bueno et al. demonstrated a 989 pixel array with an array-averaged non-optically loaded NEP of 3×10^{-19}W/$\sqrt{Hz}$ (Bueno et al., 2018). Further development is still needed for astrophysical background-limited spectroscopic observations.

1.3.3 Quantum Capacitance Detectors

Quantum capacitance detectors (QCDs) are a very recent technology, with the first proof of concept in 2010 (Bueno et al., 2010). They are based on the single Cooper-pair box, a voltage biased metal island connected via a tunnel junction to superconducting leads, which also serve as the photon absorbers. The island is sufficiently small as to have a capacitance that is sensitive to the presence or absence of single electrons. This causes a capacitance change when electrons tunnel across from the absorber after photons break Cooper pairs into individual electrons therein (Shaw et al., 2009). The Cooper-pair box is embedded in a resonant circuit that allows the measurement of the changing capacitance via a shift in the resonance frequency.

Initial demonstrations of QCDs are promising, with NEP's of 2×10^{-20} W/$\sqrt{Hz}$ and single-photon counting ability in a 25-pixel array observing at 200 μm (Echternach et al., 2013). While these results are promising, QCDs are quite far behind TES and KID arrays in terms of technological development due to their recent invention.

1.3.4 Impurity Band Conductors

Impurity band conductors (IBCs), also known as blocked impurity band conductors (BIBs), are a common detector technology in the near- and mid-infrared (with some extension into the far-infrared) from $\sim$1 μm out to $\sim$180 μm. They are comprised of a semiconductor doped with "impurity atoms" sandwiched by electrodes. Examples of impurity atoms are As, Sb, and Ga, and common semiconductors used are Si and Ge. Photons free a charge carrier (either an electron or a hole) after interacting with the impurity atom. This charge carrier travels through the semiconductor based on the electric field applied by the electrodes. The charge carriers then create a current that can be detected by the readout system. The wavelength range to which IBCs are sensitive is mainly determined by the impurity atom, with different atoms resulting in different excitation energies for the charge carriers.

Invented in 1979 (Herter, 1994), their use spans many decades of observatories. One of the most famous uses of this technology were the Si:As, Si:Sb, and Ge:Ga IBCs used for all three *Spitzer* instruments. Si:As IBCs are also slated for use on James Webb Space Telescope's MIRI instrument, sensitive to light from 5 to 28 μm (Love et al., 2005).

The biggest limitation of IBC arrays as compared to TES, KID, or QCD technologies is their limited spectral coverage, with no existing IBC sensitive to wavelengths longer than 220

μm (Cardozo). At the wavelength range to which they are sensitive, beyond their demonstrated efficacy, another benefit to IBCs is that they are operated around 1.5 K while TES and KID arrays require operation below around 1 K.

1.4 The Future of (Far)-Infrared and Submillimeter Observing

With *Herschel*'s retirement, a telescope that can carry out large surveys in the mid- and far-infrared to discover new objects to further our statistical understanding of the Universe's star formation history is needed. This telescope will need to overcome the confusion limits hit by both *Spitzer* and *Herschel*, requiring larger mirrors and more sensitive detectors. James Webb Space Telescope (JWST) will conduct this job well at wavelengths below 30 μm and ALMA excels at detailed observations above 300 μm. However, the intervening wavelengths needed to observe dust emission from local galaxies and redshifted rest-frame mid-infrared emission from galaxies at higher redshifts will be critical to develop our understanding of star formation and galaxy evolution over time.

One option that suits this goal is the Origins Space Telescope (OST), a proposed flagship mission with a 5.9 m diameter mirror that would observe from 2.8-588 μm. Its capabilities include measuring the spectra of transiting exoplanets and wide-band, deep spectroscopic surveys at 0.5 and 20 deg^2. It will require around 60,000 pixels across its six detector arrays, with NEP's as low as 3×10^{-20} W/$\sqrt{Hz}$ for its spectrometer.

Another observatory that will move the astronomy community towards this goal is the Space Infrared Telescope for Cosmology and Astrophysics (SPICA) mission. It is an upcoming space-based observatory with a 2.5m diameter mirror observing from 12 to 230 μm. Its science objectives are similar to those of the OST, though its lower sensitivities (NEP's $\sim 2\times10^{-19}$ W/$\sqrt{Hz}$) and smaller mirror mean slower mapping speeds and fewer astrophysical targets observed as compared to OST.

The Galaxy Evolution Probe (GEP) (Glenn et al., 2019) will also suit these goals. It is a space observatory budgeted at under $1B developed as part of NASA's Astrophysics Probe concept development program. The GEP is designed to survey star-forming galaxies from z=0 to greater than z=4 using spectroscopy and broadband imaging in wavelength bands from 10 to 400 μm. Its current design goal is to have sensitivities limited by photon noise from astrophysical backgrounds and not by emission from the telescope itself, which will be cryogenically cooled to 4 K to reduce infrared emission. This would require large arrays of detectors with noise equivalent powers (NEP's) on the order of 10^{-18} W/$\sqrt{Hz}$ for the imager and $\sim 10^{-19}$ W/$\sqrt{Hz}$ for the spectrometer, calculated below. Kinetic inductance detectors (KIDs) have been baselined as the detectors for the GEP due to their simple focal plane architecture as compared to transition edge sensors (TESs). TESs and their readout technologies, however, have higher technological maturity than KIDs and therefore are a viable option depending on the state of each technology during future design studies.

1.4.1 Astrophysical Background NEP Limit

When observing in the mid- and far-infrared, we must contest with two major sources of astrophysical background: IR light scattered off of dust in the plane of the solar system (zodiacal light) and galactic emission of IR which peaks in the galactic plane. These backgrounds were observed in detail by the Diffuse Infrared Background Experiment (DIRBE) on the Cosmic Background Explorer (COBE) satellite launched in 1989. Representative images of each of these backgrounds from DIRBE are shown in Figure 1.2.

The mean emission of these backgrounds can be modeled and subtracted from any large survey of the sky, but they are still subject to the absolute noise minimum that astrophysicists are limited by: photon noise. The photon noise of these infrared backgrounds is therefore the baseline limiter of any mid- or far-infrared instrument.

Ignoring photon bunching, the noise equivalent power (NEP) contribution of photon noise for light of frequency ν with power P_{sky} in W collected by a telescope with optical efficiency η_{opt} (0.2 for the GEP) is

$$NEP_{photon} = \sqrt{2h\nu P_{sky}/\eta_{opt}}. \tag{1.1}$$

Using the DIRBE maps of the flux density I_{sky} (delivered in MJy/sr) we can then calculate the limiting astrophysical background photon noise for any detector on any telescope:

$$P_{sky}[W] = I_{sky}[\frac{MJy}{sr}]\frac{10^{-20}W/m^2/Hz}{MJy}\Omega_{sky}A_{telescope}\Delta\nu, \tag{1.2}$$

where Ω_{sky} is the solid angle of sky that a detector sees in sr, $A_{telescope}$ is the collecting area of the telescope, and $\Delta\nu$ is the bandwidth of the observing band. For the GEP, its 2 m diameter mirror yields $A_{telescope}$ = π. The bandwidths, band centers, and Ω_{sky} for detectors in different observing bands are listed in Table 1.1.

Using equations 1.1 and 1.2 and the parameters from Table 1.1, we can calculate the NEP needed for the imager and spectrometer to be limited by astrophysical photon noise only. The results of this calculation are found in Figures 1.3 and 1.4. For this calculation, we average across bands of galactic and ecliptic latitude that are 10 degrees wide in order to get a representative sample of sections of the sky. As expected, the darkest areas of the sky are located at the poles (high galactic latitude), finding astrophysical background NEP's at 240 μm of 10^{-18} W/$\sqrt{Hz}$ for the imager and $\sim$$10^{-19}$ W/$\sqrt{Hz}$ for the spectrometer. The NEP requirements are much less stringent at shorter wavelengths, requiring NEP's at 25 μm of $\sim$5x10^{-18} W/$\sqrt{Hz}$ for the imager and $\sim$9x10^{-19} W/$\sqrt{Hz}$ for the spectrometer.

The results from Figures 1.3 and 1.4 can also inform where an astrophysical background-limited surveyor would do best to observe: focusing observations on galactic latitudes with absolute values greater than 50 degrees will facilitate the deepest surveys of the sky.

1.4.2 GEP-B

As discussed above, aluminum KIDs are baselined as the technology to meet the sensitivity requirements of the GEP. However, arrays of these detectors currently have a technology readiness level (TRL) of 3. Balloon-borne telescopes are well-suited to concurrently advance FIR science goals while raising the TRL of KIDs. One such proposed telescope is the GEP-Balloon (GEP-B), a balloon-borne telescope with a 1.2 m diameter mirror proposed by J. Glenn et al. in 2019. Its

Figure 1.2: DIRBE observations in bands 6 (*Left*) centered on 25 μm and 8 (*Right*) centered on 100 μm. Band 6 clearly shows the zodiacal light that follows the plane of our Solar System, while band 8 shows the infrared emission from dust in the galactic plane which peaks at longer wavelengths.

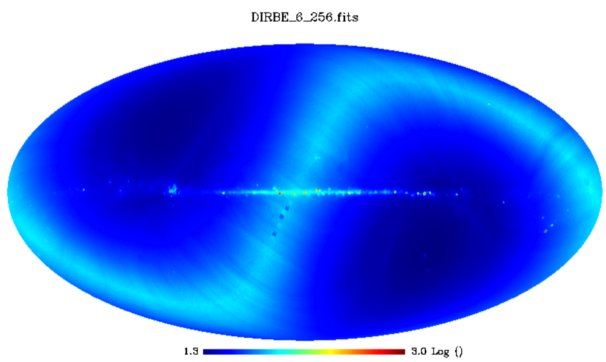

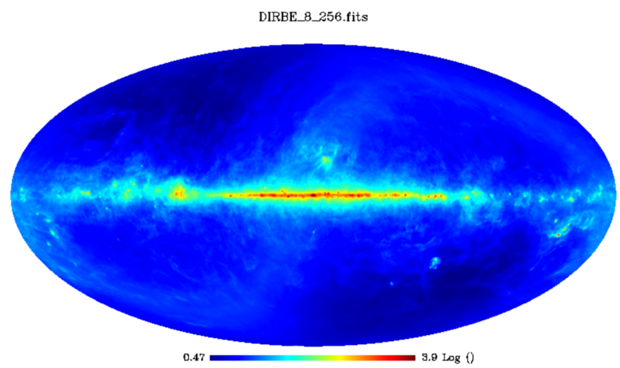

DIRBE band	DIRBE λ_{center} [μm]	GEP $\Delta\nu_{imager}$ [Hz]	GEP $\Delta\nu_{spectrometer}$ [Hz]	GEP Ω_{sky} [sr]
5	12	3.1E12	12.45E10	2.77E-10
6	25	1.2E12	5.98E10	2.77E-10
7	60	6.1E11	2.49E10	2.77E-10
8	100	7.9E11	1.5E10	7.72E-10
10	240	3.3E11	0.62E10	4.26E-9

Table 1.1: The DIRBE band name and corresponding band center, the GEP imager bandwidth covering the DIRBE band center in Hz, the GEP spectrometer bandwidth cover the DIRBE band center assuming R = 200 in Hz, and well as the area seen by a GEP pixel in each band.

Figure 1.3: NEP required by the GEP imager calculated using DIRBE observations averaged over 10 degree wide bands of galactic latitude (*Top*) and ecliptic latitude (*Bottom*) for the bands in Table 1.1. Error bars are calculated by taking the standard deviation of pixels in that band and converting them to NEP.

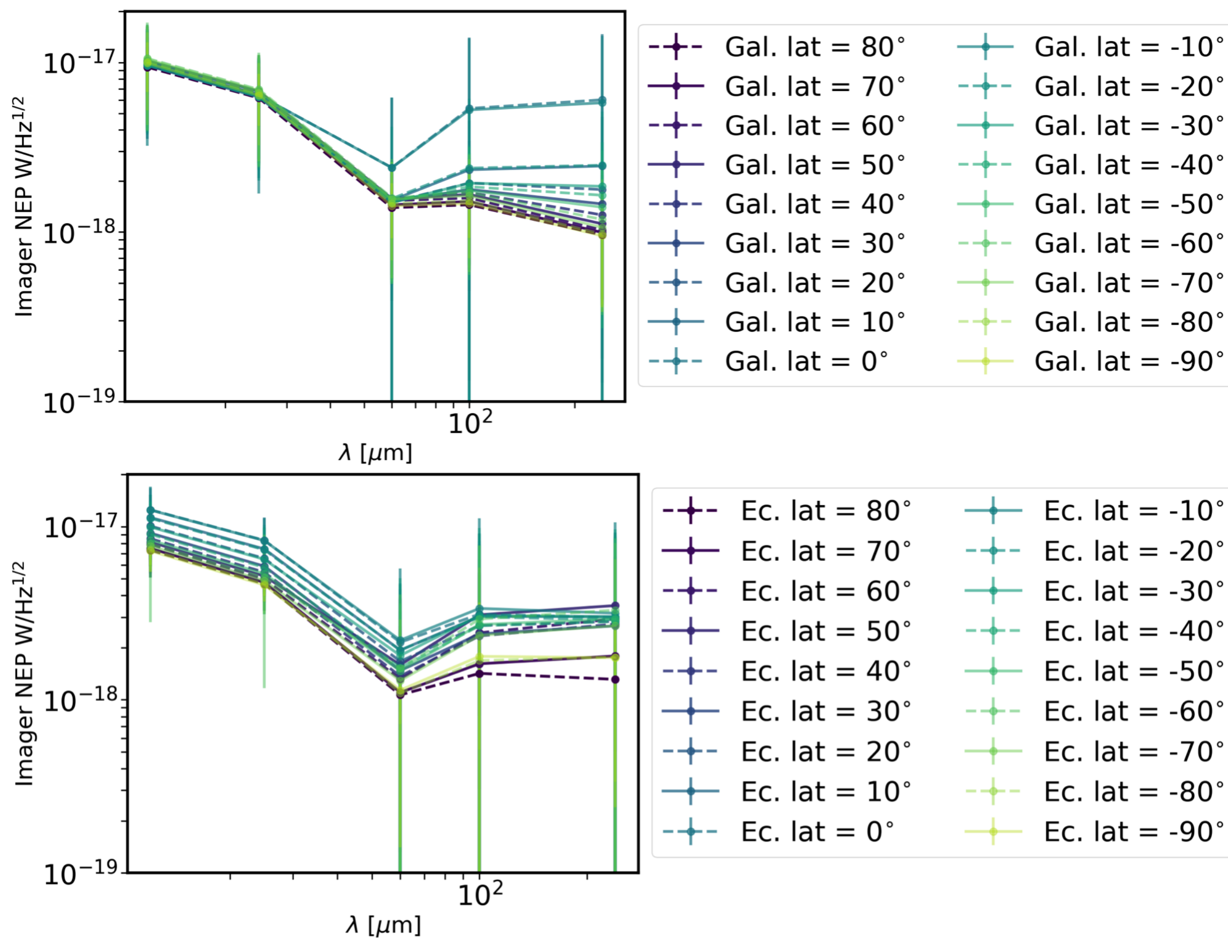

Figure 1.4: NEP required by the GEP spectrometer calculated using DIRBE observations averaged over 10 degree wide bands of galactic latitude (*Top*) and ecliptic latitude (*Bottom*) for the bands in Table 1.1. Error bars are calculated by taking the standard deviation of pixels in that band and converting them to NEP.

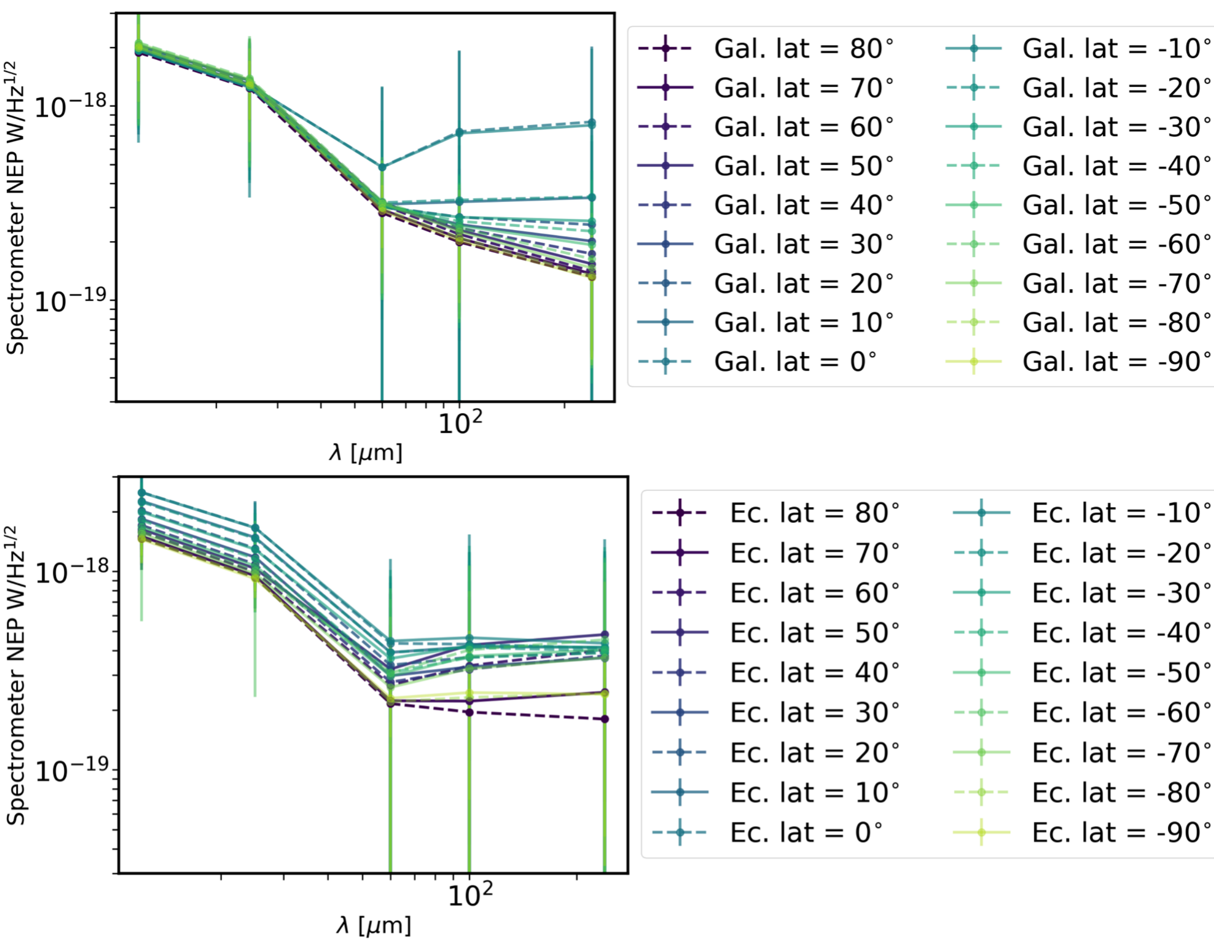

astrophysics objectives include measuring galaxy redshifts using polycyclic aromatic hydrocarbon features and measuring ISM conditions associated with star formation in both the Milky Way and in more distant galaxies.

The following was written by A. Fyhrie for the GEP-B proposal, 2019

The new deep mid- and far-IR observations to be completed by the GEP-B will require large arrays of sensitive detectors. In the past twenty years, arrays of superconducting kinetic inductance detectors (KIDs) have emerged as a competitive technology applicable to wide wavelength ranges including the 10-400 μm needed for the GEP-B(Day et al. (2003), Zmuidzinas (2012), Farrah et al. (2019)).

The detector requirements of GEP-B will be met by arrays of back-illuminated, lumped-element, silicon microlens-coupled aluminum KIDs operating at 200 mK. Antenna-coupled KIDs with sensitivity sufficient for GEP-B have been demonstrated at 350 μm (Baselmans, J. J. A. et al., 2017). For close-packed arrays operating at down to 10 μm wavelength, we have developed a resonant absorber geometry (Figure 1.5) that allows for efficient absorption over a tunable bandwidth even with a relatively high conductivity aluminum film. The minimum line/space dimension required for the absorber is 200 nm, which can be readily achieved using a stepper system. Arrays of 4 x 48 KIDs based on this absorber design and spaced on a 250 μm pitch were fabricated at JPL. Initial testing has shown high ($> 94\%$) yield and high internal quality factors (> 200k). The measured two level system noise is low enough that these devices should meet the sensitivity requirements for all science cases, and optical tests are in progress. As the 10 μm band represents the most difficult case in terms of line spacing and pixel pitch, the requirements should also be readily met in the longer wavelength bands.

The focal plane will consist of 2,864 KIDs. The spectral bands are defined by metal mesh bandpass filters mounted above KIDs designed to couple to the radiation determined by the filter.

Small alterations to the KID geometry must be made to accommodate the higher optical loading environment of the GEP-B while maintaining high quality factors required by the readout system. Simulations show that fabricating 75 nm thick aluminum devices will meet this need while still maintaining the sensitivity and noise requirements of the GEP-B.

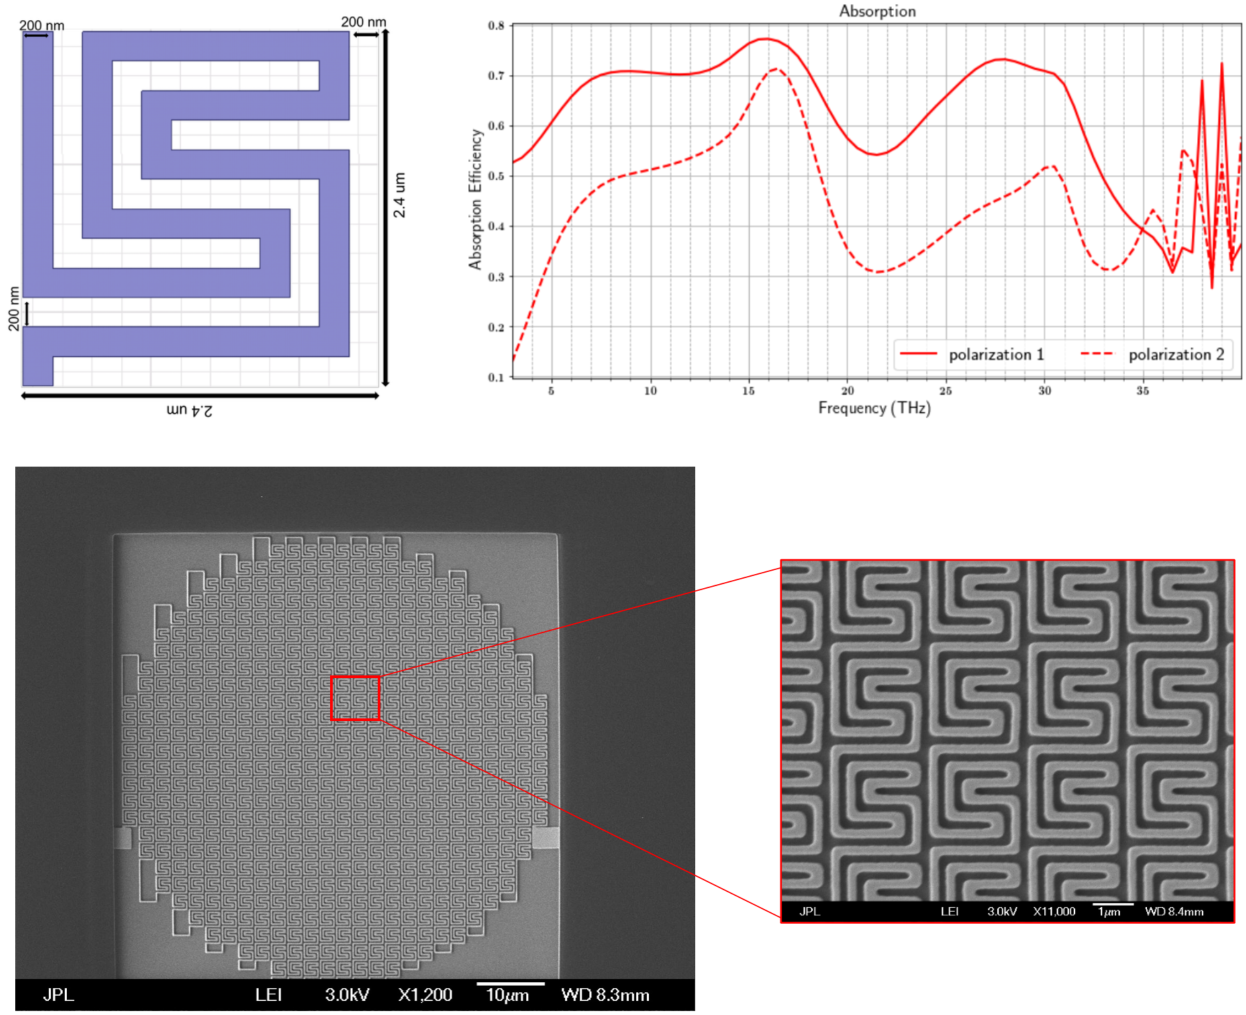

Figure 1.5: Upper Left: 10 μm KID inductor geometry. The purple represents Al with Z = (1.3 + 1.3i) Ω. The total unit cell size is 2.4 x 2.4 μm. Upper Right: HFSS-simulated absorption efficiency of 73% near 30 THz (λ = 10 μm) for absorber/inductor geometry on left. The absorber is back-illuminated. Bottom: Fabricated absorber of a KID concept for λ = 10 μm. The unit cell pattern simulated in Sonnet is repeated to cover entire absorber area, which is in the form of a circle with radius of $\approx$ 60 μm. Work completed by P. Day and J. Perido in 2019.

1.4.3 This Work: Detector Development for Future Far-IR Observatories

The extremely low NEP's required by these future space-based observatories (below 3×10^{-20} W/$\sqrt{Hz}$), along with their high pixel counts (10^5 to 10^6) at wavelengths from 10-1000 μm have not been demonstrated by any technology to date (see Section 1.3). This thesis presents work moving towards these required ultra-low NEP's. We focus on the development of KIDs due to their simple readout schemes that are well suited to large-scale cameras on space-borne missions. KIDs were baselined for both the GEP and GEP-B, and are therefore of direct import to this research group headed by the PI of those proposed missions.

We explore geometrical design parameters of KID detectors to inform future designs of ultra-low NEP KIDs. In Chapter 3, we present work on low metal volume detectors, an approach that should assist attaining low NEP's because volume is directly proportional to NEP. Chapters 4 and 5 focus on geometric design changes to elongate the signal duration in the detectors, since NEP is inversely proportional to this value. Chapter 4 investigates physically trapping the signal in the detector active region, while Chapter 5 quantifies the dependency of signal duration on metal thickness.

1.5 Layout of this Thesis

In Chapter 2 we cover the basics of the physics governing KIDs, as well as how we characterize their properties in the lab. We also discuss the essentials of theoretical KID noise. In the following chapters we explore geometrical design parameters of KID arrays to inform future designs of ultra-low NEP KIDs (see above, Section 1.4.3).

In Chapter 6 we analyze high-resolution observations of CO J = $3-2$ and $6-5$ of the nuclear region of NGC 6240 taken at the Atacama Large Millimeter Array (ALMA). Using these CO line observations, we model the kinematics of the molecular gas located between the nuclei of the progenitor galaxies. Our models suggest this gas is a tidal bridge linking the two nuclei that could fall onto the nuclei prior to second pass and feed future starbursts. We also find possible support for the presence of an AGN outflow or gravitational slingshot accelerating molecular gas away from the nuclear region. These findings shed light onto small-scale processes that can affect galaxy evolution and the corresponding star formation, with the tidal bridge depositing molecular gas onto the nuclei while other energetic processes push molecular gas further out of the nuclear region.

Chapter 2

KID Physics, Data Analysis, and Detector Readout

2.1 KID physics

Kinetic inductance detectors are circuits fabricated from superconducting metal, which can transmit DC current with no resistance when operated below their critical temperature T_c, the temperature at which the metal becomes superconducting. Therefore, to understand KIDs, a brief review of circuits and some superconducting physics is necessary.

2.1.1 The Gist: Circuits

KIDs are superconducting inductor-capacitor (LC) circuits whose resonance frequency f_0 is determined by

$$f_0 = \frac{1}{2\pi\sqrt{LC}}, \tag{2.1}$$

where L is the inductance and C is the capacitance of the circuit. A schematic of a KID's equivalent LC circuit is shown in panel B of Figure 2.1. To probe the KIDs' responses, a signal generator sends many frequencies ('tones') through a coaxial cable into the array of KIDs. Each KID absorbs its resonance frequency (Equation 2.1) and transmits frequencies far from resonance without change, resulting in an absorption feature like the solid line in panel c of Figure 2.1. The transmission of the signal through a detector is quantified via

$$S_{21} = \frac{V_{out}}{V_{in}}, \tag{2.2}$$

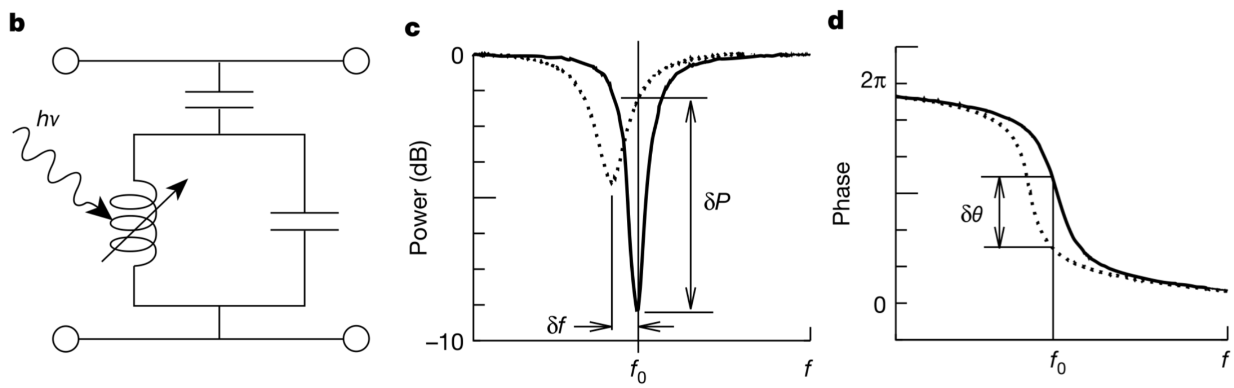

Figure 2.1: From Peter Day et al.'s nature paper Day et al. (2003), one of the first demonstration of KIDs detecting a single X-ray photon. **a**, omitted from this work. **b**, an equivalent circuit representation of a KID, with its inductor shown as variable due to absorbed light. The upper capacitor is the coupling capacitor, which allows the circuit's responses to be read out on the through lines. **c**, the power transmitted through the circuit measured in dB as a function of frequency. The KID absorbs its resonance frequency f_0, creating a dip. The solid line is the KID response before absorbing photons, and the dotted line is after absorbing photons. **d**, the relative phase of the signal transmitted through the line with respect to the input signal. Again, the solid line is the KID response before absorbing photons, and the dotted line is after absorbing photons.

which is often quantified in decibels (dB) and is further explained in Section 2.2.2. When a KID absorbs light, the shape and location of this absorption feature changes as described in Section 2.1.2.

The quality factor Q of a detector quantifies its response to the readout tones and is a statement of how effective the circuit is as a resonator, with higher Q's corresponding to deeper and narrower absorption features. Q is defined as the energy stored in the circuit divided by the energy dissipated per cycle of the driving frequency ω and can be calculated in two ways. First, as

$$Q = \frac{f_0}{\Delta f}, \tag{2.3}$$

where Δf is the width of the resonance feature S_{21} between the -3 dB points, which mark where the absorption is half of its maximum depth. This value can be readily quantified during in-lab tests by observation of the resonance feature. The second definition is

$$Q = \frac{\omega_0 L}{R}, \tag{2.4}$$

where $\omega_0 = 2\pi f_0$ is the angular resonance frequency, L is the inductance, and R is the resistance of the circuit. Practically, this definition is used less often in empirical KID characterization because L and R are theoretical quantities that are not readily measured during standard in-lab tests.

There are three quality factors that are important when understanding the properties of a KID. The first is the coupling quality factor Q_C that quantifies the losses that occur when coupling the KID response to the readout lines. This is often done via capacitive coupling as represented in panel B of Figure 2.1. The second is the internal quality factor Q_i, which quantifies the losses due to all other processes within the resonant circuit itself, for example from the resistivity of the metal film. Finally, the total quality factor Q_t that combines these two effects and is calculated with

$$\frac{1}{Q_t} = \frac{1}{Q_C} + \frac{1}{Q_i} \tag{2.5}$$

(Gao et al., 2007). The width of the resonance feature $S_{21}(f)$ is determined by this combined quality factor Q_t and can be roughly calculated with Equation 2.3.

2.1.2 Superconducting Physics

The capacitance in Equation 2.1 is dominated by the capacitor in the KID circuit, whose C is set during fabrication. There are two sources of inductance in the circuit: the geometric inductance L_{geom} that is determined by the physical geometry of the circuit and dominated by the pattern of the KID inductor. The second source is the *kinetic inductance* L_{ki} that comes from superconducting charge carriers called Cooper pairs. Cooper pairs are electrons of opposite spin bound together by an energy $2\Delta_0$ that can carry DC current with no resistance since they do not scatter off the superconductor's lattice. However, the Cooper pairs' inertia resists changes in electric field resulting from AC current – the very definition of inductance. This process occurs in normal-state metals as well but the resistive contribution to the impedance is much larger than the kinetic inductance. In superconductors, the ultra-low resistivity can result in L_{ki} dominating the impedance of the circuit. This occurs in kinetic inductance detectors by design. The contribution of L_{ki} to the total inductance $L_{tot} = L_{geom} + L_{ki}$ is quantified by the kinetic inductance fraction

$$\alpha = \frac{L_{ki}}{L_{tot}}. \tag{2.6}$$

It is important to note that L_{geom} and C are constants set by the geometry of the circuit while L_{ki} can change. This leads to the mechanism by which kinetic inductors function, described below.

The complex conductivity σ of a superconducting circuit measures its ability to conduct electric current and is defined as

$$\sigma = \sigma_1 - i\sigma_2, \tag{2.7}$$

where σ_1 is the inverse of resistance of the circuit and σ_2 is the inverse of the reactance X. The reactance describes how the phase of the current changes as it passes through the circuit. For KIDs it is provided mainly by the inductance, meaning $\sigma_2 = 1/X_L = 1/\omega L_{tot}$, where ω is the input frequency to the circuit. D. Mattis and J. Bardeen developed a theory in 1958 relating σ_1 and σ_2 to other properties of the superconductor (Mattis and Bardeen, 1958). σ_1 is linearly proportional to the the distribution function of quasiparticles (individual electrons created by the breaking of

Cooper pairs), given by the Fermi-Dirac distribution $f(E) = 1/(e^{E/kT}+1)$ when in thermal equilibrium. When operating at extremely low temperatures, σ_1 exponentially approaches zero due to this dependency (Zmuidzinas, 2012). σ_2 can be calculated using (Gao et al., 2007)

$$\frac{\sigma_2}{\sigma_n} = \frac{\pi\Delta_0}{hf}\left(1-\sqrt{\frac{2\pi k_B T}{\Delta_0}}e^{-\Delta_0/k_B T} - 2e^{-\Delta_0/k_B T}J_0(\xi)\right); \quad \xi = \frac{hf}{2k_B T}. \tag{2.8}$$

Here, σ_n is the normal state normal state conductivity, f is the microwave probe frequency, T is the temperature of the superconductor, and J_0 is a Bessel function of the first kind.

When light with energy $> 2\Delta_0$ (the binding energy of the Cooper pairs) is absorbed by superconductor, Cooper pairs can be broken into two individual electrons called 'quasiparticles'. In a KID this results in two effects:

(1) a decrease in quality factor

(2) a decrease in resonance frequency

The combined effects are shown in panel C of Figure 2.1 and are the mechanism by which we detect light with KIDs.

The first effect is because quasiparticles are resistive while Cooper pairs are not. Because Q is inversely proportional to the energy dissipated in the resonator, an increase in resistance will lead to a decrease in the quality factor. This is seen as an increase in the absorption feature's width and a decrease in its depth.

The second effect can be understood with Equation 2.1. Since L_{ki} is resultant from Cooper pairs, it is understandable that when light breaks them into quasi-particles that this portion of the inductance would change, thereby changing f_0. The relationship between L_{ki} and the number density of Cooper pairs n_s can be found by equating the kinetic energy of the Cooper pairs with the inductive energy of the circuit, and is given by

$$L_{ki} = \frac{m_e l}{2e^2 A}\frac{1}{n_s}. \tag{2.9}$$

Here, m_e is the mass of an electron, l is the length of the superconducting meander line, e is the charge of an electron, and A is the cross section of this line. When Cooper pairs are broken by incident light the kinetic inductance increases, shifting f_0 to the left by Δf according to

$$\frac{\Delta f}{f_0} = -\frac{\alpha}{2}\frac{\Delta L_{ki}}{L_{ki}}. \tag{2.10}$$

Intuitively, it may not be immediately obvious why decreasing the number of Cooper pairs will *increase* the kinetic inductance. Mathematically, we see L_{ki} is inversely proportional to the number density of Cooper pairs according to Equation 2.9. Intuitively we can understand this dependency with a thought experiment where we assume constant current. To maintain a constant current for a lower charge carrier density, each charge carrier must have a higher velocity. Higher velocity corresponds to larger momentum (due to their inertia) and, in turn, a higher kinetic inductance.

How much the detector's resonance frequency shifts as a function of absorbed power P_{opt} is measured by the responsivity

$$R_x = \frac{f_{res} - f_0}{f_0}\frac{1}{P_{opt}} = \frac{x}{P_{opt}}, \tag{2.11}$$

where $x \equiv (f_{res} - f_0)/f_0 = \frac{\Delta f}{f_0}$ is the fractional shift of the resonance frequency f_{res} under the optical load P_{opt} away from the resonance frequency under no optical load, f_0. The responsivity can also be related to properties of the detector through

$$R_x = \frac{\alpha\gamma\sigma_2\eta_o\eta_{qp}\tau_{qp}}{4N_0\Delta^2 V} \tag{2.12}$$

(Zmuidzinas, 2012). γ relates changes in σ to changes in the impedance Z of the circuit based on the thickness of superconductor. For our thin-filmed KIDs, $\gamma = 1$. η_o is the optical coupling coefficient, which quantifies the fraction of light incident on the detectors that is absorbed. The quasiparticle creation efficiency η_{qp} quantifies how many absorbed photons with $E > 2\Delta_0$ successfully break Cooper pairs into quasiparticles. τ_{qp} is the quasiparticle lifetime, how long it takes quasiparticles to recombine into Cooper pairs. N_0 is the single spin density of states, or how many quantum states of a certain spin are within a certain energy interval. Finally, V is the volume of the detector.

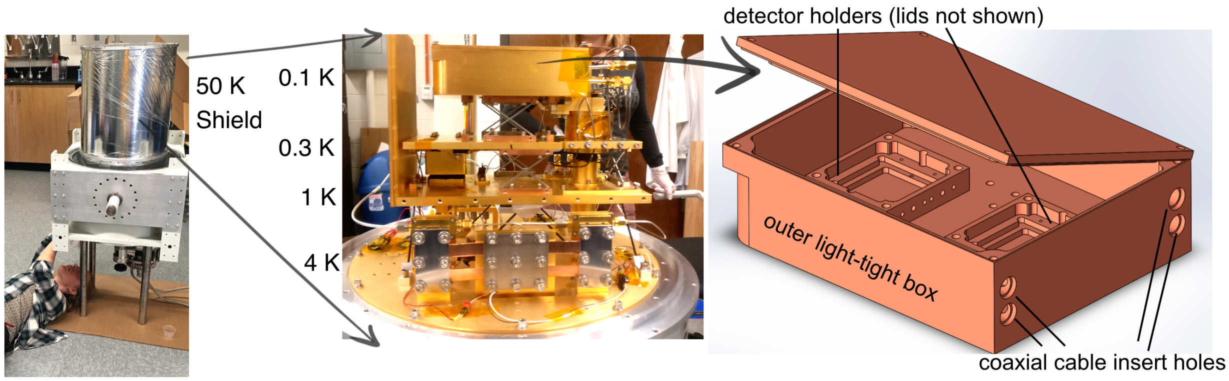

Figure 2.2: **Left:** The cryostat with its 300K vacuum shield removed, showing the first nested layer of decreasing temperature, the 50 K shield. Within this shield is a similar nested shield at 4 K. **Middle:** The internal workings of the cryostat, showing the 4 K plate, the partially disassembled 1 K box, the 0.3 K plate, and the detector stage (the 0.1 K box). **Right:** A zoom-in of the coldest stage, the 100 mK box-in-a-box setup attached to the adiabatic dilution refrigerator. This is the detector stage.

2.2 Data Analysis

2.2.1 CU Testing Facility

The superconducting KIDs need to be cryogenically cooled to below their critical temperature T_c at which they become superconducting. The best noise performance is achieved below $T_c/4$, usually a few hundred mK.

Our cryogenic testbed was initially designed by Matthew Hollister and its construction began in 2014. A schematic of the full testbed is shown in Figure 2.2. The main body of the cryostat is comprised of nesting aluminum shields held at decreasing temperatures to decrease optical loading on the coldest stages. Cooling to as low as 80 mK is achieved by three successively colder cryogenic cooling systems, described below.

The majority of the cooling is achieved by a closed helium compressor that cools from room temperature to 50 and 4 K. These stages are thermally sunk to the second and third nesting shield layers of the cryostat as well as copper plates that are used to mount and heat sink the cryostat's internal components. The second stage of cooling is from 4 K to $\sim$ 1 K and 0.3 K using a He sorption fridge from Chase cryogenics. The coldest temperature, as low as 80 mK, is attained using an adiabatic dilution refrigerator (ADR) that is attached to the detector stage. The ADR is made of a salt pill attached to a gold-plated copper finger surrounded by a superconducting solenoid magnet. The copper finger is physically connected via heat straps to the detector stage that contains the KID arrays to be tested.

The ADR's cooling is achieved by aligning the salt dipoles using the surrounding superconducting magnet powered with 10 amps, creating a $\sim$ 4 Tesla field. The salt pill is heat sunk to the 1 K stage with a heat switch while the magnetic field is powered to draw excess heat generated by dipole alignment out of the salt pill. Once the salt pill and the 1 K stage have equilibrated, the heat switch is disconnected and the salt pill and attached detector stage become completely thermally isolated from the warmer stages. At this point the magnetic field is slowly ramped down, allowing the salt dipoles to become unaligned and drawing heat out of the detector stage. When functioning at peak efficiency the cryostat can hold the detectors at 0.1 K for 8-10 hours.

The detector stage employs a box-in-a-box setup modeled after the dark, low-noise configuration at SRON (Baselmans et al., 2012). The box-in-a-box setup is rather self explanatory, with an inner gold-plated copper box that holds the detector array mounted to the base of an outer box that is designed to be light tight. The outer box is heat sunk to the salt pill's copper finger and employs lipped lids and coaxial feed-throughs for light tightness. It sits above the 0.3 K plate and inside of a 1 K box for additional light and heat shielding. The outer box is designed to be large enough to test up to four arrays at once and is currently configured for the testing of two arrays concurrently.

2.2.2 KID Readout Strategy

Figure 2.3 shows a block diagram of the single-tone readout setup used for measuring KID array properties at CU. A signal generator sends a single microwave drive tone of frequency f through a variable attenuator that allows the user to select the microwave drive power of the detector array. The signal is delivered to the array via coaxial cables that are heat sunk at 50 K, 4 K, and 1 K. At each heat sinking location the signal is also sent through an attenuator (3, 6, 10, or 20

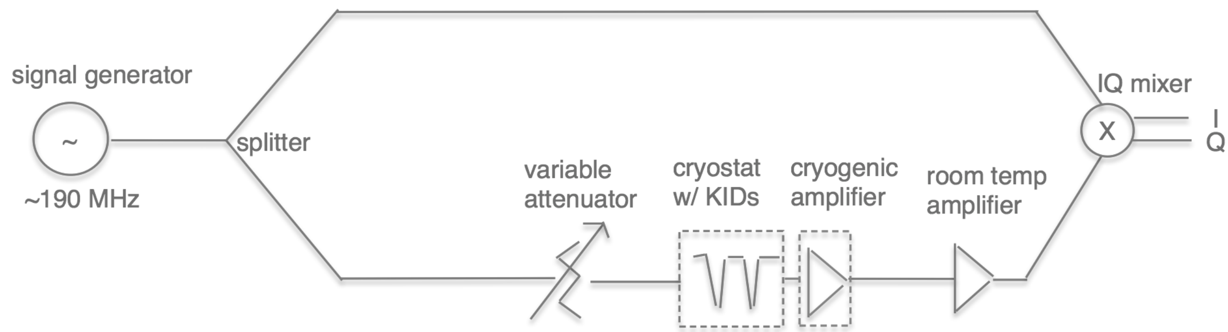

Figure 2.3: A schematic of the readout system, described in Section 2.2.2. A signal generator's tones are split into two lines, one of which is fed unchanged into the IQ mixer. The other line is sent through a variable attenuator, the cryostat containing the array of KIDs, a cryogenic and a room temperature amplifier, and then mixed with the input signal in the IQ mixer.

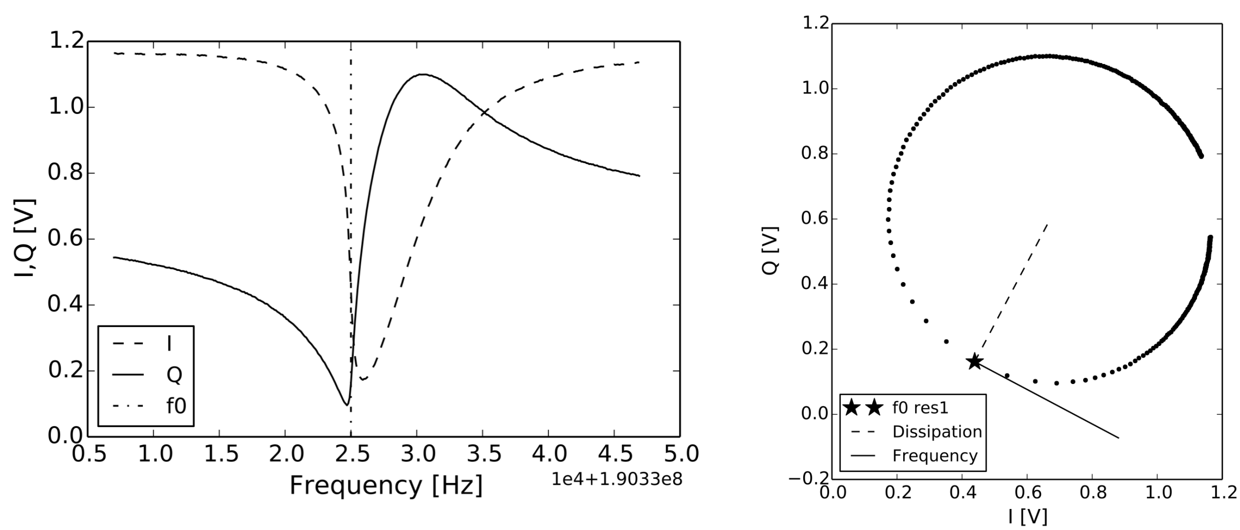

Figure 2.4: *Left:* The in-phase (I) and quadrature (Q) components as a function of frequency, as described in Section 2.2.2. The dotted-dashed vertical line marks where the resonance frequency f_0 is measured to be. *Right:* Q(f) plotted against I(f), known as a resonance circle. The input frequency is what changes from one point to the next. The radial (dissipation) and tangent (frequency) directions to the circle are shown, as described in Section 2.2.2.

dB) that reduces signal reflections at the heat sink locations. The array output is amplified by both cryogenic and room temperature amplifiers, then mixed with the original signal in an IQ mixer.

The IQ mixer breaks the detector array output ($S_{21}(f)$) into the components that are in phase (I) and out of phase (in "quadrature" Q) with the original signal from the signal generator. The total transmission through the system is $S_{21}(f) = \sqrt{I(f)^2 + Q(f)^2}$. I and Q are each plotted as a function of microwave readout frequency in the left panel of Figure 2.4. The shapes of these signals can be understood by considering the phase difference of the attenuated signal with respect to the input signal for two frequency limits in an ideal KID. Far off resonance the KID cannot absorb the power, and the signal passes unaltered through the resonator. Therefore, the output signal is either 2π or 0 radians out-of-phase (that is, fully in-phase) with the input and the entire signal is in I, which is at a maximum. Conversely, when exactly on resonance the signal is highly attenuated by the resonator meaning that the attenuated signal is 180° (π radians) out of phase, corresponding to a minimum of I and a Q of 0. The phase as a function of frequency is shown in panel d of Figure 2.1, displaying this smooth change from 2π to 0 radians with a maximal slope directly on resonance (π radians out of phase).

Plotting Q(f) vs. I(f) results in a resonance circle shown in the right panel of Figure 2.4. The quantity that changes around the perimeter of the circle is frequency. Therefore, the direction tangent to the circle is referred to as the frequency (or phase) direction. The direction perpendicular to the circle is referred to the amplitude (or dissipation) direction because the radius of the circle roughly corresponds to the depth of the absorption feature $|S_{21}|$. The resonance circle distorts as the microwave tone driving power increases because the kinetic inductance is nonlinearly dependent on current. The current is maximized on-resonance, slightly increasing the kinetic inductance and decreasing the resonant frequency. When the drive power is too high, the correspondingly high current further decreases the kinetic inductance, decreasing the resonance frequency and therefore increasing the current, leading to a runaway feedback loop (Swenson et al., 2013). This is called detector bifurcation and shows up as a sudden jump in phase in the resonance circle at f_0 at a high microwave drive power.

The resonance frequency f_0 does not correspond to a minimum in I or Q in Figure 2.4 or a minimum in $S_{21}(f)$, though one would expect maximal absorption of the signal at the resonant frequency. This shift is caused by an overall phase shift between the input signal that travels directly from the signal generator to the IQ mixer and the signal that goes through the cryostat and then into the IQ mixer. The resonance frequency is instead located at the location of maximal rate of phase change around the resonance circle with frequency. This corresponds to the maximum spacing of points on the resonance circle shown in the right panel of Figure 2.4, as well as the steepest slope of phase vs frequency in panel d of Figure 2.1.

2.2.3 Correction and conversion of data streams to fractional frequency shifts

To measure a KID's response to light (or to characterize its noise in a dark environment), a few steps are taken that are expanded upon below:

- Take a wide sweep of $I(f)$ and $Q(f)$ over a frequency range extending at least three-four times the width of the resonance feature
- Take a sweep of $I(f)$ and $Q(f)$ that finely samples the resonance feature
- Plot the finely sampled $Q(f)$ vs. $I(f)$ and determine f_0 based on the maximum spacing between data points
- Using the signal generator, provide a microwave tone of $f = f_0$ and stream $I(t)$ and $Q(t)$
- Correct $I(t)$ and $Q(t)$ for systematic effects introduced by the readout system
- Convert $I(t)$ and $Q(t)$ to phase and/or fractional frequency shift with time to prepare for data analysis.

We stream $I(t)$ and $Q(t)$ at $f = f_0$ because, as stated above, this is the point at which the resonator is the most responsive to changes in the system (the slope of phase vs frequency is steepest at this location).

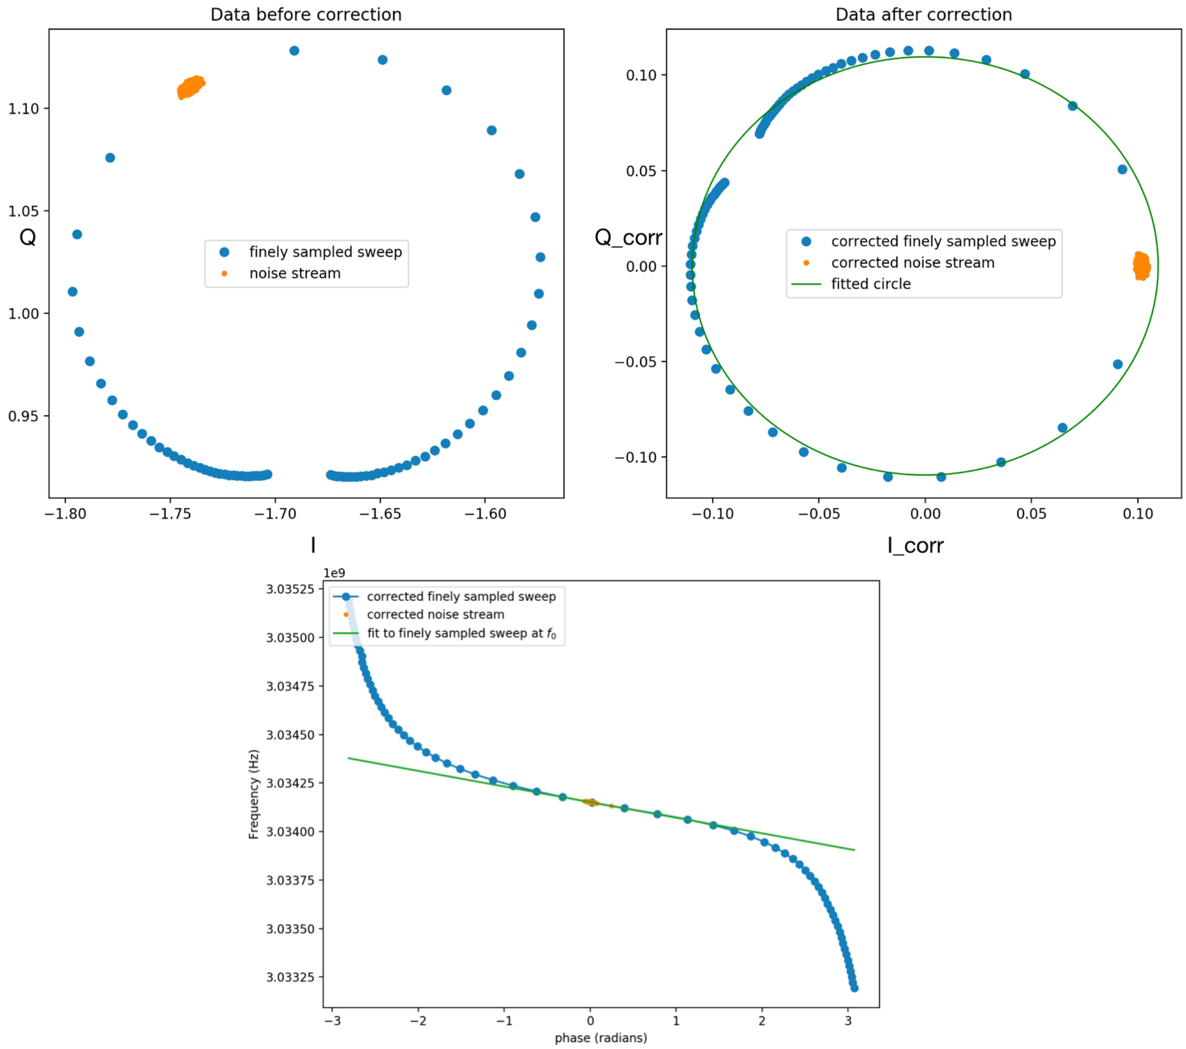

Figure 2.5: **Upper Left:** the raw data of the finely sampled resonance feature (*orange*) and the raw noise $Q(t)$ vs $I(t)$ (*blue*). **Upper Right:** the corrected finely sampled resonance feature (*blue*) and the corrected noise (*orange*) along with the fit circle that moves the resonance circle's center to (0,0). **Bottom:** the corrected finely sampled resonance feature from the upper right, but plotted as frequency vs phase ($\tan^{-1}(Q_{corr}/I_{corr})$; *blue*), *orange* is the same but for the corrected noise stream from the upper right hand plot, and *green* is the linear fit to frequency vs phase to convert the data between phase and frequency. **change color scheme of plot to match**

The systematic effects that alter I and Q as they pass through the readout system, as well as the techniques employed to correct these alterations, are well documented in Jiansong Gao's thesis (Gao et al., 2007). We briefly present them here as they pertain to the corrections we apply to streamed data.

First, we remove the slow variation of gain with frequency often caused by impedance mismatches in the system or frequency-dependent gain intrinsic to amplifiers in the system. This is completed with a polynomial fit of $S_{21}(f)$ for the wide sweep of $I(f)$ and $Q(f)$ that is then removed from the data.

We then remove the cable delay, the overall phase shift between the input signal that travels directly from the signal generator to the IQ mixer and the signal that goes through the cryostat and then into the IQ mixer. The cable delay is geometric, meaning we expect it to affect all frequencies equally and its affect is to rotate the resonance frequency off of the Q axis, with the amplitude direction fully in the I plane and the phase/frequency direction fully in the Q plane. We measure the phase shift using a linear fit to the wide frequency sweep's phase ($\tan^{-1}(Q(f)/I(f))$) vs frequency, and rotate the finely sampled resonance feature and the noise accordingly. The final systematic effect that we remove is the DC offset that simply moves the resonance circle's center away from (0,0) and is corrected simply by moving the resonance circle's center to (0,0). The result of the correction process is shown in the top two panels of Figure 2.5.

The result of these corrections is a round, finely sampled resonance feature in IQ space centered at the origin with the corrected noise $Q_{corr}(t)$ vs $I_{corr}(t)$ lying on the I axis (Gao et al., 2007). Using the corrected finely sampled resonance feature the conversion between phase $\theta = tan^{-1}(Q_{corr}(f)/I_{corr}(f)$ and frequency is found with a linear fit at the resonant frequency f_0, the frequency at which we took the noise stream. This is shown in the bottom panel of Figure 2.5. Finally, the fit line is used to convert the data streamed at $f = f_0$ from phase $\theta = tan^{-1}(Q_{corr}(t)/I_{corr}(t)$ to frequency $f(t)$. The fractional shift away from the resonance frequency $\Delta f/f_0 = (f(t) - f_0)/f_0$ over time can be calculated, a standard quantity for measuring a KID's response to light or its environment over time.

2.2.4 Q fitting

The resonator's complex transmission $S_{21}(f)$ near resonance is

$$S_{21} = \Big(1 - \frac{Q_t}{Q_c}\frac{e^{j\phi}}{1+2jQ_t x}\Big), \tag{2.13}$$

(Gao et al., 2007) where x is the fractional frequency shift away from resonance, Q_t is the quality factor of the resonator, and Q_c is the coupling quality factor as defined in Section 2.1.1. ϕ is the phase shift between the input signal that travels directly from the signal generator to the IQ mixer and the signal that goes through the cryostat and then into the IQ mixer described in Section 2.2.2.

Equation 2.13 describes the response of a linear resonator, but as the resonator is driven at higher power its response becomes nonlinear as it approaches bifurcation. Swenson et al. (2013) derive the implications on S_{21} of this nonlinear effect caused by a nonlinear kinetic inductance. The microwave power P sent into the resonator is related to the dissipated microwave power via $P_{diss} = P(1-|S_{11}|^2 - |S_{21}|^2)$, with $S_{11} = S_{21} - 1$ quantifying the reflected microwave power. Carrying out the algebra, they find that the nonlinear frequency shift scaled by the quality factor $y = Q_t x$ is related to the scaled frequency shift for a linear resonator y_0 via $y_0 = y - \frac{a}{1+4y^2}$, where a is a nonlinearity parameter between 0 and 1. Finding the roots of this equation yields the nonlinear y that is used in the stead of $Q_t x$ in the denominator of Equation 2.13 to fit a detector when it is close to bifurcation. For our resonators driven at a few dB below bifurcation, this nonlinear y is used to fit $S_{21}(f)$ and quantify the quality factors of our resonators.

2.2.5 Finding T_c and α

To constrain T_c and α of our detectors, we first find the fractional resonant frequency shift $\Delta f / f_0$ as a function of detector temperature and fit it using

$$\frac{\Delta f}{f} = 0.5\alpha\Big(\frac{\sigma_2 - max(\sigma_2)}{max(\sigma_2)}\Big), \tag{2.14}$$

(Zmuidzinas, 2012). Because σ_2 depends on $\Delta_0 \approx 1.76 k_B T_c$ (T $<<$ T$_c$) this measurement will simultaneously constrain both T_c and α.

T_c can also be measured directly using four-wire resistance measurements of the array (or an equivalent metal sheet) and observing where the resistance of the metal decreases to nearly zero. This is useful as it avoids the intrinsic degeneracy between T_c and α when fitting the above equation. However, direct measurements in this way disallow other tests to be conducted on the array while it is cooled.

2.2.6 Noise Streaming and Power Spectral Densities

Noise measurements that are converted into noise spectral densities S are taken for a single detector by providing a microwave input tone with $f = f_0$ and measuring the resultant $I(t)$ and $Q(t)$, as explained in Section 2.2.2. $I(t)$ and $Q(t)$ are then broken into their frequency and dissipation directions. The total noise in the dissipation (or frequency) direction can be found by adding the components of $I(t)$ and $Q(t)$ in the dissipation (or frequency) direction together. The power spectral densities (PSDs) of these noise quantities are defined as the power spectrum of their Fourier transforms,

$$\begin{aligned} S_{freq} &= |FT(freq(t))|^2 \\ S_{diss} &= |FT(diss(t))|^2, \end{aligned} \tag{2.15}$$

and are in units of V^2/Hz. Different sources of noise contribute to these noise components differently. An example PSD is plotted in Figure 2.6. This figure shows on-resonance noise spectra (input frequency = f_0) and off-resonance noise spectra. The off-resonance noise characterizes the noise from the system, meaning any excess noise of the on-resonance measurement above the off-resonance originates from the detector array. Noise that is taken in an environment without light shining on the detectors is referred to as “dark noise”.

The PSDs must be converted from noise power in V per frequency bin (V^2/Hz) to noise power in Hz per frequency bin. This is because, ultimately, shifts in V are not what we measure in our readout system, but rather shifts in frequency space (i.e. we measure $\Delta f/f_0$ as described in Section 2.2.3). This conversion is completed by scaling the PSDs by the readout voltage’s

Figure 2.6: An example power spectral density in both the frequency (*black*) and dissipation (*orange*) directions, as defined in Equation 2.15. The PSD's from noise streamed off-resonance are plotted as dashed lines. The detector noise roll off in the frequency direction is marked with a vertical pink line.

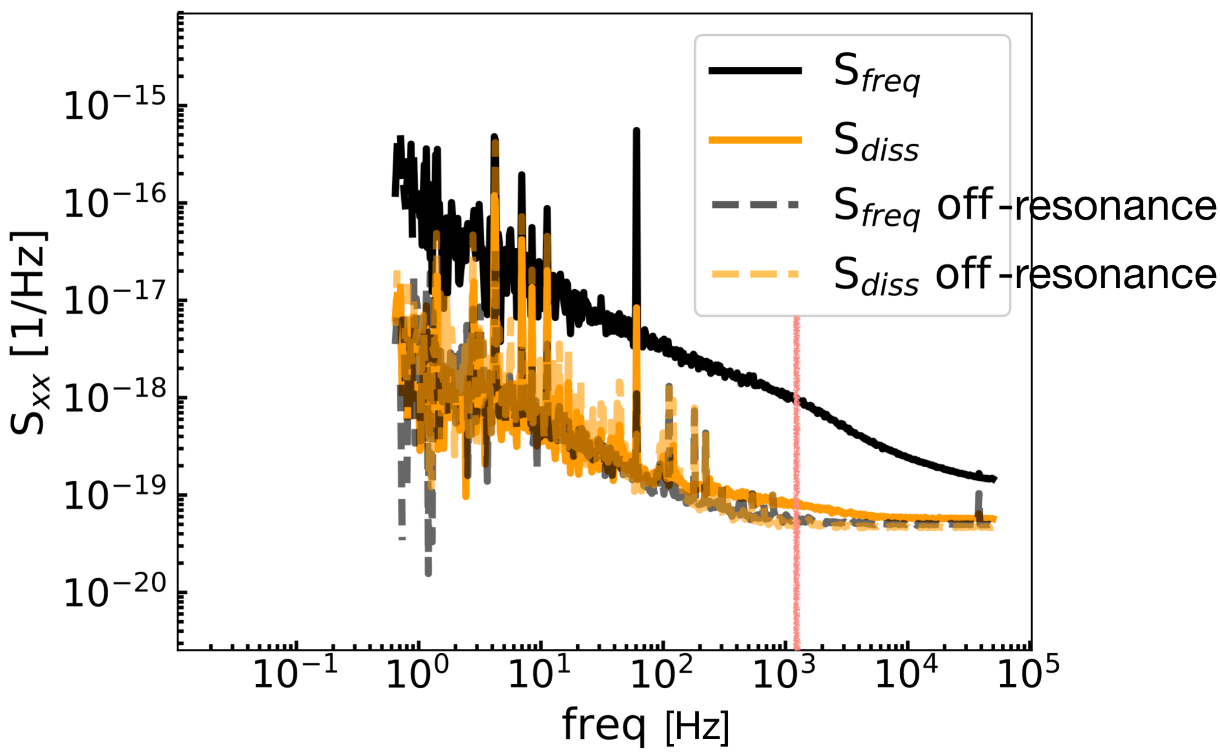

dependence on microwave frequency. Practically, this is done by dividing the PSDs by the $|\text{slope}|^2$ of the voltage $V(\vec{f})$ of I and Q near resonance,

$$\begin{aligned} \left|\frac{d\vec{V}}{df}\right|^2 &= \left|\frac{d}{df}(\vec{Q}+\vec{I})\right|^2 \\ &= \left|\frac{dI}{df}\hat{I}+\frac{dQ}{df}\hat{Q}\right|^2. \end{aligned} \tag{2.16}$$

I and Q both change linearly with frequency near resonance. Therefore, their derivatives are the slopes with frequency (m_I and m_Q, respectively) and

$$\left|\frac{d\vec{V}}{df}\right|^2 = m_I^2 + m_Q^2, \tag{2.17}$$

where both m_I and m_Q come from line fits to $I(f)$ and $Q(f)$ near resonance. By dividing the PSDs (units of V^2/Hz) by this conversion factor $\left|\frac{d\vec{V}}{df}\right|^2$ we have S_{freq} and S_{diss} in units of 1/Hz.

2.2.7 Quasiparticle Recombination Time From Noise PSD

The roll off marked with a vertical line in Figure 2.6 is a signature of detector noise PSDs and is Lorentzian in nature. A Lorentzian roll off can be caused by readout system properties or by the detector bandwidth $f_0/2Q$. It can also be caused by the detector's inability to react faster than its slowest timescale. One possible timescale is the resonator ring time τ_{res} that is intrinsic to the LC circuit's quality factor Q_{tot}. This ring time is calculated with

$$\tau_{res} = \frac{Q_{tot}}{\pi f_0} = \frac{1}{\pi \Delta f} \tag{2.18}$$

and is a few tens of μs for the resonators studied throughout this body of work. This timescale reflects that the more energy a circuit can store, the larger its Q_{tot} and the longer it takes to 'relax' (dissipate its energy). The other dominant timescale in the system is the quasiparticle recombination time τ_{qp}. The longest timescale in the system can be found by fitting a Lorentzian$L(f)$ to the rolloff

$$L(f) = \frac{A}{1+(2\pi\tau f)^2}, \tag{2.19}$$

where A is an amplitude in V^2/Hz, f is the frequency, and τ is the longest relaxation timescale in the system.

2.2.8 Noise Equivalent Power

The oft-used measure used to quantify the level of noise present in KIDs is the noise equivalent power (NEP), defined as

$$NEP = \frac{\sqrt{S}}{R_x} W/\sqrt{Hz}, \tag{2.20}$$

where $\sqrt{S}$ is the noise spectral density in 1/Hz, defined above, and R_x is the responsivity as defined in Equation 2.12 in Hz/W. An equivalent definition of NEP is the power at which the signal to noise ratio is one over a bandwidth of 1 Hz.

2.3 KID Noise

NEP contributions from different sources of noise can be theoretically derived, and each adds in quadrature to the overall NEP of the system. Fundamental sources of noise and their contribution to the NEP are described below.

2.3.1 Generation-Recombination (G-R) Noise

The most fundamental source of noise within a superconductor comes from the random generation and recombination of quasiparticles. Generation happens from thermal fluctuations and hence decreases exponentially with the operating temperature of the KID. The NEP contribution of generation-recombination noise is

$$NEP_{g-r} = 2\Delta\sqrt{\frac{N_{eq}}{\tau_{qp}}}, \tag{2.21}$$

(Zmuidzinas, 2012), where N_{eq} is the equilibrium number of quasiparticles in the detector and τ_{qp} is the lifetime of a quasiparticle before recombination. G-R noise shows up as white noise, contributing equally to all frequencies below the frequency corresponding to the quasiparticle recombination time τ_{qp}.

2.3.2 Photon Noise

Photons contribute random noise to the system, with an NEP contribution of

$$NEP_{photon} = \sqrt{2P_o h\nu(1+n_o)}, \tag{2.22}$$

(Boyd, 1982) where P_o is the optical power, ν is the frequency of the incoming radiation, and n_o is the occupation number of the photons. Higher optical powers or photon energies contribute more noise. n_o accounts for the fact that photons are bosons and tend to bunch and arrive together.

2.3.3 Two Level System Noise

Another common contributor to KID noise is two level system (TLS) noise that comes not from the superconducting circuit but from the substrate upon which it sits. Incomplete bonds between atoms or groups of atoms in this substrate (in our case silicon) lead to quantum states that are close to each other in energy. These quantum states are each local minima, with one state slightly higher than the other in energy. Thermal energy can cause atoms or groups of atoms to tunnel between these states, changing the dielectric constant ε_r of the substrate and altering the capacitance of the circuit. This change in capacitance is due to the fact that $C \propto \varepsilon_r$. f_0 in turn changes according to Equation 2.1, while the Q of the detector remains almost unchanged. This means that TLS noise is almost entirely in the frequency direction.

TLS noise's power spectral density S_{TLS} has a 1/f shape (Gao et al., 2007) due to spectral diffusion of energy between the TLS states caused by spin-spin coupling (Black and Halperin, 1977). This spectral diffusion causes a logarithmic broadening of the TLS energy states in time. S_{TLS} at a given frequency decreases with microwave drive power according to $S_{TLS} \propto P_{microwave}^{-1/2}$. Physically this occurs because the higher drive power suppresses TLS noise by saturating the atoms in their higher energy state and prevents them from oscillating. Similarly, higher detector temperatures suppress TLS noise because the thermal energy saturates atoms in their higher energy states, presenting as $S_{TLS} \propto T^{-2}$.

TLS noise can be mitigated in the design process with the use of an interdigitated capacitor (IDC, pictured on the right in Figure 3.1). The many tines of the IDC reduce the electric field density by increasing the capacitor's area, thereby decreasing the energy available to assist the electrons in tunneling between states (Gao et al., 2007).

Chapter 3

Low Volume detectors

3.1 Abstract

Arrays of tens of thousands of sensitive far-infrared detectors coupled to a cryogenic 4 - 6 meter class orbital telescope are needed to trace the assembly of galaxies over cosmic time. The sensitivity of a 4 Kelvin telescope observing in the far-infrared (30-300 μm) would be limited by zodiacal light and Galactic interstellar dust emission, and require broadband detector noise equivalent powers (NEPs) in the range of 3×10^{-19} W/$\sqrt{Hz}$. We are fabricating and testing 96 element arrays of lumped-element kinetic inductance detectors (LEKIDs) designed to reach NEPs near this level in a low-background laboratory environment. The LEKIDs are fabricated with aluminum: the low normal-state resistivity of Al permits the use of very thin wire-grid absorber lines (150 nm) for efficient absorption of radiation, while the small volumes enable high sensitivities because quasiparticles densities are high. Such narrow absorption lines present a fabrication challenge, but we deposit TiN atop the Al to increase the robustness of the detectors and achieve a 95% yield. We present the design of these Al/TiN bilayer LEKIDs and preliminary sensitivity measurements at 350 μm optically loaded by cold blackbody radiation.

3.2 Preface: Division of Work

Fabrication and testing of devices was done at the Jet Propulsion Laboratory (JPL) in Pasadena, CA, while data analysis was conducted primarily at CU Boulder. Fabrication was done by Henry Leduc at the microdevices laboratory. Testing was done by Byeong Ho Eom and Peter Day with as-

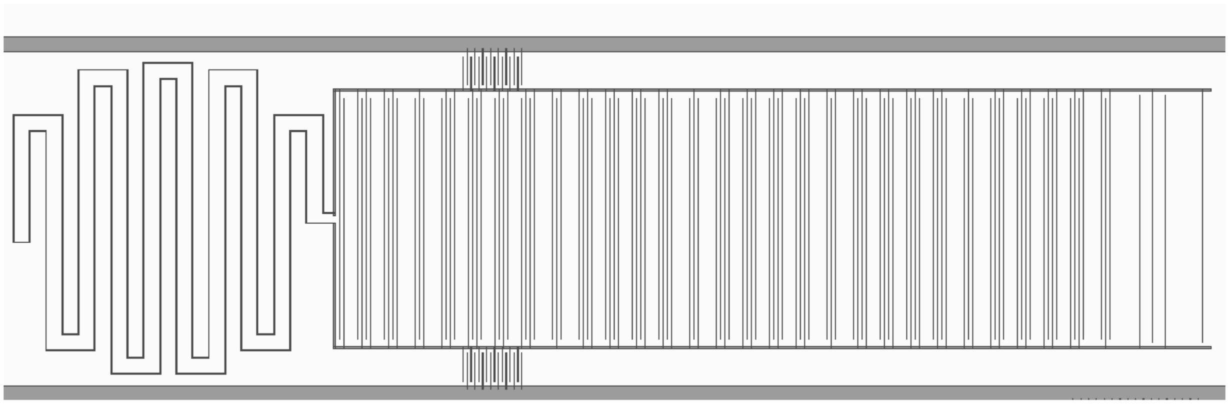

Figure 3.1: A single detector. The inductor (which is also the absorber) is the meandered structure in an approximately circular pattern on the left. The structure on the right is the interdigitated capacitor, which extends past the edge of the image. For clarity, not all of its tines are shown in the image.

sistance for one weekend's worth of data by Adalyn Fyhrie. The data analysis chain was developed primarily by Adalyn Fyhrie under the direction of Jason Glenn and Peter Day.

3.3 Introduction

In this work, we quantify the superconducting detector parameters and noise properties of one KID array that was fabricated at JPL as a first attempt at attaining the low NEPs needed by the aforementioned Surveyor. This array was designed to have a low NEP by using extremely small volumes of superconducting metal (explained in Section 3.4), a technique that was only recently attainable due to improved fabrication capabilities at JPL.

3.4 Detector design

Our approach for achieving low NEP (a main goal of this project) is to minimize the detector's active volume, to which the NEP is directly proportional according to Equations 2.12 and 2.20. The active volume V is computed by

$$V = wlt = \frac{A\rho}{Z}, \tag{3.1}$$

where w is the width of the active area of the resonator microstrip lines, l is the total length of the lines, t is the thickness of the lines, ρ is the normal state resistivity of the superconductor, A is the active detector area, and Z is the impedance that is fixed to match the silicon substrate. This impedance matching reduces the reflection of incident light, which must first pass through the Si before reaching the detector in our usual back-illuminated case. A is minimized by coupling the detectors to incoming radiation with a silicon lenslet array, making A (and, in turn, l) completely determined by the diffraction spot size of the lenslets. The lenslet array consists of many small spherical lenses, each of which focuses light onto a detectors' inductor that doubles as the absorber. The devices were fabricated from aluminum due to its low ρ, which yields a smaller V for this fixed A while still impedance matching with the silicon substrate.

One consequence is that the 150 nm wide Al microstrips necessitated by the low volume are extremely delicate. Therefore, to increase the robustness of the array we deposit TiN on top of the Al to create a more resilient bilayer. The TiN serves a dual purpose: it is much more durable than Al and thus is less likely to be damaged during the fabrication process, and it also prevents the Al from oxidizing. In principle, the superconducting properties of Al should still dominate the response of the detectors because its low ρ causes it to 'short out' the TiN.

An array of 96 bilayer detectors was fabricated at the Microdevices Laboratory at JPL on a Si substrate, with a detector pitch of 1 mm. Each detector (shown in Fig. 3.1) consists of a meandered inductor (that doubles as the absorber) with a roughly circular profile that allows for effective optical coupling to the lenslet array and an extended IDC, the virtues of which are explained in Section 2.3. The microstrip line widths are 150 nm with a total thickness of 40 nm (20 nm of Al below 20 nm of TiN). A yield of 95% was achieved, demonstrating the efficacy of using TiN on top of Al to create resilient devices. The resonances were designed to be between 190 and 360 MHz with quality factors between 8×10^3 and 8×10^4.

3.5 Detector Characterization

Five measurements were obtained to characterize our detectors, each of which is enumerated below, with the data analysis and results described in the following section. For tests under optical load, the cryostat contained a blackbody separated from the detectors by two 350 μm bandpass filters, a 58 cm^{-1} low pass filter, and a Zitex filter. The coupling efficiency of the light to the detector ε_{opt} quantifies the the fraction of light that is available to be absorbed by the KIDs after passing through these filters and the lenslet array. The transmission of each component is: bandpass filters, 0.87^2 (two filters). Lowpass filters, 0.9. Zitex filter, 0.9. Lack of antireflection coating on the lenslet array, 0.7. Lack of backshort on the detectors, 0.7. Multiplying these together, the total ε_{opt} is found to be 0.3.

(1) A sheet of the Al/TiN bilayer was fabricated with the same thickness as the detectors (20 nm Al/20 nm TiN). Its T_c was determined using a 4-wire resistance measurement.

(2) Frequency sweeps with a vector network analyzer (VNA) yielded 5 detectors' resonance frequency shifts in response to detector temperatures between 0.02 and 0.8 K in a light-tight environment. This measurement yields information on the detectors' kinetic inductance fraction, $\alpha = L_{kin}/L_{tot}$.

(3) Noise measurements that yield NEP and quasiparticle recombination time τ_{qp} were taken at a bath temperature of 0.2 K. These were repeated for four signal attenuations and blackbody temperatures of 40 K, 28 K, 16 K, and 4.2 K, which correspond to calculated optical loads of 3.8 pW, 2.0 pW, 0.57 pW, and 0.4 fW.

(4) Measurements of $S_{21}(f)$ and the detectors' quality factors were made via frequency sweeps with a VNA at each of the attenuations and optical loads listed above.

(5) The optical responsivity was quantified for 5 detectors using a VNA by measuring 5 detectors' resonance frequency shifts in response to blackbody temperatures between 4 and 40 K.

Quantity	Measurement	Quantity	Measurement
Δ_0	2.87×10^{-4} eV	α	0.9 ± 0.05
Q_i	$1.7\times10^4 \pm 300$	T_c	1.9 K
Q_t	$9.3\times10^3 \pm 200$	τ_{qp}	$100 \pm 10\ \mu s$
Q_c	$2\times10^4 \pm 500$		

Table 3.1: Detector parameters for the bilayer array. Δ_0 is the gap parameter, Q_i the internal quality factor, Q_t the total quality factor, Q_c the coupling quality factor of the detector to the readout line, T_c the critical temperature, and τ_{qp} the quasiparticle lifetime.

3.6 Data Analysis and Results

3.6.1 Detector Parameter Results

Measuring Critical Temperature: The critical temperature T_c of our film was measured to be 1.9 K, corresponding to a $\Delta_0 = 1.76 k_B T_c = 2.87 \times 10^{-4}$ eV Gao et al. (2007). This T_c is approximately twice what we were expecting.

Measuring Quality Factors: Quality factors are found from fitting data from measurement 4 using Equation 4.5. An example of one of these fits is shown in Figure 3.2. Table 3.1 reports the Q's that correspond to our best NEP measurements (shown in Figure 3.7).

Measuring Kinetic Inductance Fraction: To find the α of our detectors, we fit the data from measurement 2 with Equation 5.13. For this measurement (shown in Figure 3.3a), we fix Δ_0 to its measured value of 2.87×10^{-4} eV and fit only for α with a χ^2 minimization algorithm. We do this for five different detectors, and find an average α of 0.9 for the five resonators explored in depth in this study.

A $\Delta\chi^2$ plot in which both Δ_0 and α are allowed to vary is shown in Figure 3.3b. The contours display a high correlation between the two variables - by fixing Δ_0 to its measured value we constrain α to a certain value on the correlation line. Interestingly, the global minimum χ^2 value yields unphysical α's greater than one.

Measuring Responsivity: Figure 3.4 shows the temperature sweep results (measurement 5 from Section 3.5), which is fit with Equation 2.12 to determine the responsivity of the detectors.

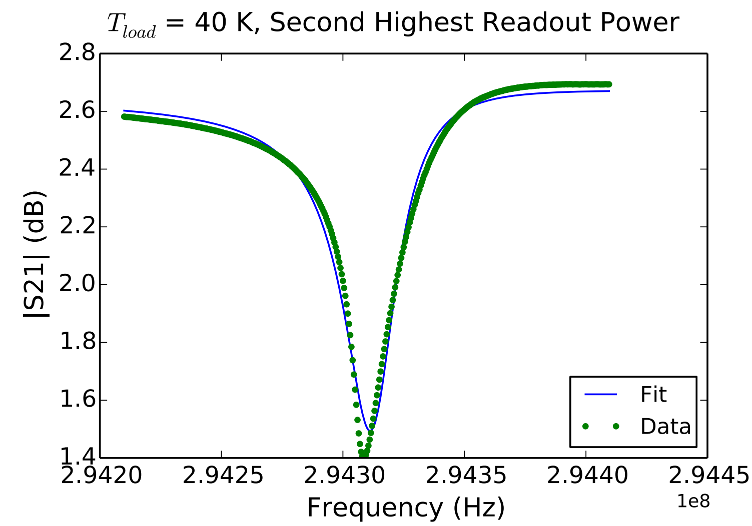

Figure 3.2: A fit to $|S_{21}(f)|$ for the highest blackbody temperature of 40 K, and the second highest readout power.

Optical Power [pW]	R_x [W^{-1}]	R_{diss} [W^{-1}]
3.8	0.63×10^8	4.85×10^6
2.0	0.77×10^8	5.92×10^6
0.6	0.99×10^8	7.62×10^6
4×10^{-4}	1.14×10^8	8.77×10^6

Table 3.2: For each of the calculated optical loads, the measured responsivity in the frequency direction R_x and the responsivity in the dissipation direction R_{diss}, as described in Chapter 2.

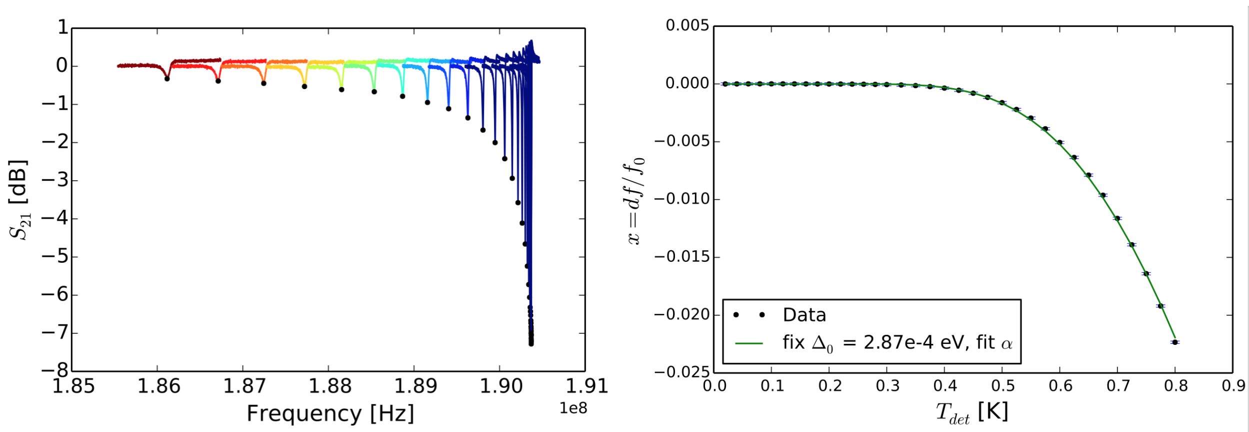

(a) *Left*: Results from measurement 2 from Section 3.5. The response of one detector for detector temperatures T_{det} between 0.02 K (*blue*, the rightmost curve) and 0.8 K (*red*, the leftmost curve). Resonance frequencies for each T_{det} are shown as black dots. *Right*: The data from the lefthand panel translated to resonance frequency shifts away from the lowest bath temperature $T_{det} = 0.02K$ as a function of T_{det}.

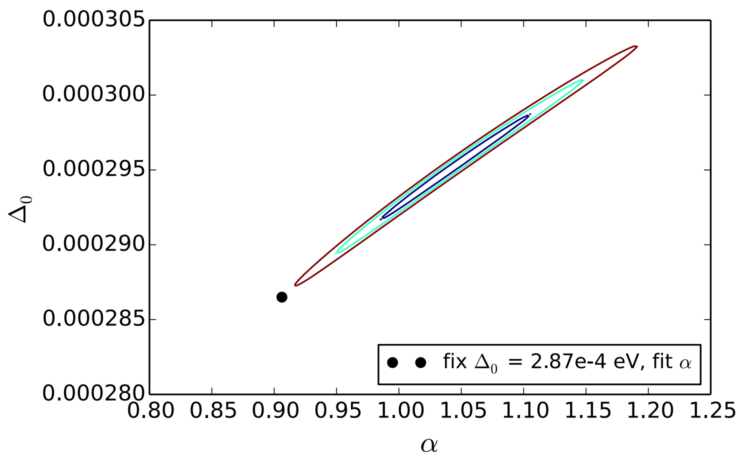

(b) $\Delta\chi^2$ contours for the fit to the data above. The contours drawn are $\Delta\chi^2 = 2.3$, 6.2, and 11.8 corresponding to 1, 2, and 3σ confidence levels for 2 degrees of freedom. The black dot corresponds to the fit shown in Figure 3.3a, for which Δ_0 is fixed and a minimum χ^2 value of α is fit. This fit sits on the same line that the $\Delta\chi^2$ contours make, showing a high correlation between the two variables. The global minimum χ^2 gives an unphysical $\alpha > 1$, a behavior that is not yet understood. However, a prior is applied enforcing $\alpha < 1$ for all interpretations of the data in the text. Additionally, the minimum χ^2 value of α for the fixed value of Δ_0 is < 1, indicating the model results in physically possible values when Δ_0 is constrained according to independent measures of T_c.

Figure 3.3: Measurement of α for one detector.

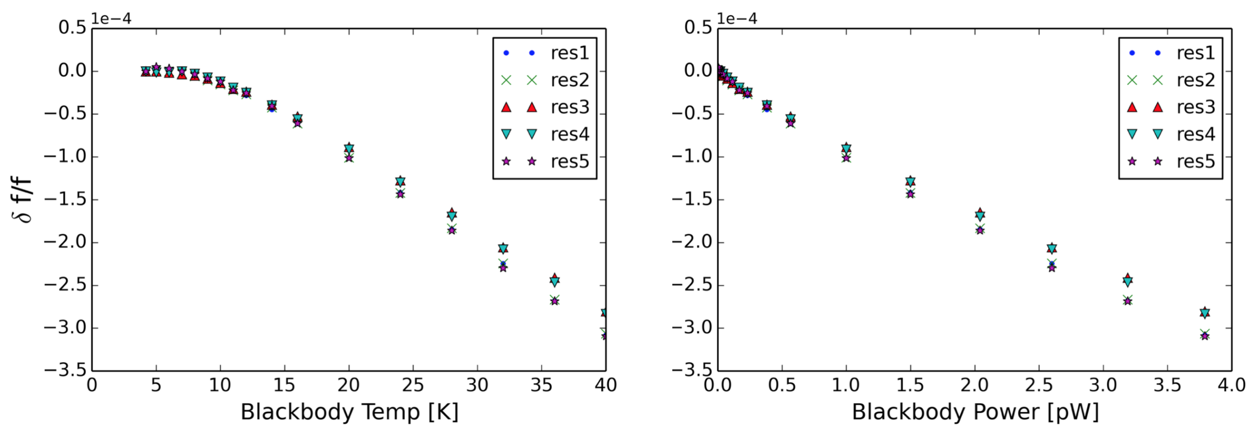

Figure 3.4: Five resonators' frequency shifts (relative to no optical load) as a function of blackbody temperature (*Left*), and the same plot with the x-axis converted to optical load power in picowatts (*Right*). These data are fit to find R_x, the responsivity in the frequency direction.

To make this measurement the blackbody temperature T_{BB} must be converted to the power incident on the detectors P_{BB}. A few parameters are required to do this conversion: First, the number of modes n_{modes} of blackbody radiation that can make it through the lenslet

$$n_{modes} = A_{lenslet} \frac{\Omega}{\lambda^2}, \tag{3.2}$$

with $A_{lenslet}$ the physical size of the lenslet, Ω the solid angle subtended by the blackbody, and λ the wavelength of light for which we are computing this quantity. ε_{opt}, calculated in Section 3.5 quantifies how much light can make it through the optical components. Finally, the average energy per mode from a blackbody with temperature T is

$$\langle E \rangle = kT \frac{h\nu/kT}{exp(h\nu/kT)-1}. \tag{3.3}$$

In total, the blackbody power is the product of these three quantities – the number of modes times the energy in each of those modes that reaches the detectors per second. Stated mathematically,

$$P_{bb}(T)[Watts] = n_{modes} \times \varepsilon_{opt} \times \frac{(k_B T)^2}{h} \int_{\nu_1=27\times 3e10}^{\nu_2=30\times 3e10} \left(\frac{h\nu/k_B T}{e^{\frac{h\nu}{k_B T}} - 1} \right) d\nu, \tag{3.4}$$

where $\{\nu_1, \nu_2\}$ is the bandpass defined by the filters.

Responsivity data (the right panel of Figure 3.4) were fit to Equation 2.12, using $\gamma = 1$ for the thin film limit, $\eta_o = 0.13$, $\eta_{qp} = 0.4$, $N_0 = 1.7 \times 10^7$ eV$^{-1}\mu$m^{-3}, and $V = 94$ μm^3. η_o is set to what was needed to fit the responsivity data, while η_{qp} and N_0 have been measured for Al. V is the calculated volumed of superconductor based on geometry. The responsivities for the four optical loads are enumerated in Table 3.2.

The responsivities R_x found above are in the frequency direction - they quantify how much the resonators' f_0's shift as a function of blackbody power. To quantify the change in quality factor with blackbody power, responsivities in the dissipation direction R_{diss} are needed. Since R_{diss} has to do with the dissipation (1/Q) and R_{freq} has to do with shifts in f_0, R_{diss} can be found by scaling R_x by the relationship between (1/Q) and x:

$$R_{diss} = 2R_x \left| \frac{d(1/Q_i)}{dx} \right|. \tag{3.5}$$

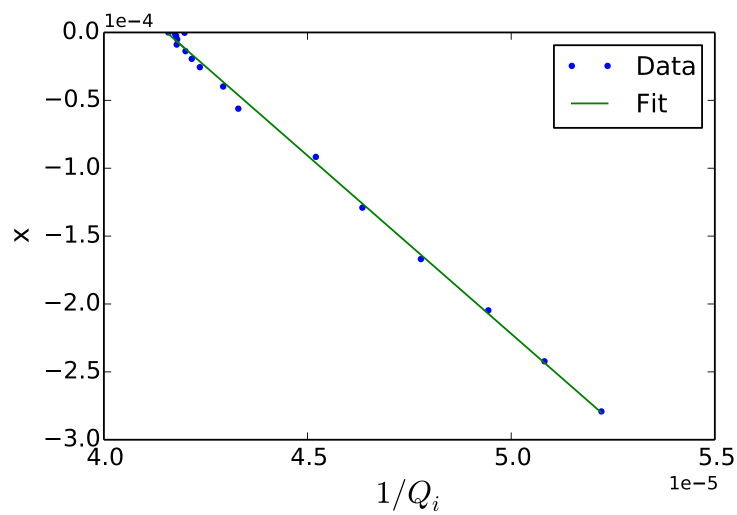

Figure 3.5: frequency shift vs shift in internal quality factor. data (*blue circles*) and linear fit (*green line*). used to convert from responsivity in the frequency direction to responsivity in amplitude direction

The factor of 2 comes from differentiating Equation 4.5 and comparing the voltage response to changes in x and (1/Q):

$$\begin{aligned} \frac{dV}{dx} &= V_0 \frac{dS_{21}}{dx} = V_0 2iQ_t^2/Q_c \\ \frac{dV}{d(1/Q_i)} &= V_0 \frac{dS_{21}}{d(1/Q_i)} = V_0 Q_t^2/Q_c, \end{aligned} \tag{3.6}$$

leading to the factor of 2 (and a phase shift) between $dx/d(1/Q)$ and the response ratio in the two directions. Figure 3.5 shows the data used in this conversion, which come from fitting S_{21} at each blackbody power yielding $Q_i(T_{bb})$ and $x(T_{BB})$ (as shown in Figure 3.2). A line is fit to these data, resulting in a measured slope $\frac{d(1/Q_i)}{dx} = \frac{1}{26}$, meaning $R_{diss} = R_x \times \frac{2}{26}$.

3.6.2 Noise Measurement Results

Measuring Quasiparticle Recombination Time: The PSD of the noise measurements have a rolloff at the high frequency end, shown in Figure 3.6. The detector bandwidth ($f_0/2Q$) for these detectors is $\sim 1.5 \times 10^4$ Hz, far above the observed rolloff around 10^3 Hz. Therefore, we must assume the feature is caused by the detector's inability to react faster than its slowest timescale. We calculate τ_{res} according to Equation 2.18 to be between 10 and 20 μs based on the Q's measured for our detectors. A second possible relaxation timescale that could be causing the rolloff is the quasiparticle recombination time τ_{qp}.

A fit of Equation 2.19 to the data in Figure 3.6 results in a relaxation timescale $\tau \approx 100$ μs. Since τ measures the longest timescale but $\tau_{res} < \tau$, we conclude that τ_{qp} is the longest timescale in the system and is 100 μs.

We convert our power spectral densities to NEPs according to Equation 2.20 in Chapter 2. Fig. 3.7 shows the NEP$_{freq}$ under the four optical loads. 1/f noise dominates below 10 Hz, which is especially apparent under the lowest optical load. The 1/f noise is not common mode and therefore can not be rejected by a correlation analysis. This same figure shows that these devices are photon-noise limited at optical loads above 0.6 pW, but that there is excess noise evident for the lowest load.

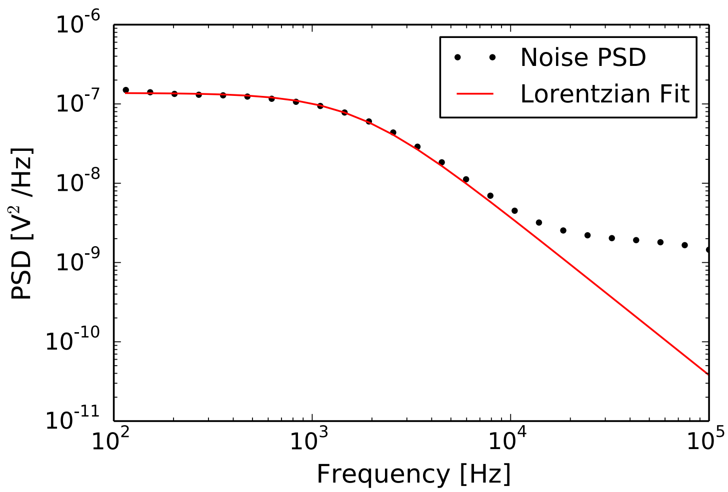

Figure 3.6: The noise power spectral density (PSD) as a function of frequency, and a fit to the rolloff that results in a measurement of τ_{qp}, as described in Chapter 2, Section 2.2.7. The fit deviates from the data at the highest frequencies because the noise is becoming dominated by amplifier noise, which is a different noise source than what is being fit for at the lower frequencies.

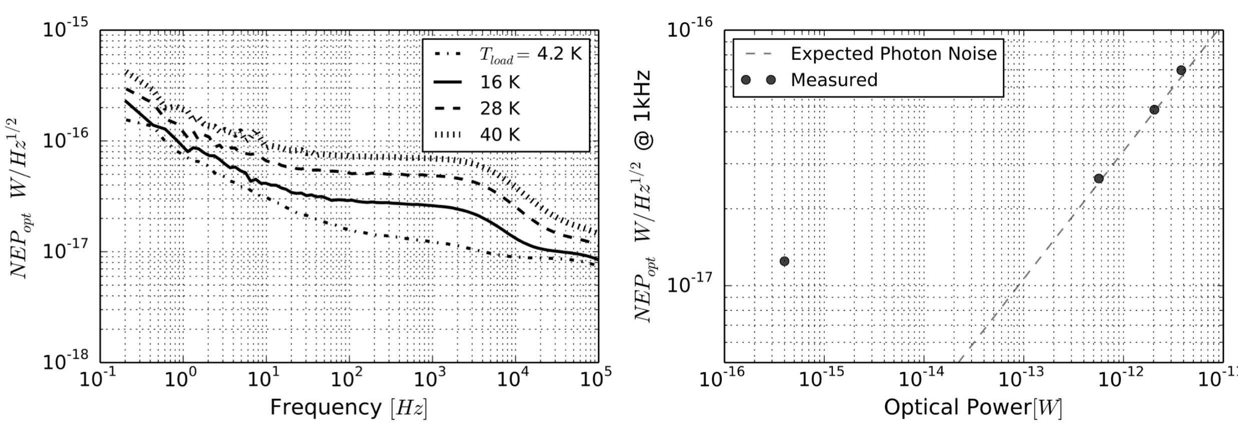

Figure 3.7: *Left*: Measured optical NEP in the frequency direction for the four optical loads described in Section 3.5 for our second highest readout power. For reference, the highest readout power was within a factor of 2 of bifurcation. *Right*: A comparison of the measured noise at 1 kHz to the theoretical photon noise. Our detectors are photon noise limited above optical powers of 0.6 pW.

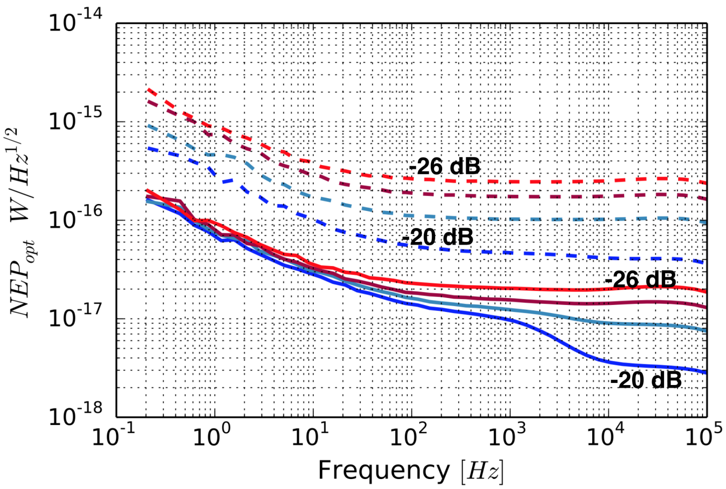

Figure 3.8: The amplitude (*dashed*) and phase (*solid*) noise spectra for the lowest blackbody temperature of 4 K and all readout power dissipations, in steps of 2 dB. The phase noise is much lower than the amplitude noise, but the amplitude noise decreases with increasing readout power, while the phase noise stays nearly constant.

Figure 3.8 explores the effects of readout power on both the amplitude and phase noise for the lowest optical load of 0.4 fW. For a given optical load the NEP in the dissipation direction (NEP_{diss}) is between five and ten times higher than the NEP in the frequency direction (NEP_{freq}), depending on the readout power. NEP_{diss} decreases with increasing readout power, while NEP_{freq} stays nearly constant with readout power.

3.6.3 Uncertainties and Implications

The calculated optical responsivity depends on an accurate conversion from blackbody temperature to absorbed power. A major uncertainty in this conversion is the optical coupling coefficient ε_{opt}, to which the blackbody power is directly proportional. Other uncertainties affect only the expected responsivity. For example, the assumed active area of the absorber determines the active volume of the detector, which affects the responsivity according to Eq.(2.12).

3.7 Discussion of Results

The expectation for this bilayer was that the Al would dominate the superconducting properties, but the τ_{qp} is more consistent with that of TiN; Al lifetimes of several ms have been measured under small optical loads Baselmans and Yates (2009), while the measured τ_{qp} was much less than this.

Another unexpected result was the measured bilayer T_c, which was higher than that of single Al (1.2 K) or TiN (1 K) films deposited under identical conditions. Additionally, the responsivity was smaller than expected by about 20 times, likely due to the short quasiparticle lifetimes. Absorption of stray light by the detectors could cause the shortened τ_{qp}, but it is also possible that this is simply due to the higher quasiparticle density that comes with the elevated T_c. The low quality factors could also indicate absorption of stray light in the system.

The behavior of NEP_{diss} and NEP_{freq} with readout power gives a good indication of what is the main source of noise for these detectors in each direction. There are two main sources of noise that are dependent on readout power: amplifier noise and TLS noise.

Amplifiers cause noise mostly in the dissipation direction. For increasing readout power, the signal strength into the amplifier is also increased due to the heightened response of the KIDs to a stronger driving force. This increase would in turn improve the signal to noise ratio and drive down the NEP. It is likely that noise in the dissipation direction is determined by noise from the amplifier, since NEP_{diss} decreases with readout power but cannot be determined by TLS noise which is almost entirely in the frequency direction.

If TLS noise were an important contributor to this detector, then a decrease in NEP_{freq} for increasing readout power would be expected (see Section 2.3.3). However, this is not what is observed, with the detectors' NEP_{freq} staying about constant for all readout powers. For our three highest optical load powers this behavior is expected because the detectors are photon-noise limited and NEP_{photon} does not depend on readout power (see Section 2.3.2). For the lowest optical load this behavior is not well understood yet. Its noise lies well above the photon noise limit (see Figure 3.7), but its NEP_{freq} is still independent of readout power. It is possible that there is stray light that has not been accounted for in the system, which would make the calculated optical power for this lowest blackbody temperature underestimated. If this is the case, then the photon noise limit for this power would be underestimated as well.

3.8 Summary

The main success of this work is demonstrating that extremely low-volume (94 μm^3) Al microstrip lines can be made robust by depositing TiN atop the lines, resulting in a 95% yield for an array of 96 detectors. Our KIDs reached the photon noise limit for optical loads > 0.6 pW for modulation frequencies above 10 Hz, but were far above this limit for the lowest optical load. The lowest NEP reached was 1.2×10^{-17} W/$\sqrt{Hz}$.

Chapter 4

Phonon Recycling KIDs

4.1 Introduction

In this chapter we introduce the concept of phonon recycling, an effort to boost responsivity and reduce NEP. In short, the mechanism of "phonon recycling" refers to trapping the phonons created by photo-produced quasiparticles recombining into Cooper pairs in the detector's absorber/inductor area. These recombination phonons can be reabsorbed by the absorber/inductor, breaking more Cooper pairs and thereby enhancing the kinetic inductance response. Trapping the recombination phonons elongates the quasiparticle lifetime τ_{qp}, to which the responsivity is directly proportional (Equation 2.12). We first define the concept of phonon recycling and describe the geometric changes that can be made to a KID to achieve phonon recycling. We then present our Monte-Carlo simulations that predict the amount of signal boost expected from phonon recycling given flexible detector geometries and illumination schemes. Finally, we describe attempts to measure the phonon recycling effect on arrays fabricated in the Microdevices Laboratory at the Jet Propulsion Laboratory.

4.2 Phonon Recycling Theory

When Cooper pairs broken by incident radiation recombine, they release a recombination phonon of energy $\geq 2\ \Delta_0$. These phonons are energetic enough to break additional Cooper pairs, though in traditional KID geometries the majority are lost to the phonon bath and do not contribute to the detected signal. Phonon recycling is the process of making geometrical changes to the

KID array in order to trap these recombination phonons in the detector active area, facilitating re-absorption and further Cooper pair breaking.

The trapped recombination phonon has a chance to break an additional Cooper pair, which itself will eventually re-combine after an average time τ_{qp} and create a new recombination phonon, which can again be re-absorbed and break another Cooper pair. The number of times the phonon is emitted and re-absorbed is P_{rf}-1 , with P_{rf} called the "phonon recycling factor". This emission and re-absorption increases the effective quasiparticle recombination time from $\tau_{qp,0}$ to

$$\tau_{qp,rec} = P_{rf} \times \tau_{qp,0}, \tag{4.1}$$

thereby increasing the responsivity $R_x \propto \tau_{qp}$. This re-absorption and emission process can in theory happen many times until the phonon anharmonically down converts or escapes the active inductor area.

Phonon trapping can be achieved by a few physical changes to the KID array: one, thin the substrate below the detector active area (the inductor/absorber). One could also etch away a portion of the perimeter around the active area. In extreme geometries, this would leave the inductor/absorber connected to the surrounding substrate by only thin legs, mimicking the geometry of a bolometer. Thinning the substrate increases the number of reflections a phonon will undergo before escaping the active area, thereby improving chances for re-absorption in the inductor. Similarly, suspending the inductor on legs decreases the probability for a phonon to escape from the active area before being re-absorbed. Given a number of times a phonon is absorbed in the inductor N_{abs}, the phonon recycling factor is determined by

$$P_{rf} = 1 + N_{abs}. \tag{4.2}$$

Back-thinned detectors are fabricated on a silicon-on-insulator (SOI) wafer, a "sandwich" of silicon bonded together by a thin oxide layer. The lower layer of silicon is etched away in a circular pattern on the backside of the KID absorbers, with the etching halted at the oxide layer. This results in inductors atop a thinner layer of silicon than the rest of the array. The upper layer of silicon of a SOI wafer can be anywhere from 100 nm to 300 μm thick, with the lower layer of silicon on the order of 300 μm thick.

The presence of an oxide in the substrate should not introduce additional noise. Initially this was a concern, because TLS noise is exacerbated by oxide layers (Gao et al., 2007). However, this is only true of oxides on the surface where the TLS states exist. TLS states predominantly occur on oxide layers present on the KID deposition layer, not on the backside of the wafer where the SOI oxide would be present.

An added benefit of phonon recycling arrays is the isolation of detectors when cosmic rays are absorbed by detectors on the array. Cosmic arrays have been a source of major data loss in previous space-based missions, for example affecting 12-20% of data on the Planck mission (Catalano et al., 2014). Bueno et al. (2018) demonstrated the efficacy of back-thinned devices in mitigating the effect of cosmic ray events in a 989 pixel array. For the array with a solid substrate a single cosmic ray event affects the entire array, causing full array downtime. The back-thinned array confines cosmic ray events to much fewer KIDs at a single time, with a calculated downtime of only ~0.6% for an L2 orbit.

4.3 Simulations of Phonon Recycled Devices

4.3.1 Simulation Geometry

To investigate the influence of these geometric changes on the phonon recycling factor, we simulate the phonon reflection and re-absorption process in a KID using a Monte-Carlo style simulation. The KID geometry is the same as that in Figure 3.1 (Chapter 3) with a circular inductor connected to a rectangular capacitor. We simulate detectors made fully of TiN, though the simulation can easily be altered for other superconductors as needed. The phonons scatter in three dimensions within the silicon below a circular inductor with radius R_L made of TiN of thickness t_{TiN} deposited on a circle of back-thinned Si with radius $R_{thin} > R_L$ and thickness t_{thin}. The thinned silicon cylinder is suspended along its perimeter via four equally spaced legs with variable size de-

Parameter	Description	Default value
t_{TiN}	thickness of inductors TiN	20 nm
R_L	radius of inductor	90 μm
t_{thin}	thickness of Si below inductor	2 μm
R_{thin}	radius of thinned Si, $R_{thin} > R_L$	100 μm
η_{fill}	filling fraction of inductor	$\frac{1}{3}$
l_{pb}	phonon mean free path to pairbreaking	1 μm
l_{anharm}	phonon mean free path to anharmonic decay	1 cm
p	perimeter filling fraction of Si island	0.7

Table 4.1: User-defined variables for the phonon recycling simulations and their default values.

termined by the perimeter filling fraction $p = \{0, 1]$. A p of 1 indicates a fully connected perimeter while small p's indicate thin legs suspending the back-thinned silicon island. All user-defined parameters, a brief description, and default values are listed in Table 4.1.

The final simulation geometry does not include the KID capacitors because Cooper pairs broken in the capacitor do not contribute to detectable signal. Initial simulations were conducted including the rectangular capacitor deposited on non-thinned silicon to investigate whether phonons created in the capacitor area could travel into the thinned silicon below the inductor, thereby contributing to the phonon recycling signal. We found $< 0.6\%$ of capacitor phonons successfully make this journey. Intuitively this makes sense given the extremely small solid angle seen by phonons traveling from the thick (350 μm) to thin (2 μm) silicon. Given the negligible contributions of capacitor phonons to the phonon recycling signal, we eliminated the capacitor from the simulations to improve simulation speed.

4.3.2 Phonon Behavior and Simulation of P_{rf}

Phonons are initiated in random positions within the inductor's circular profile at the top of the detector cell, as if they had been created by Cooper pairs recombining within the TiN of the inductor. Phonon mean free paths in Si have been measured to be 80 μm at 77 K and 12 μm

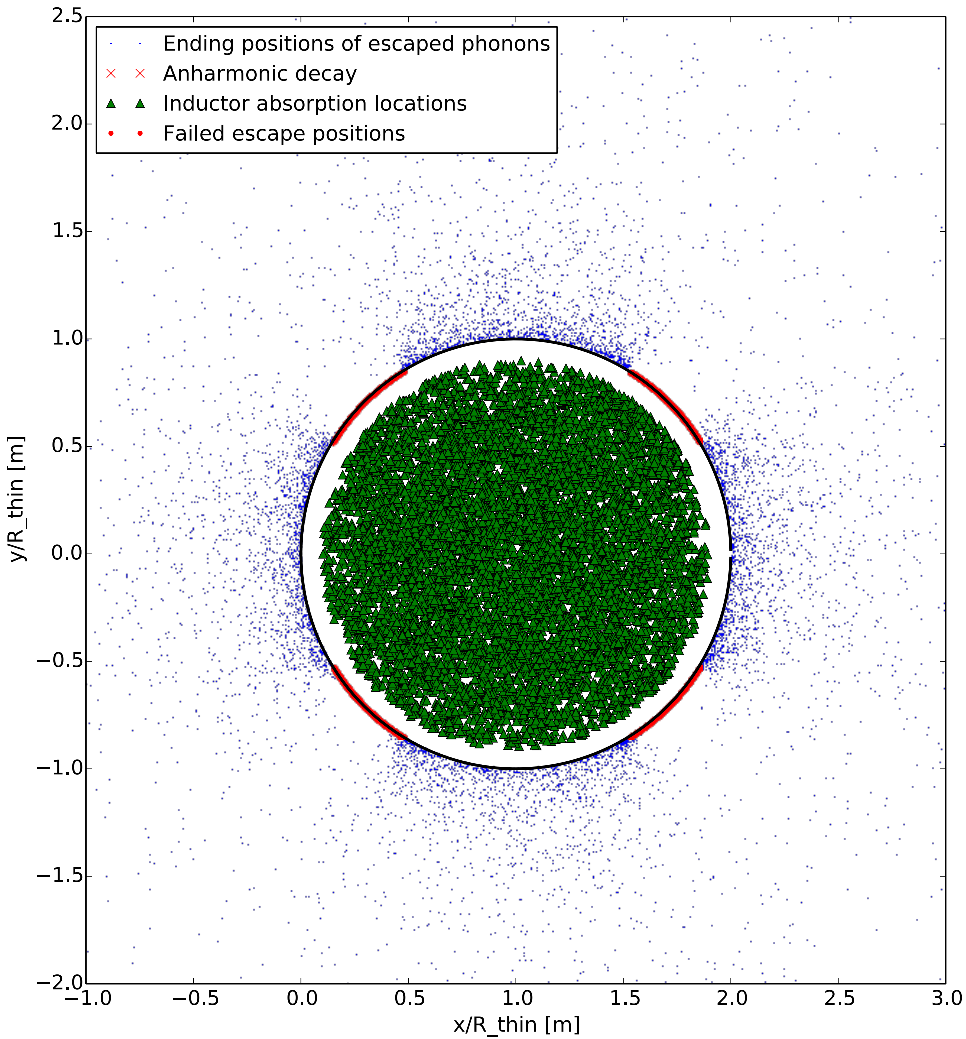

Figure 4.1: The recorded ending positions (*blue dots*), locations of re-absorptions in the inductor (*green triangles*), re-initiation points after failed escape attempts (*red dots*), and locations of anharmonic decay of phonons (*red exes*, none visible in this plot) given the default parameters in Table 4.1 . The effect of isolation of phonons via legs is shown by the ending positions and failed exit positions of the phonons: where legs exist, the phonons are able to escape and spread out of the thinned silicon island. Where legs are not present, the phonons are stopped from escaping and re-initiated at the locations of the red dots.

at 250 K Gereth and Hubner (1964). Simulations extending to the low temperatures relevant to our experiments show mean free paths exponentially increasing past 0.01 cm at 25 K (Kukita and Kamakura, 2013), suggesting mean free paths in excess of 1 cm at temperatures of 0.1 K. Measurements by Marx and W. Eisenmenger in 1982 Marx and Eisenmenger (1982) found that 96% of phonons reflect isotropically ("diffusively") from surfaces and found negligible anharmonic decay at surfaces upon scattering.

This physics is captured in the phonon behavior programmed into the simulation. The long mean free paths mean the phonons travel in a straight trajectory through the thin Si. Upon incidence with a Si surface, the phonons scatter isotropically without energy loss. Phonons travel and reflect in this manner until they leave the thinned silicon area through the legs (or travel a total distance of 1 , assumed to have anharmonically decayed, and are removed from the simulation).

Any time the phonon hits the top of the detector cell, it has a probability of re-absorption η_{rec} by the TiN on the order of 0.01, given by

$$\eta_{rec} = T_{acoustic}\eta_{fill}(2t_{TiN}/cos(\theta))/l_{pb}, \tag{4.3}$$

where $T_{acoustic}$ is the acoustic matching factor of order unity (Kaplan et al., 1976), η_{fill} is the filling fraction of the inductors TiN, θ is the entry angle of the phonon into the Si, t_{TiN} is the thickness of the TiN, and l_{pb} is the mean free path to pair breaking in the superconductor. l_{pb}=$v_{phonon}\tau_{pb}$ is not well measured for TiN, however, Kaplan et al. (1976) report that the phonon lifetime against pairbreaking τ_{pb} is $\propto T_c$. Therefore, Al and TiN's comparable T_c's should have similar pair breaking lengths (to first order), and we use Al's value of $l_{pb} \approx 1$ μm as an approximation for l_{pb} in these simulations. Note that η_{rec} is different on each phonon encounter with the TiN due to the randomly generated θ associated with each reflection. Each absorption is counted, and the average number of absorptions per phonon is calculated at the end of the run, yielding the recycling factor according to Equation 4.2.

After reabsorption in the TiN, the phonon is then re-emitted in a random direction and can continue to travel between the top and bottom of the silicon at random angles with opportunities for re-absorption and emission until one of two things occur:

(1) The phonon successfully leaves the thinned silicon area through the legs on the perimeter. If a phonon attempts to leave the thinned Si in a location outside of the legs, it is re-initiated at the perimeter and continues its path between the top and bottom surfaces, with possibility for re-absorption upon every encounter with the inductor.

(2) The phonon travels beyond the anharmonic decay length of phonons in Si, assumed to be $\sim$ 1 cm (Simulations extending to the low temperatures relevant to our experiments show mean free paths exponentially increasing past 0.01 cm at 25 K (Kukita and Kamakura, 2013), suggesting mean free paths in excess of 1 cm at temperatures of 0.1 K). If this occurs, the phonon records its position and then ceases its travels. This occurrence is rare but happens more frequently as the perimeter filling fraction approaches 0. On average a simulated phonon will reflect 700 times for the smallest simulated perimeter coverage fraction of 0.01 and a Si thickness of 2 μm, a total approximate travel length of $\sim$1.4 mm. If the phonon decay length is $<<$ 1 cm as is assumed in the simulations, the phonon would decay sooner and eliminate it from the phonon recycling signal. This would decrease the phonon recycling signal approximately proportionally given the simulation geometry, since anharmonic decay lengths less than 1.4 mm would result in a proportionally smaller number of phonons contributing to the recycling signal. This caveat is only relevant for the smallest simulated perimeter coverage fractions – For a perimeter coverage fraction of 1 the phonon only bounces an average of 30 times, and a p of 0.1 only around 120 times.

The phonon ending positions, locations of re-absorptions, re-initiation points after failed escape attempts, and locations of anharmonic decay are shown in Figure 4.1 for the default simulation values in Table 4.1.

4.3.3 Simulation Results and Predictions

The simulation explored the influence of three geometrical parameters on the phonon recycling factor P_{rf}: the perimeter filling fraction p (Figure 4.2), the ratio of the thinned silicon to the radius of the inductor t_{thin}/R_L (Figure 4.3), and the ratio of the inductor radius to the radius of the circle of thinned silicon it is deposited on R_L/R_{thin} (Figure 4.4).

The simulations were run over parameter spaces that are reasonably achievable for our fabrication processes. Small perimeter filling fractions and extremely thin silicon are both hard to fabricate and delicate. t_{thin}/R_L could be decreased by increasing the size of the inductor rather than thinning the silicon, but large inductor radii contribute to worse performance due to the larger active detector volume. t_{thin} is limited by SOI wafer geometry available on the market, with the thinnest upper layers being a mere 100 nm thick corresponding to t_{thin}/R_L of 1E-3 for an inductor radius of 90 μm.

The steepest dependency of P_{rf} on a geometric parameter results from small perimeter coverage fractions (Figure 4.2). The simulation shows an exponential increase of P_{rf} with decreasing p, with $P_{rf} = 1.23\ p^{-0.61}$. For the smallest simulated p of 10^{-2}, simulations predict a recycling factor of 20.

The reliance of P_{rf} upon the ratio of t_{thin}/R_L is steeper than an exponential, but does not achieve as high of recycling factors as reducing the perimeter coverage fraction.

The effect of increasing the radius of the thinned silicon patches below the inductor is quite modest (Figure 4.4). Even with thinned silicon patches with twice the radius of the inductor, a P_{rf} of only 2 is simulated. R_{thin} cannot be further increased with respect to R_L without thinned patches merging between adjacent inductors, increasing the probability of undesirable cross-talk between detectors.

Using these simulations, we can select a geometry that should maximize phonon recycling effects: a small perimeter filling fraction, either an extremely thin silicon backlayer or a large inductor, and as large of a thinned silicon area as compared to the radius of the inductor as possible.

As discussed previously, taking these parameters to the extreme can cause fabrication and noise issues. Small perimeter filling fractions and thin silicon are more difficult to fabricate and are more delicate, while large inductor radii increase noise.

4.4 Back-thinned TiN Array

Rick LeDuc at the JPL microdevices lab fabricated a 96 pixel array of 20 nm thick TiN deposited on a SOI wafer. Each pixel consists of a meandered inductor/radiation absorber with a roughly circular outline and an interdigitated capacitor to mitigate two-level-system noise (Gao et al., 2007). A picture of a single detector is shown in Figure 4.5. The readout frequencies were designed to be a few hundred MHz with quality factors on the order of 3×10^4. To trap recombination phonons below the detector active area (the inductor) and achieve phonon recycling, the SOI is etched away in a circular pattern below the inductor so that the inductors are suspended on a 2 μm thick Si membrane. Half of the detectors are back-thinned and half of the detectors remain atop the full thickness of the SOI wafer. Rows of thinned/non-thinned detectors alternate on the array. For a single detector with this geometry, a recycling factor of R = 1.49 is simulated according to Section 4.3.

4.5 Measuring P_{rf}

All tests were completed before the testbed at CU was fully functional in generous, unpaid collaboration with Christopher McKenney and Jiansong Gao at the National Institute of Standards and Technology (NIST). At NIST, the array was flashed with a laser diode (1.5 μm telecom laser) and the decay was fit with a single exponential to measure the quasiparticle recombination time. According to 4.1 and the simulations run above, we should measure a 50% increase in τ_{qp} for the resonators with back-thinned Si that comprise half of the array.

When the array is exposed to radiation, the resonant frequencies decrease and the quality factors degrade (Chapter 2). The quasiparticles generated by the radiation then recombine, resulting

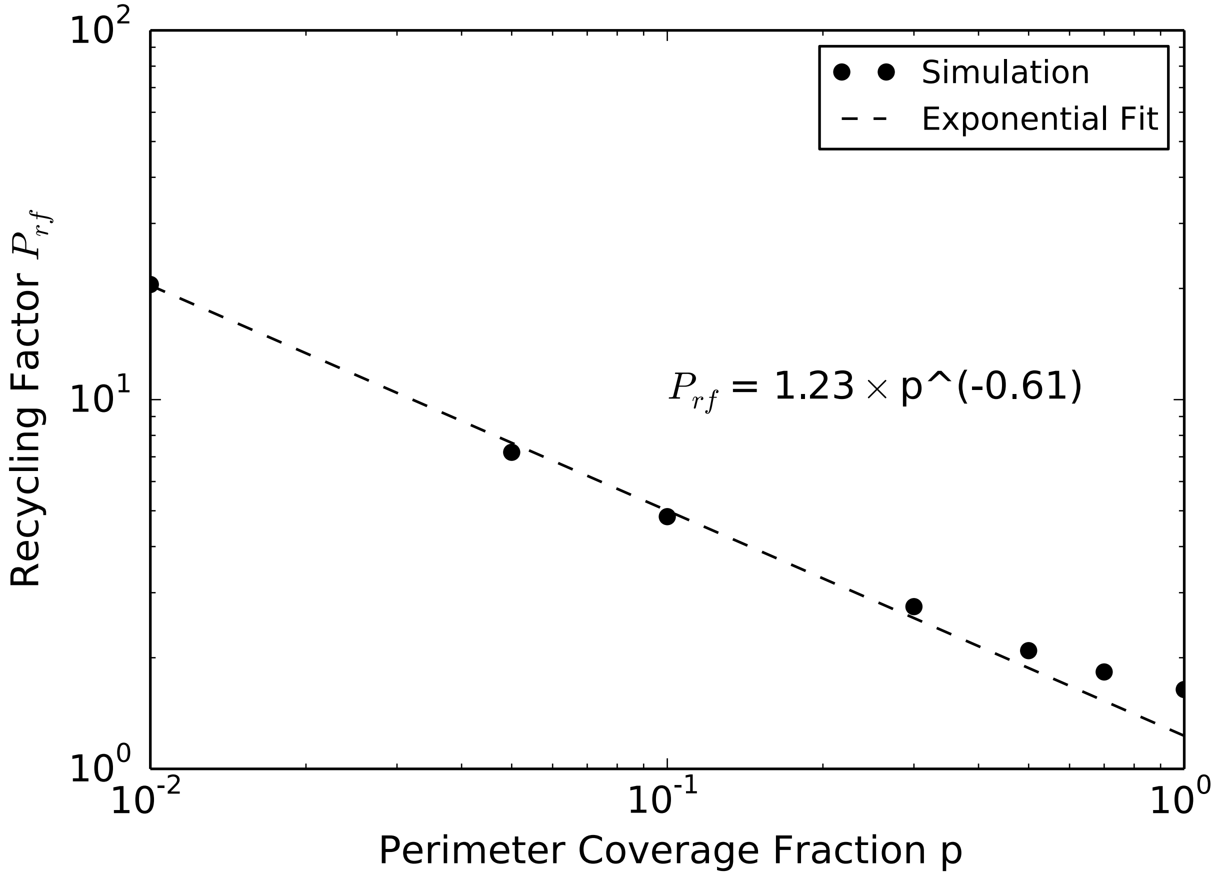

Figure 4.2: The simulated P_{rf} for a variety of filling fractions p, fit with a power function. As expected, lower filling fractions (smaller legs) lead to higher recycling factors due to the increased difficulty of phonon escape.

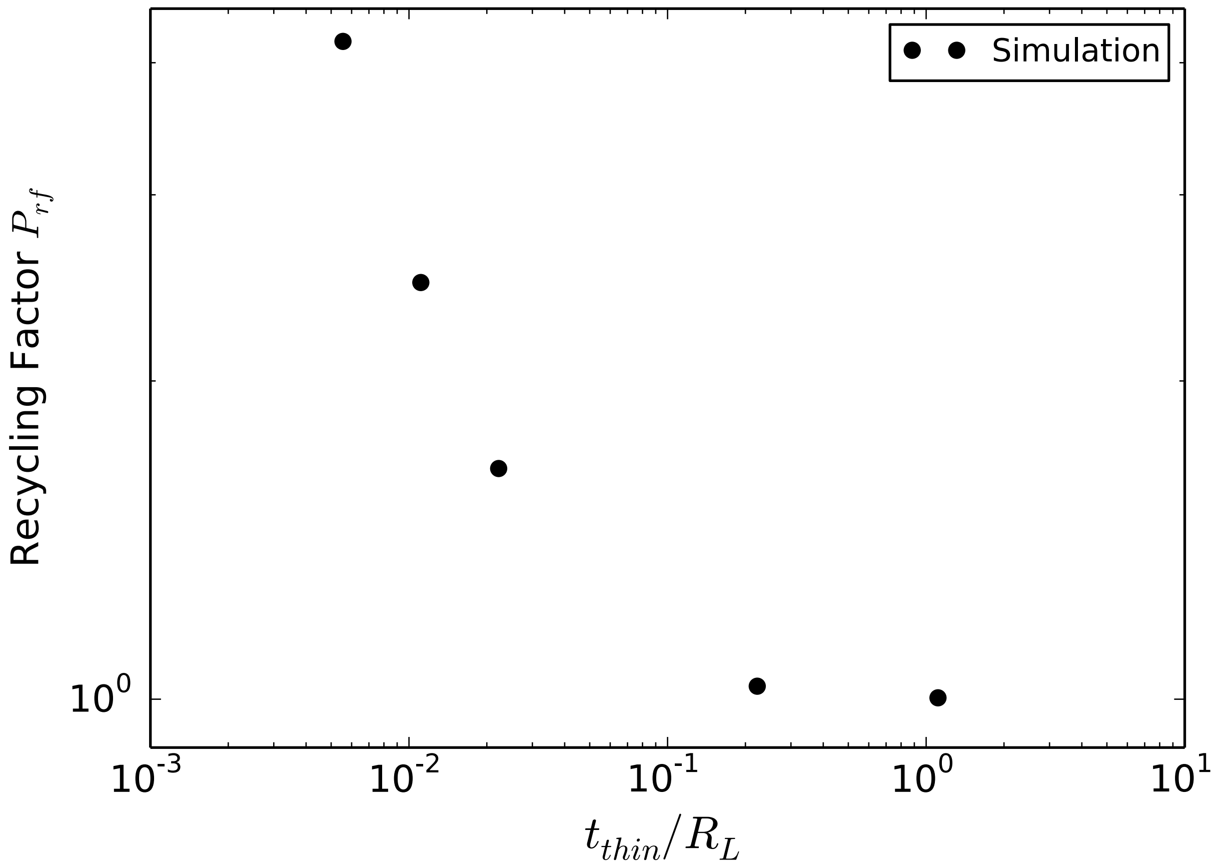

Figure 4.3: The simulated P_{rf} for a variety of t_{thin}/R_L. As the ratio decreases, the filling fraction increases sharply due to either thinner silicon resulting in more bounces and more opportunities for re-absorption, or larger inductors resulting in a greater areas and hence larger probabilities for re-absorption.

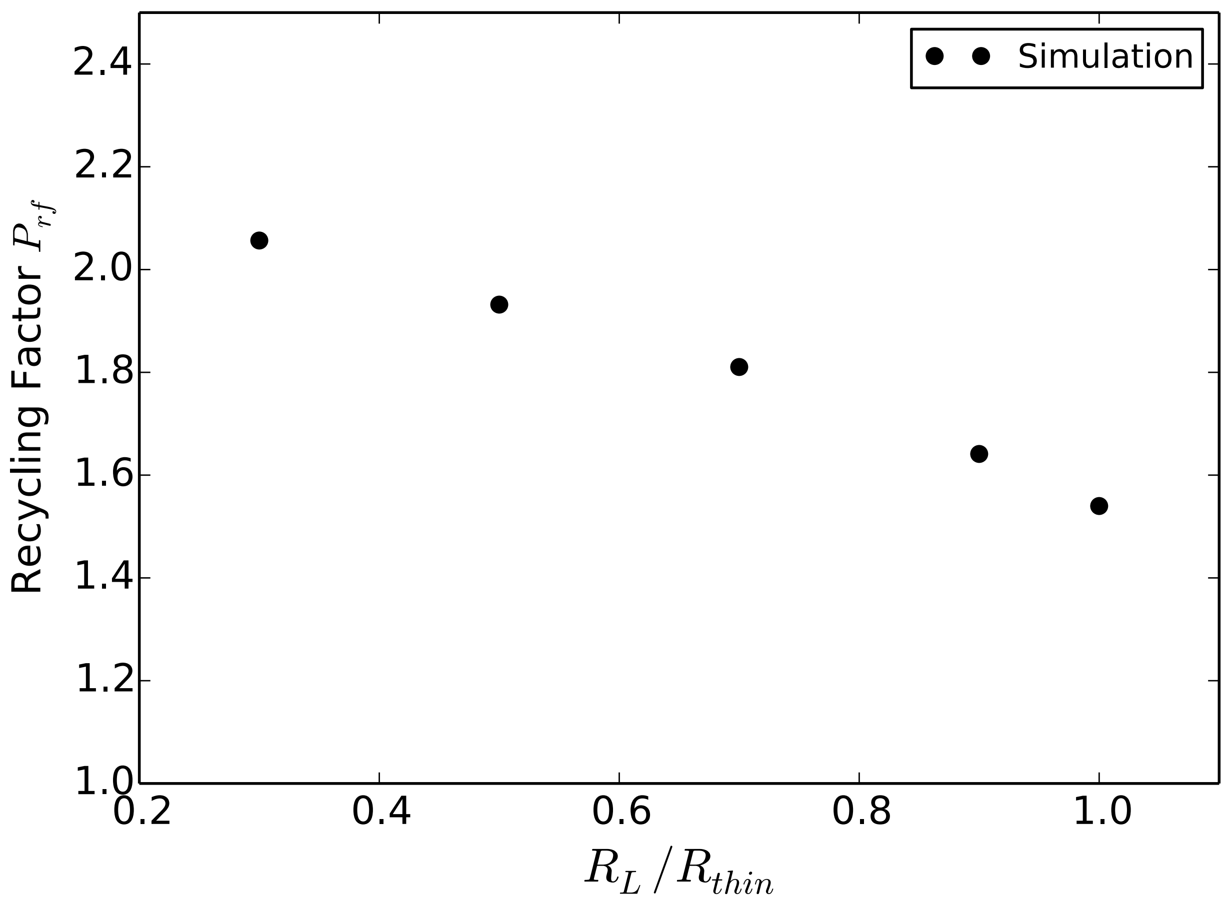

Figure 4.4: The simulated P_{rf} for a variety of R_L/R_{thin} – larger values mean a smaller annulus of bare, thinned Si around the inductor, and smaller values mean a larger annulus of bare, thinned Si.

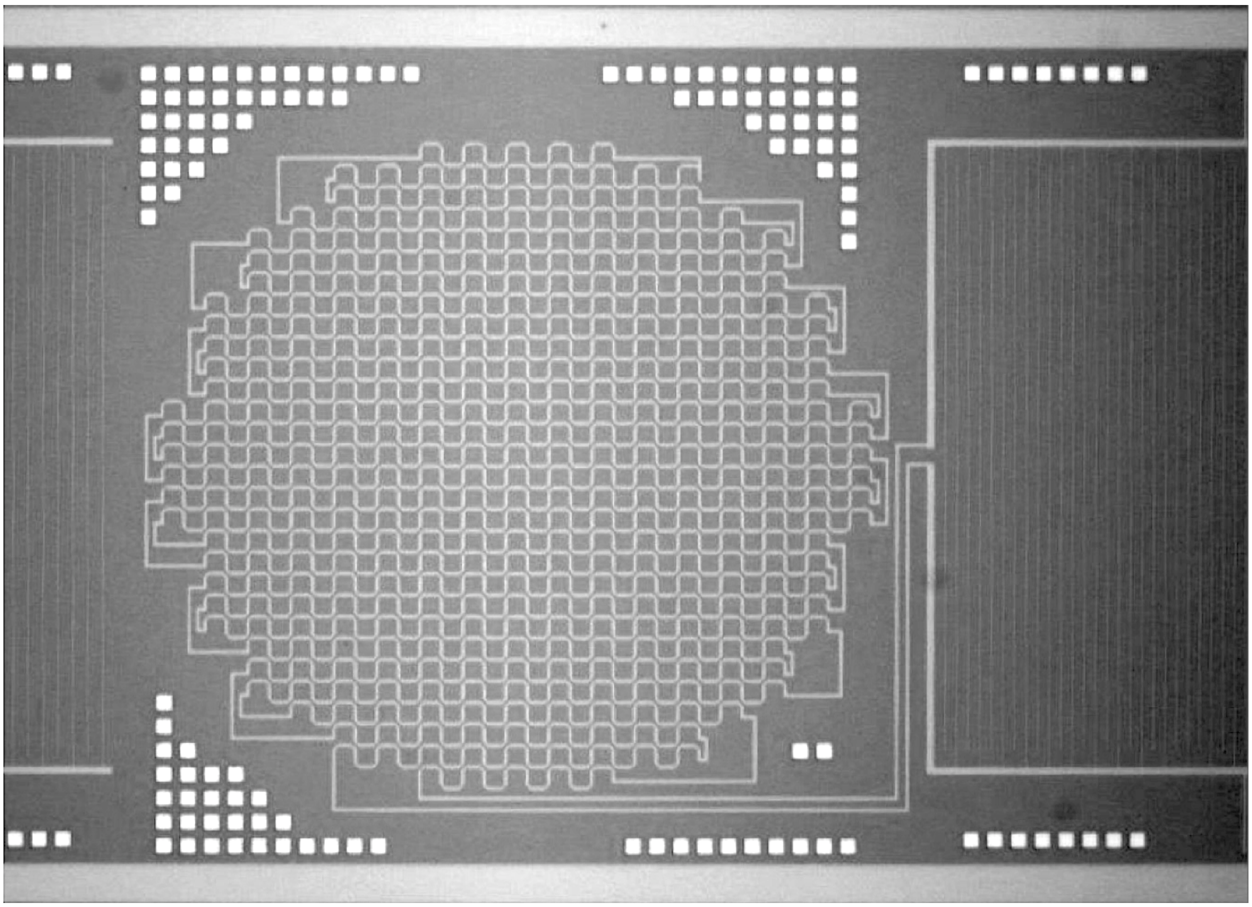

Figure 4.5: A picture of a single non-thinned detector with the same inductor and capacitor layout as the back-thinned devices. The roughly circular TiN meander is the inductor which also acts as the optical absorber, and is fully visible. For scale, the diameter of the inductor is 0.4 mm. The rectangular interdigitated capacitor, also made of TiN, is only partially visible on the right hand side of the image. In a phonon recycled device, the silicon in the circular area below the inductor would be thinned to allow trapping of phonons below the inductor, which is the detector's active area.

in an approximately exponential return of the resonant feature to its non-illuminated state. In later years, we formulated a more complex and physically motivated equation for fitting these decays (Chapter 5). However, for these early tests we used the accepted standard of the time of fitting exponentials to the decays.

The exponential decay of the fractional frequency shift away from resonance will be dominated by the longest time constant present in the system, of which there are several. The dominant time constants governing the exponential decay are the resonator ring time and the quasiparticle recombination time when other effects such as thermal influences are not present. The resonator ring time

$$\tau_{ring} = Q_t/(\pi f_0), \tag{4.4}$$

where Q_t is the total quality factor for a KID with a resonant frequency f_0, corresponds to how long it takes energy to dissipate from the resonator. Q_t is found by fitting

$$S_{21} = |A\big(1 - \frac{Q_t}{Q_c}\frac{e^{j\phi}}{1+2jQ_t dx}\big)| \tag{4.5}$$

as described in Chapter 2, Section 2.2.4.

We calculate τ_{ring} for each resonator and compare the exponential decay time τ_{exp} to this quantity. Values of $\tau_{exp}/\tau_{ring} > 1$ indicate we are not limited by the resonator ring time. If we find the exponential decay time is equal to τ_{ring}, then we are *not* measuring the quasiparticle recombination time which must be shorter than τ_{ring} and hence unmeasurable.

4.5.1 Non-uniform Illumination: The Importance of Lenslets and Diffusers

For our first measurement of τ_{exp}/τ_{ring}, the LED diode was located approximately 2 inches from the test array and was heat sunk to the 4 K stage of the cryostat. The LED diode was at the maximum distance possible from the array given the NIST cryostat's geometry.

Figure 4.6 shows the measured values of τ_{exp}/τ_{ring} for each resonator. A beam map completed by Chris McKenney at NIST prior to the author's involvement in the project provides the locations of each resonator. Due to the symmetry of the array, we cannot know whether rows one

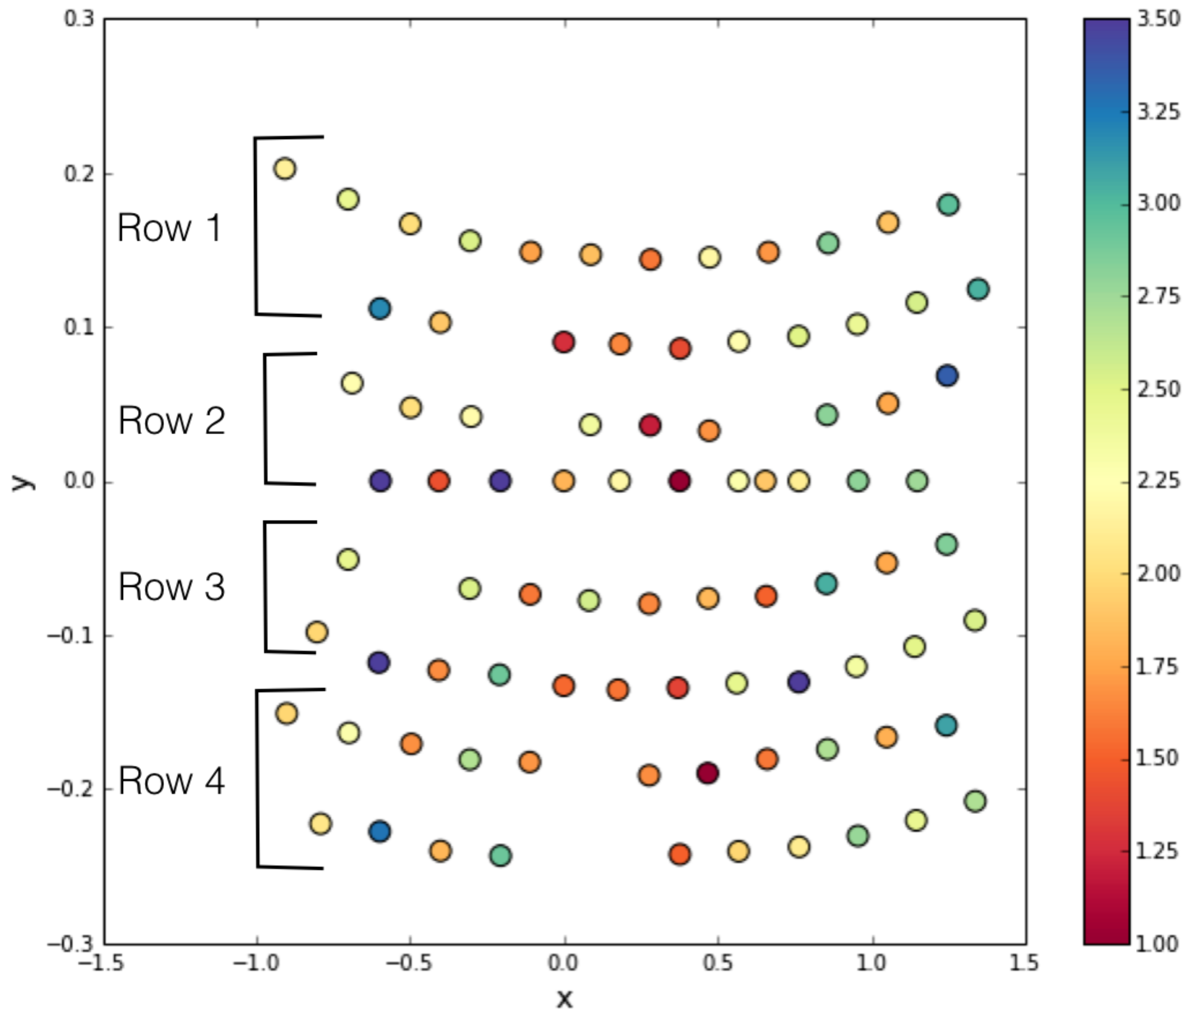

Figure 4.6: τ_{exp}/τ_{ring} for each resonator for the first test of the back-thinned array at NIST. Either rows 1 and 3 or rows 2 and 4 are back-thinned devices. All values are > 1 meaning we are measuring a time constant besides the ring time. However, we do not see a pattern of either rows 1 and 3 or rows 2 and 4 with elevated values.

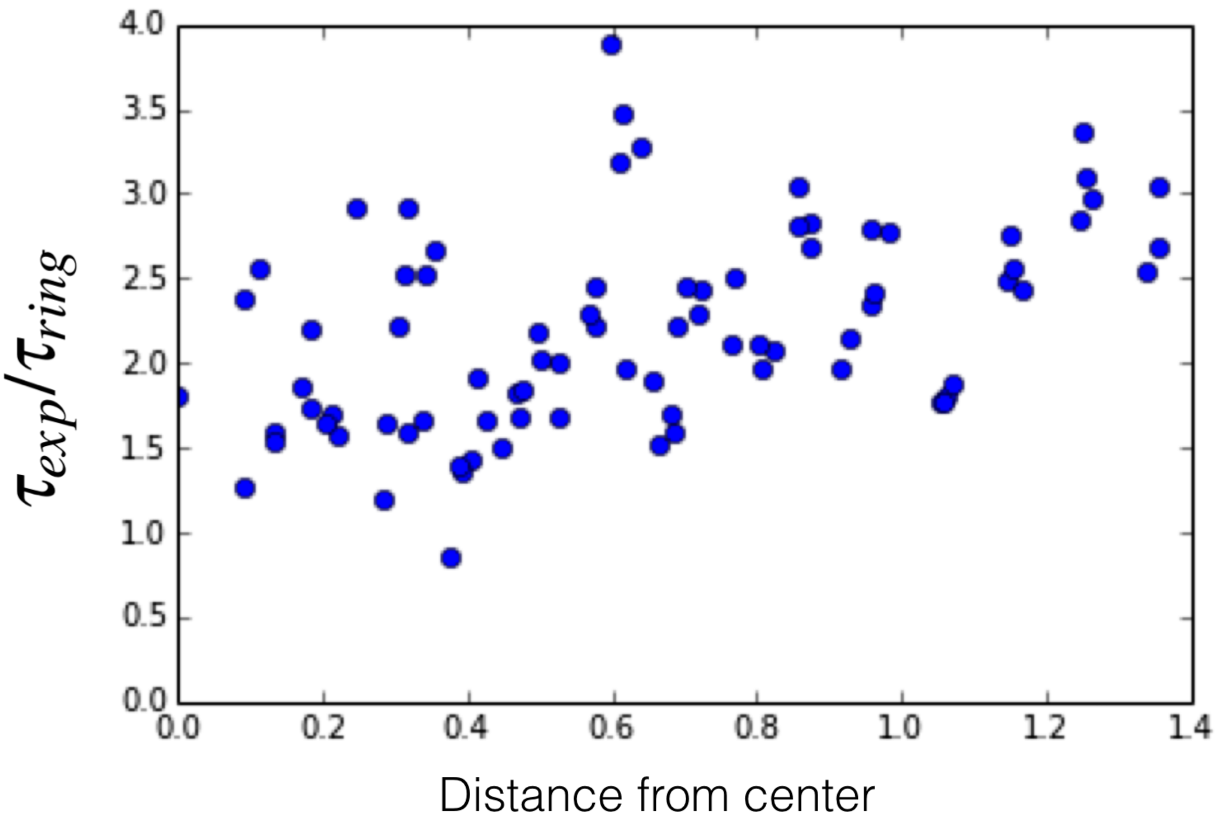

Figure 4.7: τ_{exp}/τ_{ring} vs unitless, normalized distance from the center of the array for every resonator on the array.

and three or rows two and four are back-thinned devices. Values of τ_{exp}/τ_{ring} are above one for every resonator, indicating we are measuring a time constant other than the resonator ring time. However, no systematic difference between the two sets of rows is seen, with each row having a mean $\tau_{exp}/\tau_{ring} \sim 2$. Despite this, histograms of τ_{exp}/τ_{ring} showed a bimodality, as we would expect if half of the resonators were successfully showing phonon recycled elongated τ_{qp}.

The question remained of what was causing the bimodality if not the phonon recycling. Figure 4.7 shows the measured values of τ_{exp}/τ_{ring} for every resonator on the array as a function of unitless distance from the center of the array. The exponential decay time increases as a function of distance from the center of the array, indicating non-uniform illumination of the array by the LED is the cause of elevated τ_{exp}/τ_{ring} for this measurement.

An attempt to mitigate the non-uniform illumination was made by fabricating a front plate with holes the size of inductors (shown in Figure 4.8). The front plate was aligned manually to the array inductors using a microscope. However, the non-uniform illumination remained an issue and again no augmented τ_{qp} was observed for the back-thinned devices compared to the standard inductors.

4.6 New Fabrication: Released Inductors

After the non-detection of τ_{qp} boosting and the completion of simulations that showed the current geometry would only augment τ_{qp} by 50%, a new geometry was devised to dramatically boost the phonon recycling effect.

As opposed to a circular pattern of thinned Si below the inductor, the Si was etched away in the same pattern as the inductor's meander lines. Each line of the inductor is therefore suspended on a thin bridge of Si, creating a fully released, suspended spiderweb of an inductor on Si. The spiderweb-like inductors are shown in Figure 4.9. Unlike the back-thinned devices for which the array was half back-thinned and half traditional inductors, every detector on this array was fully released. The complicated fabrication process required for this design had a high risk of damaging the detector's transmission lines and decreasing overall yield. Therefore, to maximize

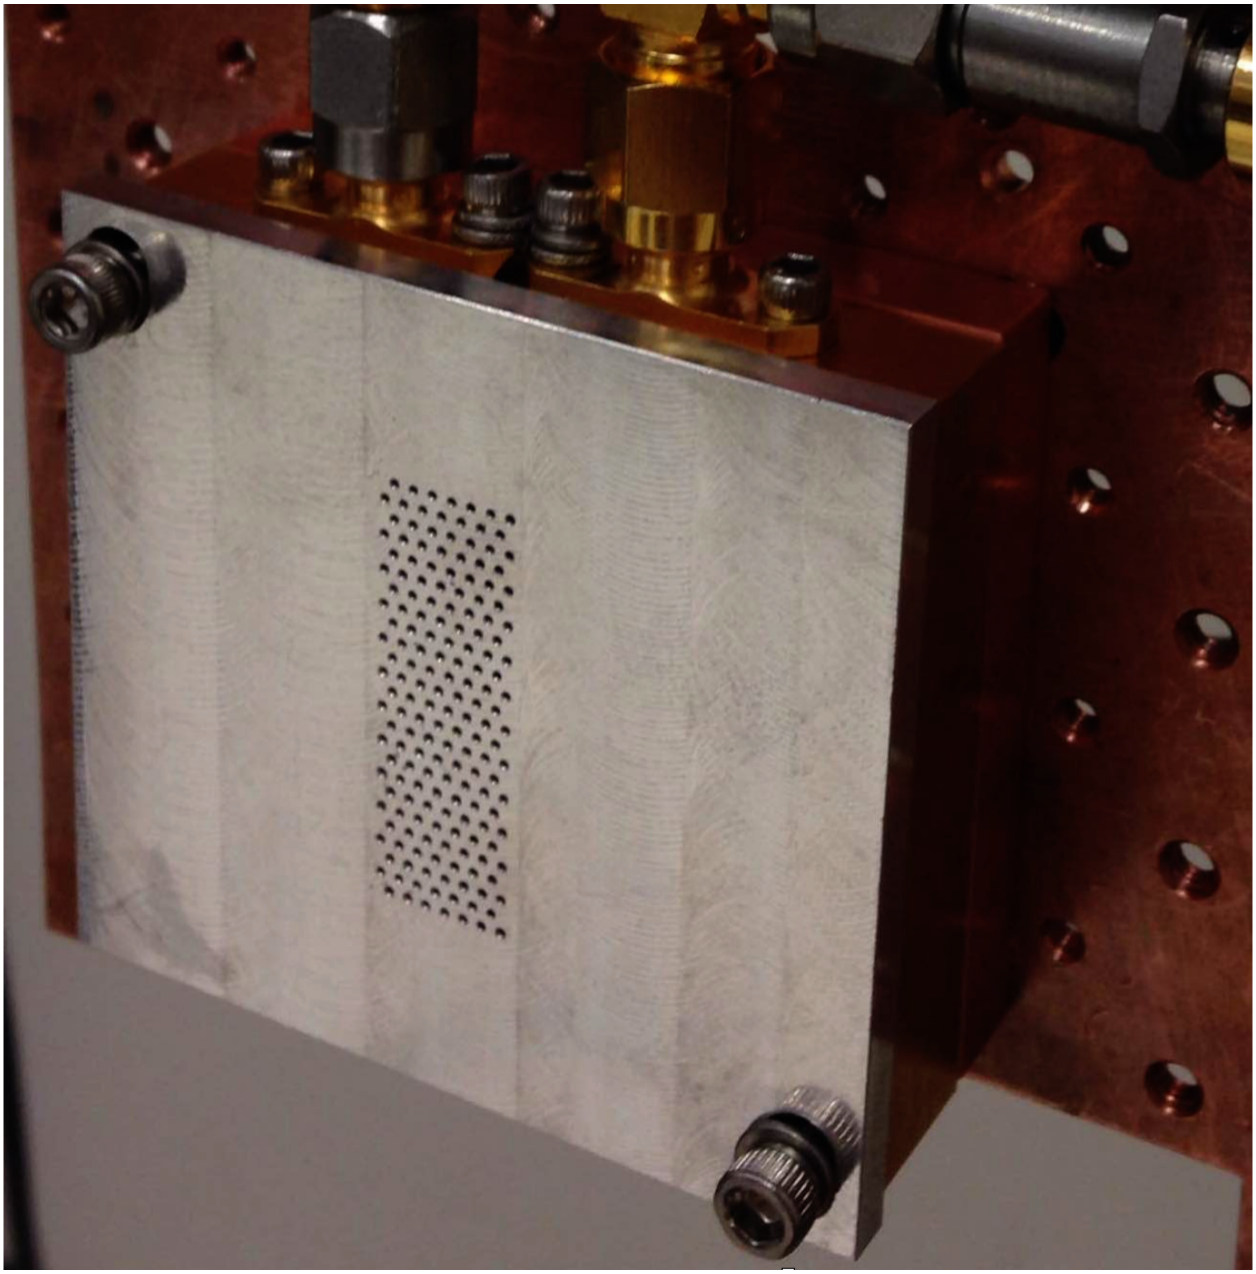

Figure 4.8: An image of the front panel designed in an attempt to mitigate the non-uniform illumination issue of the primary tests at NIST of the back-thinned array.

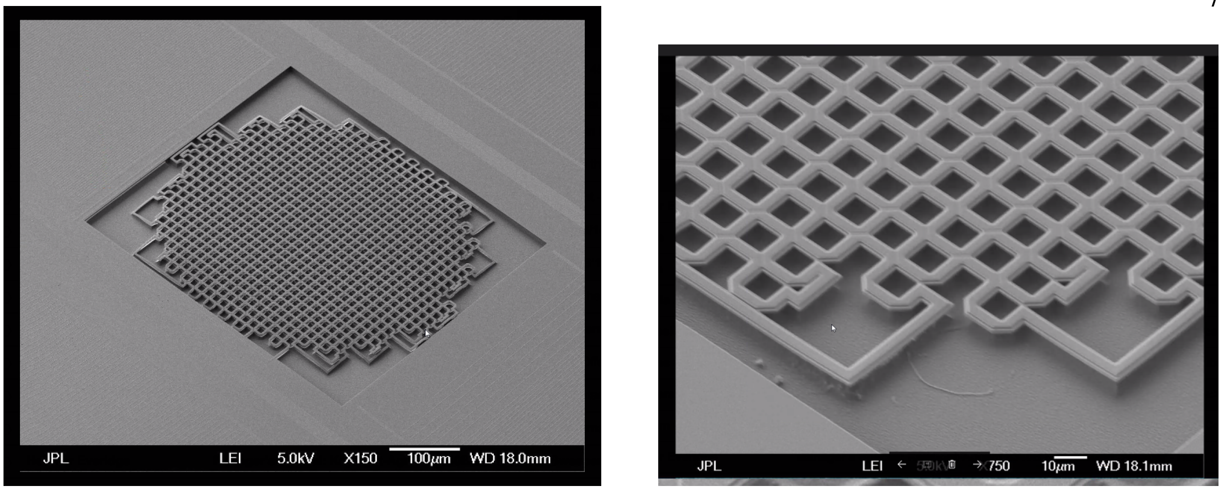

Figure 4.9: Electron microscope pictures of the released inductor taken at JPL. *Left* shows a picture of the full inductor, while *right* zooms in on a corner of the inductor.

the possibility of having at least one functional detector, no detectors were left in the un-released geometry.

The basic geometry of the phonon recycling simulations do not apply to this new geometry. Therefore, we can not predict the expected phonon recycling factor using our model. A lower limit is provided by noting the filling fraction is less than 10^{-2}, predicting a phonon recycling factor greater than 20.

The first fabrication of this array showed no resonances upon testing, likely because the TiN was damaged during the intricate fabrication processes. Upon second fabrication, initial tests of the array at JPL completed by Peter Day and Byeong-Ho Eom revealed shallow resonators. This test was conducted in a cryostat setup wherein two different arrays are read out over a single set of coaxial lines. The input signal is split inside of the cryostat, with one line transmitting signal to the released array and its cryogenic amplifier. The second line transmits signal to a separate KID array for testing and its associated cryogenic amplifier. After amplification, the signals are re-combined and transmitted to the data acquisition system. We suspected that the observed shallow resonators, with measured Q's $\ll$ designed, were in fact *reflected* transmission features from the second array.

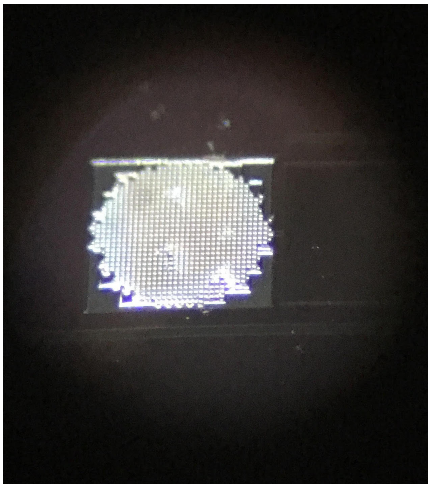

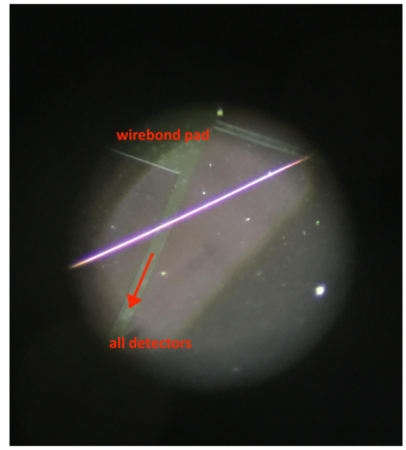

Figure 4.10: Images of the damage to the suspended inductor array. The *left* shows an image of the detector taken through a microscope at CU, showing dents and malformations on the inductor surface. The *right* shows a crack (the thin, bright line) running across the transmission line (the lighter color green).

4.6.1 The Importance of Good Shipping Materials

The suspended inductor array was sent to CU for testing once our cryostat was fully functional in late 2017. We intended to test whether the shallow resonances were intrinsic to the array or were caused by reflected transmission features from the other array. Upon shipment, the delicate array dislodged from its detector box and was damaged. There were cracks in the Si across the feedline and clear warps and dents on the released inductors (Figure 4.10). Despite the damage, the array still appeared to transmit signal. Upon cooling the chip no resonances were observed despite observed transmission through the chip, confirming that the shallow resonances observed at JPL were indeed likely due to reflections from the other array.

There was no funding remaining after this second chip fabrication to make another attempt, and as such we continued along other avenues of NEP reduction for which funding remained.

4.7 Summary

The simulations of back-thinned phonon recycling devices are extremely promising, with exponential or beyond exponential dependencies of P_{rf} on perimeter coverage fraction or Si thickness below the inductor. Boosts in τ_{qp} of over 20 times should be possible through geometric changes well within current fabrication possibilities.

Initial back-thinned devices did not show any boosting in τ_{qp} resulting from the changed geometry. Rather, the testing setup resulted in non-uniform illumination of the array that prevented measurement of the simulated 50% boost in τ_{qp} given this geometry.

Chapter 5

Dependency of Quasiparticle Recombination Time on Metal Film Thickness for Al CPW KIDs

5.1 Introduction

In this work, we focus on increasing the responsivity R_x (Equations 2.11 and 2.12) of low-volume aluminum KIDs as a step towards low NEPs with an approach that is extensible to detectors sensitive to shorter wavelengths of light. Our approach is to minimize the detector's active volume V while maximizing the quasiparticle recombination time τ_{qp} since $R_x \propto \tau_{qp}/V$. Time constants as long as a few ms have been observed in Al KIDs (Barends et al. (2008), de Visser et al. (2011); thicknesses of 250 nm and 40 nm on sapphire, respectively), yet the minimum Al thickness required to achieve these lifetimes has not been quantified directly. Therefore, to optimize our design of low volume and long τ_{qp} we aim to characterize τ_{qp} as a function of metal film thickness. For this characterization, we measure τ_{qp} for Al $\lambda/4$ grounded coplanar waveguide resonator (CPW) test arrays (17 resonators each), with metal film thicknesses of 20, 30, 40, and 50 nm ($\pm$ 10%). To test the influence of quasiparticle diffusion into the ground plane we measure τ_{qp} for an ungrounded 40 nm thick $\lambda/2$ CPW array.

5.2 Testing Layout, Data Analysis, and Time Constant Fitting

5.2.1 Testing Layout

To measure the quasiparticle recombination times τ_{qp} of the detectors, we flash the array with light emitted by a room temperature infrared LED (Thorlabs 1200L and 1600L; λ_{peak} = 1200 and 1600 nm) and observe the resultant decay in frequency shift $\Delta f/f_0$. We use an infrared LED (Thorlabs 1600L) to ensure that the silicon electrons are not excited into the conduction band. These photoelectrons, excited by light with $\lambda < 1100$ nm, can affect measurements of τ_{qp} due to their own intrinsic time constant after which they recombine with a hole and produce phonons. The light from the LED is transmitted to the detectors through a fiber optic cable (Thorlabs M53L01) strung from 300 K to 100 mK with thermal sinking at 4 K and 1 K. The fiber optic setup allows for flexibility in testing due to the ability to mount any room-temperature LED to the fiber optic.

Our cryostat employs a box-in-a-box setup modeled after the dark, low-noise configuration at SRON (Baselmans et al., 2012) described in Section 2.2.1. Two arrays are mounted in separate detector boxes that are then attached to the base of a larger 100 mK box which is designed to be light tight. The fiber optic attaches to the lid of the outer 100 mK box $\sim$5 cm away from the resonators, which are not optically coupled to the light by lenses or feedhorns. Therefore, we assume uniform illumination of the array and the ground plane, which should prevent the diffusion of quasiparticles into the surrounding ground plane from affecting our τ_{qp} measurements. This assumption is corroborated by measurements of τ_{qp} that reveal both grounded and un-grounded resonators have the same value of τ_{qp} (Section 5.4).

For sufficient signal-to-noise ratio (S/N), we flash the detectors with 180 sequential pulses of light and average the resonator response. To ensure an accurate accounting of the decay time, we take 2-4 measurements for all combinations- of metal film thickness, resonator f_0 and corresponding resonator geometry, LED voltage, LED type, detector temperature T_{det}, and microwave drive power.

Prior to fitting the resonator response to the LED pulse we convert the streams of $I(t)$ and $Q(t)$ to fractional frequency shifts away from the resonance ($\Delta f/f_0$). This conversion is described in Chapter 2, Section 2.2.3. In brief, the first step is to subtract out the gain variation (slow drifts of gain with time) and cable delay (difference in phase caused by different cable lengths of I and Q cables) using wide and fine sweeps about the resonance feature. The resulting corrected $I_{corr}(t)$ vs $Q_{corr}(t)$ is a round resonance feature in IQ space centered at the origin with the on-resonance noise lying on the I axis. Using the corrected finely sampled resonance feature the conversion between phase $\theta = tan^{-1}(Q_{corr}/I_{corr})$ and frequency is found with a linear fit at the resonant frequency f_0, the frequency at which we took the noise stream. Finally, the fit line is used to convert the data streamed during the LED pulses from phase to frequency. The fractional shift away from resonance $\Delta f/f_0$ over time is then plotted for the averaged pulse.

5.2.2 Mathematical Model for Non-Exponential Decay

When a pulse of light is absorbed by a KID, the expected signal after the LED turns off is an exponential decay to zero of fractional frequency shift ($\Delta f/f$) over time of the form $e^{-t/\tau_{qp}}$. However, in practice this is altered because the rate of quasiparticle recombination is dependent on the quasiparticle density. High densities of quasiparticles result in faster recombination times because quasiparticles are able to find a partner to recombine with more easily, and the inverse is true for lower densities of quasiparticles. Consequently, as quasiparticles recombine after a pulse of light, the quasiparticle density naturally decreases and hence the signal can appear to decay more quickly at first before it slows due to the lower quasiparticle density. This results in a deviation from a single exponential decay of $\Delta f/f$ over time.

Our measurements showed a deviation from a single exponential (see Figure 5.1), showing instead fast and then slow decay, motivating the creation of a new mathematical model for fitting for τ_{qp}.

To characterize decays in fractional frequency shift ($\Delta f/f$) in KIDs that absorb pulses of light, we begin with the rate equation governing quasiparticle recombination:

$$\frac{dN_{qp}(t)}{dt} = -\gamma_r N_{qp}(t) + I_{qp}(t)\,, \tag{5.1}$$

where $N_{qp}(t)$ is the number of quasiparticles with time, γ_r is the recombination rate, and $I_{qp}(t)$ is the quasiparticle generation rate with time.

The recombination rate γ_r, as mentioned previously, increases with the number of quasiparticles. Additionally, experiments show that it is never zero – rather, it decreases to some fiducial

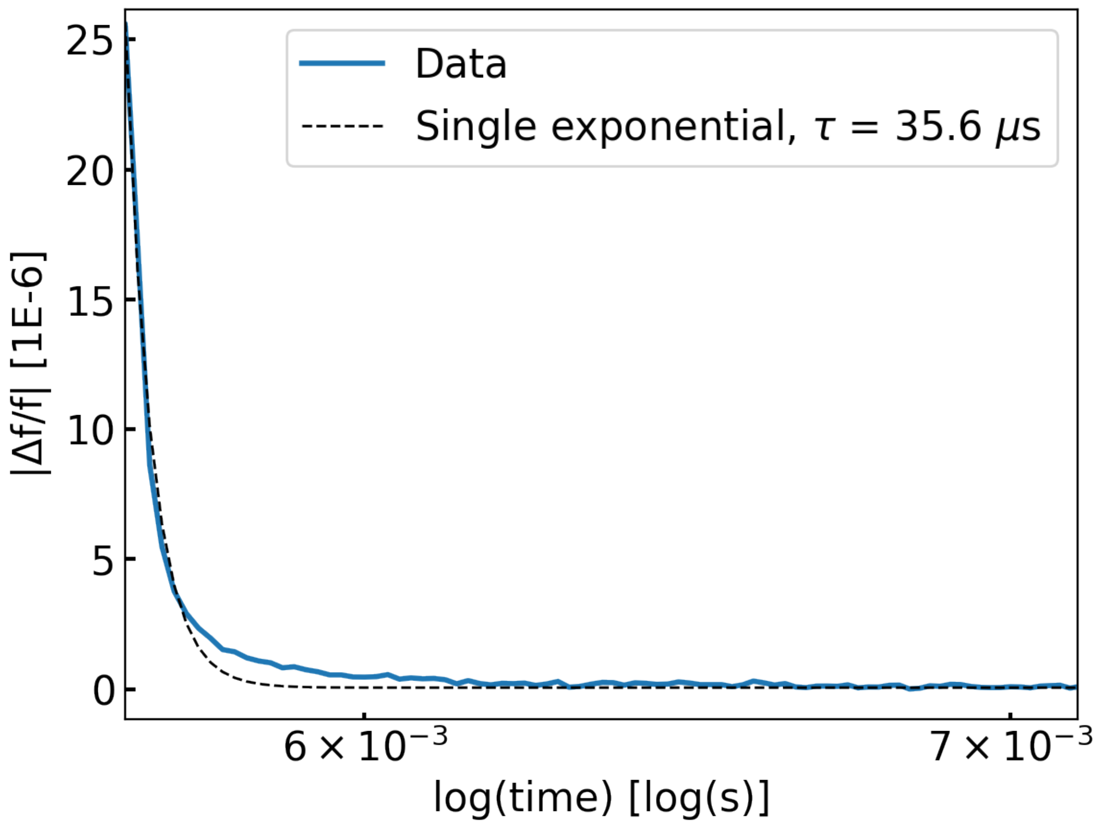

Figure 5.1: Data (*blue*) showing a typical decay of $\Delta f/f$ vs time for a detector at 100 mK, and a single exponential fit (*dashed black*). The exponential captures the quick decay at the beginning of the pulse, but fails to capture the slower decay at the tail end.

value at a characteristic quasiparticle density. This behavior is similar to that of τ_{qp} observed in experiments (Barends et al., 2008), which follows the form

$$\tau_{qp} = \frac{\tau_{max}}{1 + n_{qp}/n_*}. \tag{5.2}$$

Here, n_{qp} is the thermal quasiparticle density and n_* is a fiducial density below which τ_{qp} = τ_{max}. Using this observed behavior of τ_{qp}, we can write

$$\gamma_r = \frac{1}{\tau_{max}} \left(1 + \frac{n_{qp}}{2n_*}\right). \tag{5.3}$$

The factor of 2 in the denominator is introduced for consistent definitions of n_* between Equations 5.2 and 5.3. It can be found by combining Equations 5.3 and 5.1 and integrating for a small perturbation ΔN to a steady-state quasiparticle population of N.

Noting that n_{qp}/n_* = N_{qp}/N_* where n_{qp} and n_* are densities in a volume V and N_{qp} and N_* are total numbers in a volume V, we can combine Equations 5.3 and 5.1 to write

$$\frac{dN_{qp}(t)}{dt} = -\frac{1}{\tau_{max}}[1 + \frac{N_{qp}(t)}{2N_*}]N_{qp}(t) + I_{qp}(t). \tag{5.4}$$

In our testing scenario, we deliver a quick pulse of light to a test device and then return to a dark state, meaning I_{qp} = 0 during the decay of $\Delta f/f_0$ back to zero. By setting I_{qp} = 0 in Equation 5.4 and introducing the variable $x_{qp} = N_{qp}(t)/2N_*$ for ease of calculation, we come to our final governing rate equation of

$$\frac{dx_{qp}(t)}{dt} = -\frac{1}{\tau_{max}}[1 + x_{qp}(t)]x_{qp}(t). \tag{5.5}$$

By integrating and rearranging terms the rate equation can be rewritten as

$$x_{qp}(t) = \frac{1}{[1 + 1/x_{qp}(0)]\exp(t/\tau_{max}) - 1}. \tag{5.6}$$

The physical meaning of the initial value $\mathrm{x}_{qp}(0)$ that is a byproduct of integration can be verified by examining the limits of Equation 5.6 for $\mathrm{x}_{qp}(0) \ll 1$ and $\gg 1$. When $\mathrm{x}_{qp}(0) \ll 1$ we can approximate Equation 5.6 as an exponential decay with a characteristic decay time τ_{max}. This is the same form that results from Equation 5.4 for low quasiparticle density. When $x_{qp}(0) \gg 1$ we can approximate Equation 5.6 as

$$x_{qp}(t) \approx \frac{1}{\exp(t/\tau_{max}) - 1} \,. \tag{5.7}$$

This behaves approximately as $1/t$ for times $t \ll \tau_{max}$ and as an exponential decay for $t \gg \tau_{max}$.

Interpreting these limiting cases, it appears that our interpretation of the integration constant $x_{qp}(0)$ as a measure of the starting values of quasiparticle densities is correct. For high values ($\gg$ 1), the decay will deviate from an exponential. For low values ($\ll$ 1) the decay will behave almost exactly like an exponential. Importantly, the characteristic time τ_{max} is the same in both cases, indicating that τ_{max} is the recombination time τ_{qp}. Therefore, by fitting Equation 5.6 to the decay of signal in a KID that has absorbed a pulse of light, we can characterize the recombination time.

When fitting Equation 5.6 to measured data, we must convert from the physical values of $x_{qp}(t)$ to a measured parameter X(t) (in our case $\Delta f/f$) which should be a linear conversion, meaning

$$X(t) = \frac{A}{[1 + 1/x_{qp}(0)]\exp(t/\tau_{qp}) - 1} \,. \tag{5.8}$$

However, setting t = 0 in this equation we see that X(0) = $x_{qp}(0)$ A, meaning the intial amplitude of the observed pulse constrains the product of $x_{qp}(0)$ and A but not each independently. Our initial analysis found this equation to produce highly correlated and poorly constrained values for both $x_{qp}(0)$ and A. To mitigate the effects of the correlation, we instead define $B \equiv x_{qp}(0)$ A and come to our final equation that we use for all fits to data

$$X(t) = \frac{B}{[1 + x_{qp}(0)]\exp(t/\tau_{qp}) - x_{qp}(0)} \,. \tag{5.9}$$

We find that this reduces the correlation between the fit parameters B and $x_{qp}(0)$ as well as the error on B as compared to the parameterization presented in Equation 5.8. Correlations between parameters are shown in Figure 5.2, displaying the lack of correlation between $x_{qp}(0)$ and B as well as between τ_{qp} and B. A correlation exists between τ_{qp} and $x_{qp}(0)$, but the values of each are well constrained within their parameter spaces. We find that this parameterization captures the decays much more accurately than a single exponential.

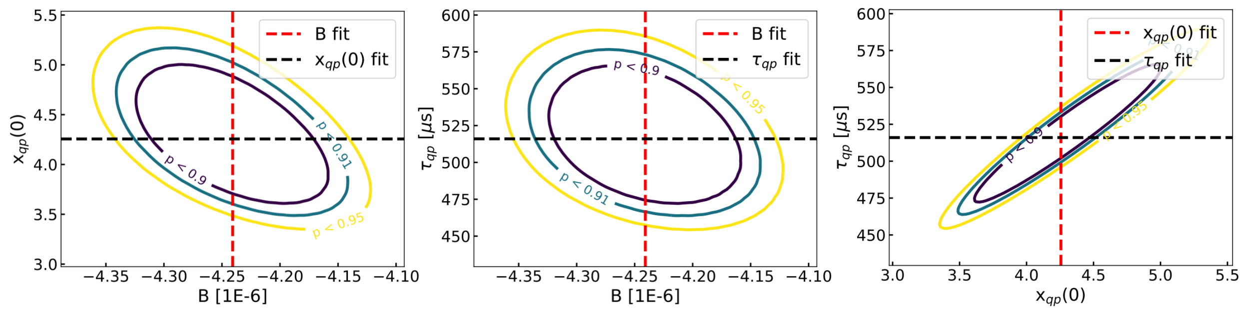

Figure 5.2: Results of a $\Delta\chi^2$ analysis for fitting Equation 5.9 to data from a 100 mK resonator. χ^2 is minimized over over τ_{qp} (*left*), $x_{qp}(0)$ (*center*), and B (*right*). Contours show the $\Delta\chi^2$ values corresponding to one, two, and three σ uncertainties in the parameters. Fit values for parameters are shown as dashed lines.

Daniel Flanigan's thesis (Flanigan, 2018) also presents a non-exponential parameterization to fit the optical decays to find τ_{qp}. We find no difference in the fit τ_{qp} between Equation 5.9 and this parameterization, indicating that they are probing the same physics.

5.3 Simulated Data and Fitting Technique

To find biases in our fitting technique, we generate simulated decays mimicking observations and analyze them using our X(t) fitting Python code. The simulated datasets also test our understanding of the underlying signal and the noise properties of the real decays.

The simulated decays are created using Equation 5.9 with user-defined $x_{qp}(0)$, B, and τ_{qp} and Gaussian-distributed noise, a 60 Hz sine wave, and pink (1/f) noise added to match the observed noise. The noise amplitudes and standard deviations σ are user-defined multiples of the measured $\sigma_{\Delta f/f_0}$ from the tail of a real decay, chosen to match observed noise levels. An example simulated decay is shown and compared to real data in Figure 5.3. The simulated data shows excellent correspondence to the real data in both time streams and power spectral densities (PSDs) indicating accurate and representative simulated data.

Figure 5.3 also shows an example fit to the simulated data. The fitting code in this example case does not accurately re-capture the simulated τ_{qp}, finding a best fit value of 5.4 ms for a

Figure 5.3: *Top Left:* An example of simulated data with the final simulated decay shown in orange. The base of the signal is a generated X(t) (Equation 5.9) shown in blue with the contributing random noise sources added to this base X(t) shown in other colors. Each curve is offset for visualization purposes on this plot. *Top Right:* An example of the Python code's fit of X(t) to the generated data from the left hand plot along with the residuals between the data and the fit. *Bottom Left:* a comparison between a real decay (in maroon) and 200 simulated data sets (in faded colors), along with their average (in dashed black) for a 40 nm thick resonator at 140 mK. *Bottom Right:* a comparison of the power spectral density (PSD) of the real and simulated decays plotted in the bottom left panel.

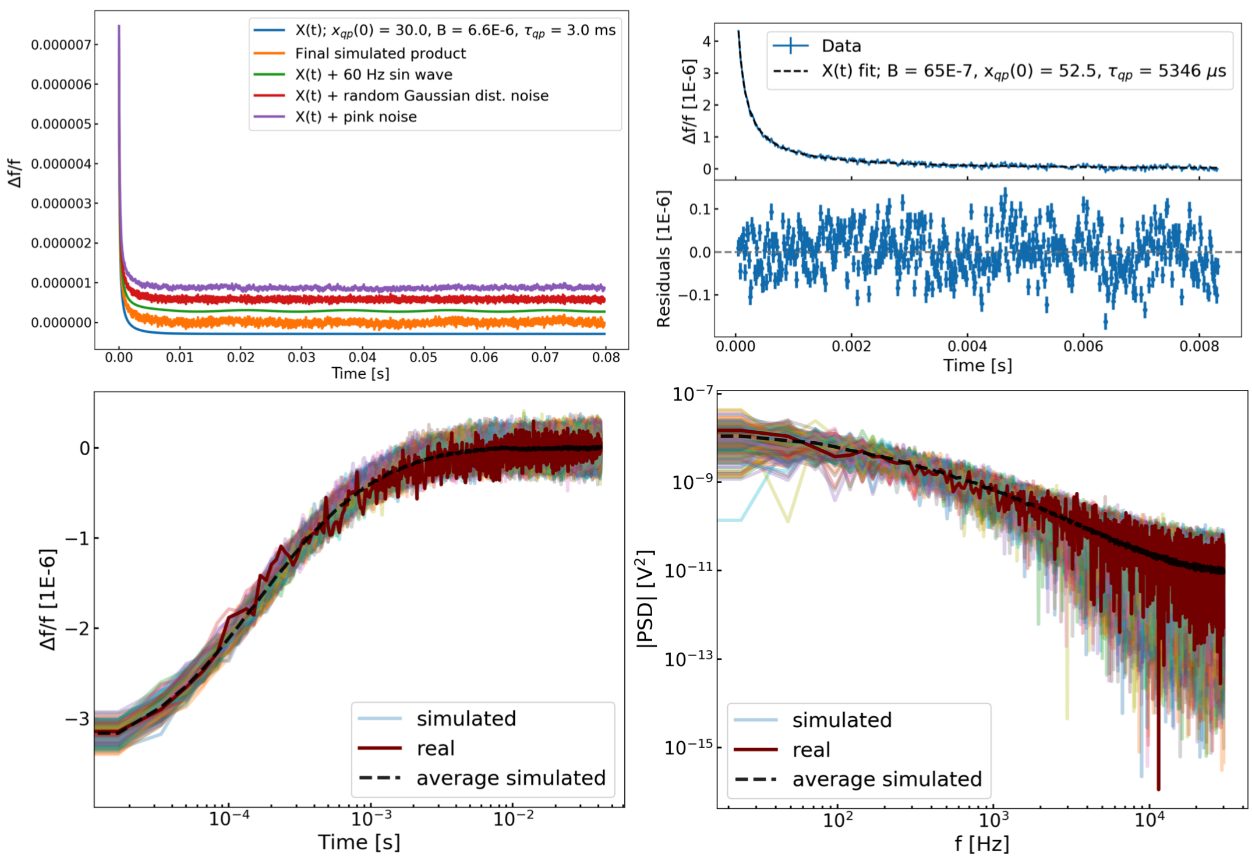

Figure 5.4: *Left:* X(t) fit parameters for three simulated datasets for a variety of fit lengths (number of data points used in the fit). The data are generated with τ_{sim} = 2.5 ms, $x_{qp}(0)_{sim}$ = 20, and B_{sim} = -3.6E-6 to match the real data, as shown in the lower left of Figure 5.3. *Right:* the same, but for three real decays of a 40 nm thick resonator at 140 mK. *Black dashed line* shows where the data (real and simulated) begins to be 1/f noise dominated, demonstrated by the increased dispersion between data sets as the number of data points used in the fit increases.

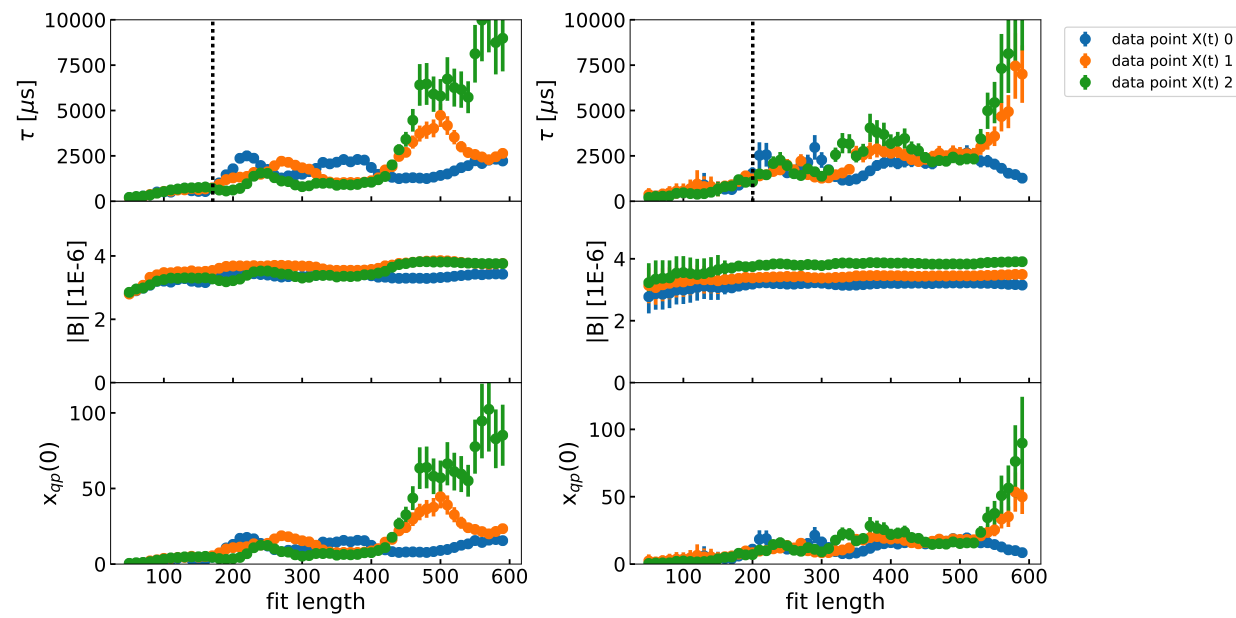

Figure 5.5: Histograms of the recovered τ_{qp} for three different simulated τ_{qp} = [1, 2, 3] ms. τ_{meas} is determined using the location where 1/f noise begins to dominate in the τ_{qp} vs fit length plots (vertical black dashed lines in Figure 5.4).

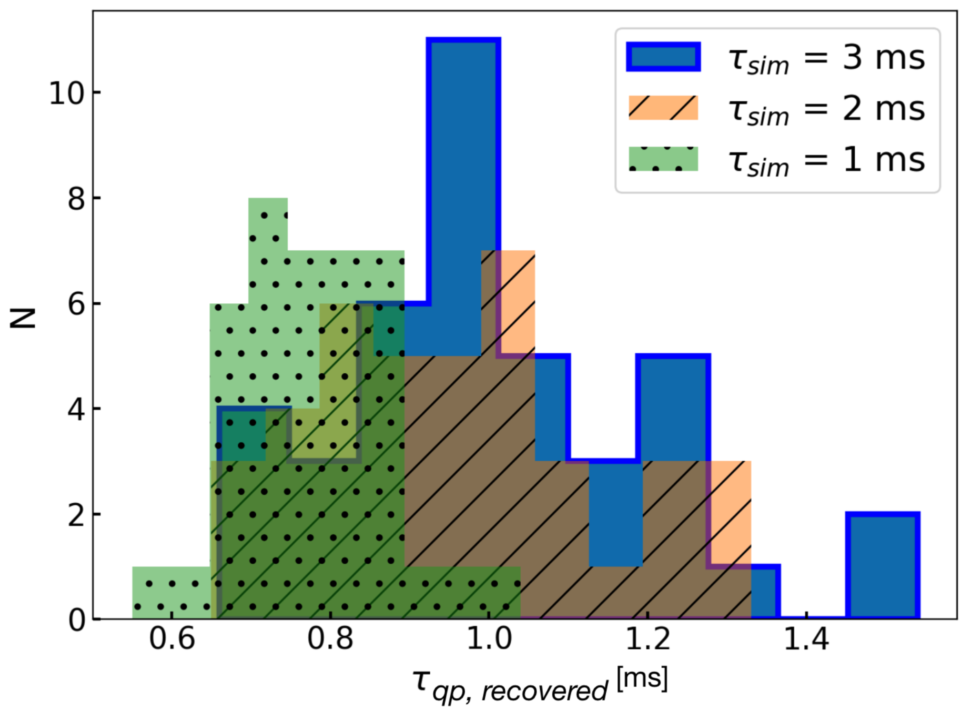

Figure 5.6: The median fit τ_{qp} (*left*) and mean fit τ_{qp} (*right*) for a simulated τ_{qp} = [1, 2, 3] ms as a function of the number of data points used in the fits. For each fit length, we generate 500 simulated decays and fit X(t) to each of these decays to find $\tau_{qp,recovered}$ for each. We then calculate the mean and median $\tau_{qp,recovered}$ for a given fit length for the 500 data sets and error bars representing one standard deviation of the distribution of $\tau_{qp,recovered}$. This plot shows that we need $\gtrsim$ 400 data points to accurately recover longer simulated values of τ_{qp}. The error bars on these values of $\tau_{qp,recovered}$, in the limit of many measurements, should capture the expected error bar sizes on physical measurements of τ_{qp}.

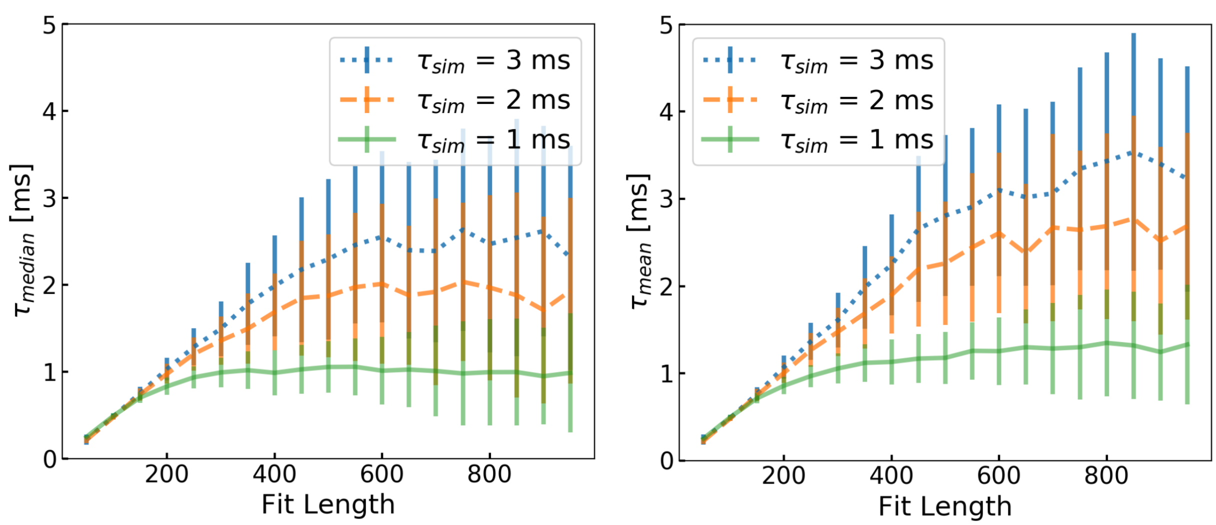

Figure 5.7: Histograms of $\tau_{qp,recovered}$ for three different values of $\tau_{qp,simulated}$ = [1, 2, 3] ms. $\tau_{qp,recovered}$ is determined using the location where dispersion is minimized for maximal fit length in the τ_{qp} vs fit length plots (Figure 5.4).

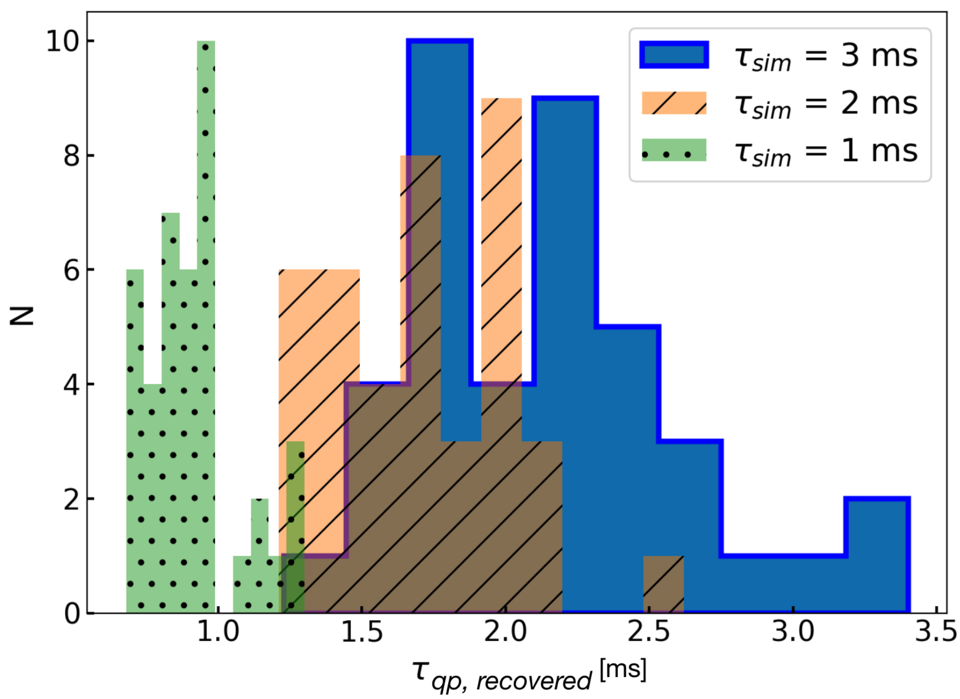

Figure 5.8: The $\tau_{qp,recovered}$ and $x_{qp}(0)_{recovered}$ from measurements of simulations with three different values of $\tau_{qp,simulated}$: [1,2,3] ms (from *left* to *right*). The values of $\tau_{qp,simulated}$ and $x_{qp}(0)_{simulated}$ for each plot are marked with *black lines*. *Gray dots*: recovered parameters for simulations including all possible noise types (pink, Gaussian, and sine wave), with *gray dashed lines* marking the mean recovered values for these simulations. Comparing the mean recovered values to the simulated values, it is clear that both $\tau_{qp,recovered}$ and $x_{qp}(0)_{recovered}$ are lower than the simulated values, with the discrepancy greatest for the highest values of $\tau_{qp,simulated}$. We also explore the influence of the three different noise types on $\tau_{qp,recovered}$ and $x_{qp}(0)_{recovered}$ by including only one noise type at a time in the simulations – *blue stars* show the recovered parameters when including pink noise only, *orange squares* including sine wave noise only, and *green triangles* including Gaussian noise only. The pink noise (*blue stars*) appears as the noise type most responsible for suppressing $\tau_{qp,recovered}$ and $x_{qp}(0)_{recovered}$ with respect to $\tau_{qp,simulated}$ and $x_{qp}(0)_{simulated}$.

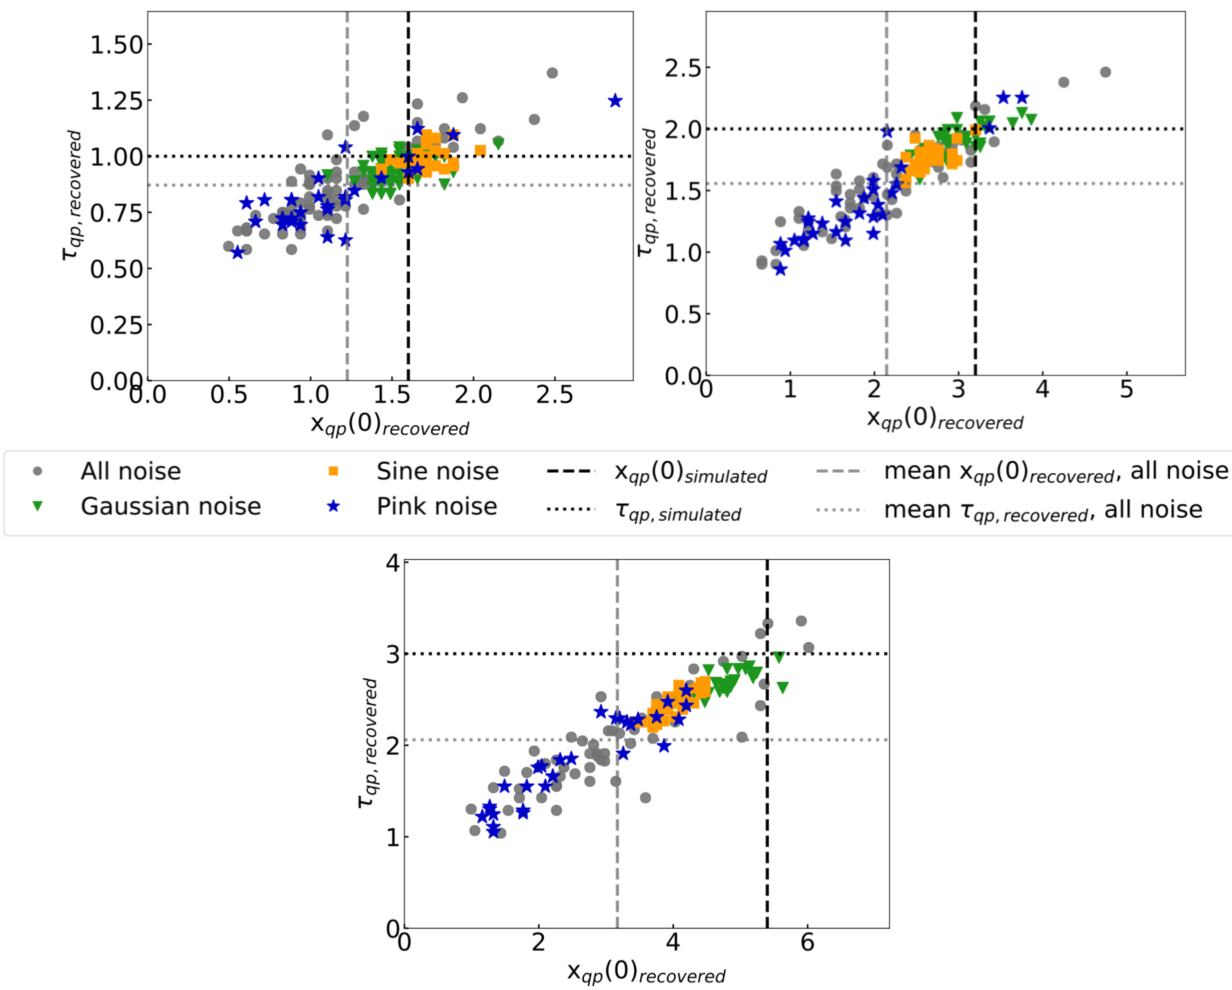

simulated τ_{qp} of 3 ms when we use 500 data points to fit the decay. It was common for the fitter to not accurately re-capture the simulated τ_{qp}, like the example shown in this figure. Therefore we use the simulated data to understand the biases of the fitting technique and to choose the fitting technique that most accurately re-captures the simulated fit parameters. The parameter that most influences the recovered time constant $\tau_{qp,recovered}$ (the τ_{qp} measured from the simulated data) is the length of time stream data included in the fit. Below, we use the simulated data to motivate a strategy for selecting the proper number of data points to use when fitting X(t). Once we choose a strategy for choosing the number of data points included in the fit, we characterize how accurately this fitting technique recreates the simulated decay parameters.

Figure 5.4 shows the fit parameters τ_{qp}, $x_{qp}(0)$, and $|B|$ as a function of the number of data points included in the fit (what we call "fit length") for three data sets of both simulated and real data. We observe similar behavior between most sets of simulated data and most sets of real data: for any set of decays taken under the same conditions we see tightly correlated fit τ_{qp}'s that then begin to diverge as a function of fit length around 150-200 data points. $x_{qp}(0)$ is tightly correlated with τ_{qp} and therefore demonstrates a similar behavior. The fit parameter $|B|$ remains nearly constant as a function of fit length. We postulate this divergence of fit τ_{qp}'s around 150 data points is due to the 1/f noise becoming dominant after this point in the decay. We denote the 1/f noise dominance point in Figure 5.4 with a vertical dashed black line.

Using this observed behavior of fit parameters with fit length we devise and test a fitting strategy: take multiple (3-5) data sets and choose a fit length just prior to the divergence where 1/f noise begins to dominate the decay. This methodology should decrease the influence of the 1/f noise properties on the recovered τ_{qp} and its scatter between data sets. However, when we use this methodology for three simulated values of $\tau_{qp,simulated}$ = [1, 2, 3] ms we find that regardless of the true value of $\tau_{qp,simulated}$ the resultant recovered $\tau_{qp,recovered}$ is only $\sim$1 ms. This effect is illustrated in Figure 5.5, a histogram of $\tau_{qp,recovered}$ for 40 randomly generated decays with $\tau_{qp,simulated}$ = [1, 2, 3] ms using this methodology. We therefore conclude that while this method should reduce the influence of the 1/f noise, it does not use a sufficient number of data points to recapture the true underlying τ_{qp} of the decay.

As a next step in devising how to choose the proper fit lengths for the decays we calculate and plot the mean and median fit value for τ_{qp} as a function of fit length in Figure 5.6 for 500 simulated decays with τ_{qp} = [1, 2, 3] ms. The median fit τ_{qp} accurately recovers the simulated time constants of 1 and 2 ms, while the mean fit τ_{qp} has more accuracy for the longest τ_{sim} of 3 ms. It can be seen from this plot that to more accurately recover $\tau_{qp,simulated} > 1$ ms, we must use longer fit lengths of around 400-600 data points. Notably the typical location of the 1/f dominance is around 200 data points, much shorter than required to accurately recover the simulated τ_{qp}.

For our real data we cannot average between 500 decays to find the mean and median τ_{qp} for a given fit length because each decay takes five minutes to acquire. Each dataset includes a wide and finely sampled sweep of $I(f)$ and $Q(f)$ as described in Section 2.2.3 ($3-4$ minutes total) to allow the conversion of $I(t)$ and $Q(t)$ (40 seconds $-$ 1 minute long) to $\Delta f/f_0$. We have instead $3-4$ data points per temperature, LED voltage, and attenuation for each resonator. We therefore choose our fit lengths based on the τ_{qp} vs fit length as follows: minimize the dispersion between measurements while maximizing the number of data points used in the fit. We choose a maximum possible number of data points of 800 because we found that there was no added constraint of τ_{qp} for fit lengths in excess of this value. This methodology avoids the issue of severely under-predicted τ_{qp} that arises from choosing the data length based on the 1/f dominance point while simultaneously reducing the scatter in τ_{qp} that tends to dominate at longer fit lengths.

Figure 5.7 shows the distributions of $\tau_{qp,recovered}$ for $\tau_{qp,simulated}$ = [1, 2, 3] ms using this methodology with the corresponding mean and standard deviation of $\tau_{qp,recovered}$ for each $\tau_{qp,simulated}$ reported in Table 5.1. The histograms show that this methodology breaks the degeneracy between simulated τ_{qp} compared to choosing the 1/f noise dominance location (Figure 5.5).

Figure 5.8 shows the correlation between $\tau_{qp,recovered}$ and $x_{qp}(0)_{recovered}$ for simulated time constants $\tau_{qp,simulated}$ = [1, 2, 3] ms. In these plots, the influence of the different types of noise is also explored by including only one noise type at a time in the simulation. Our method is shown to consistently under predict longer $\tau_{qp,simulated}$ primarily due to the pink (1/f) noise. The

$\tau_{qp,simulated}$ [ms]	$\mu_{\tau_{qp,recovered}}$ [ms]	$\sigma_{\tau_{qp,recovered}}$ [ms]	$x_{qp}(0)_{simulated}$	$\mu_{x_{qp}(0)_{recovered}}$	$\sigma_{x_{qp}(0)_{recovered}}$
1	0.91	0.16	1.6	1.04	0.56
2	1.70	0.29	3.2	2.07	0.91
3	2.14	0.48	5.4	2.88	1.33

Table 5.1: The simulated τ_{qp} and $x_{qp}(0)$, mean recovered τ_{qp} and $x_{qp}(0)$, $\mu_{\tau_{qp,recovered}}$ and $\mu_{x_{qp}(0)_{recovered}}$, and standard deviation of $\tau_{recovered}$ and $x_{qp}(0)$, $\sigma_{\tau_{qp,recovered}}$ and $\sigma_{x_{qp}(0)_{recovered}}$ using our fitting technique.

Gaussian noise increases scatter in $\tau_{qp,recovered}$ and the 60 Hz sine wave decreases $\tau_{qp,recovered}$ by only $\sim$0.1$-$0.5 ms at the observed level of contamination. We found that increased levels of noise increase both the magnitude of under prediction as well as the scatter in $\tau_{qp,recovered}$.

The under prediction presents as a bias towards lower $\tau_{qp,recovered}$ compared to $\tau_{qp,simulated}$ by as much as $\sim$1 ms in our measurements. The correlation between τ_{qp} and $x_{qp}(0)$ also leads to an under prediction of $x_{qp}(0)$. Similar to the findings in Figure 5.6, we observe that the mean $\tau_{qp,recovered}$ is less accurate for longest τ_{sim} and that the scatter increases with τ_{sim}. This behavior is likely caused by the 1/f noise that will influence long and slow decays more than faster decays.

We do not quantitatively correct for this bias because it is highly dependent on the level of noise, which varies greatly between measurements. In the simulated data presented we took care to match the noise levels (and types) as closely as possible to a set of measured data. However, this matching is not done via quantitative measurements of the noise, but rather through iterative visual comparisons to recreate the measured signal, an inherently qualitative process. It is possible that the bias could be quantitatively corrected via a careful accounting of the bias values based on quantitatively measured noise levels of the data.

5.4 Time constants

We fit X(t) to the measured decays for 20, 30, 40, and 50 nm thick $\lambda/4$ resonators as a function of detector temperature with results shown in Figures 5.9, 5.10, 5.11, and 5.12. Figure 5.13 shows $\tau_{qp}(T_{det})$ for 40 nm thick $\lambda/2$ CPW resonators. The observed behavior of the 20, 30, and 50

Figure 5.9: X(t) fit parameters for two 20 nm thick resonators. Both resonators have a center conductor width of 0.6 μm. Drive powers: 0.25 & 2.5 aW. Each data point represents a decay.

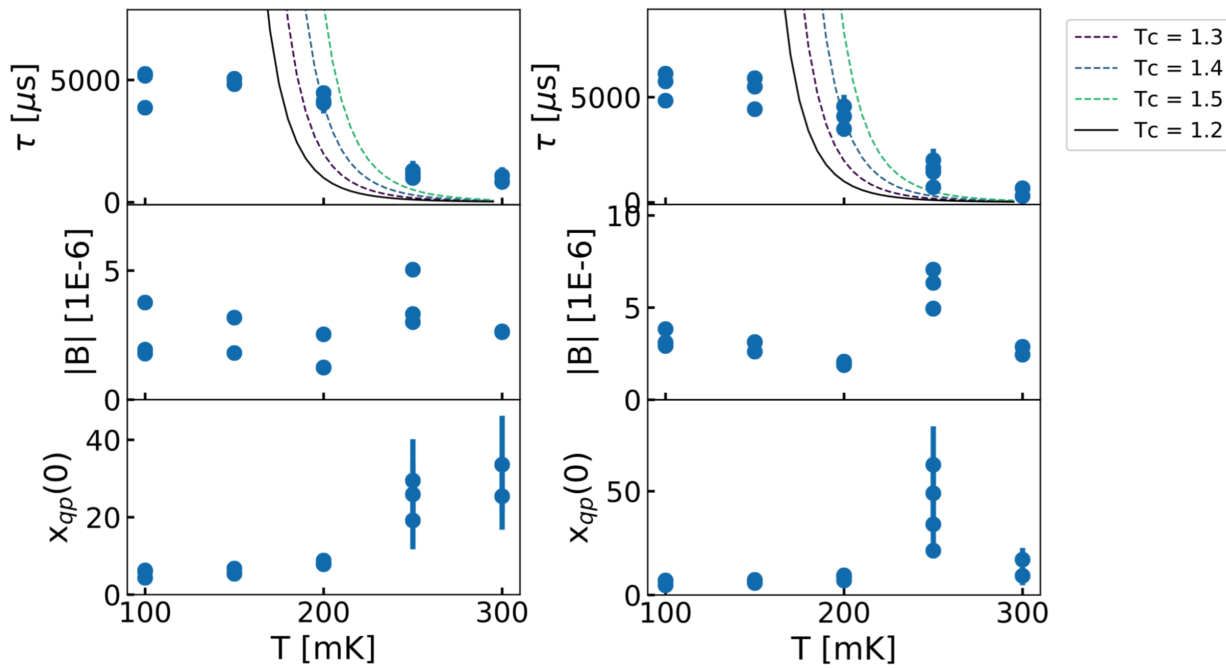

Figure 5.10: X(t) fit parameters for two 30 nm thick resonators. Both resonators have a center conductor width of 0.6 μm. Drive powers: 0.4 aW. Each data point represents a decay.

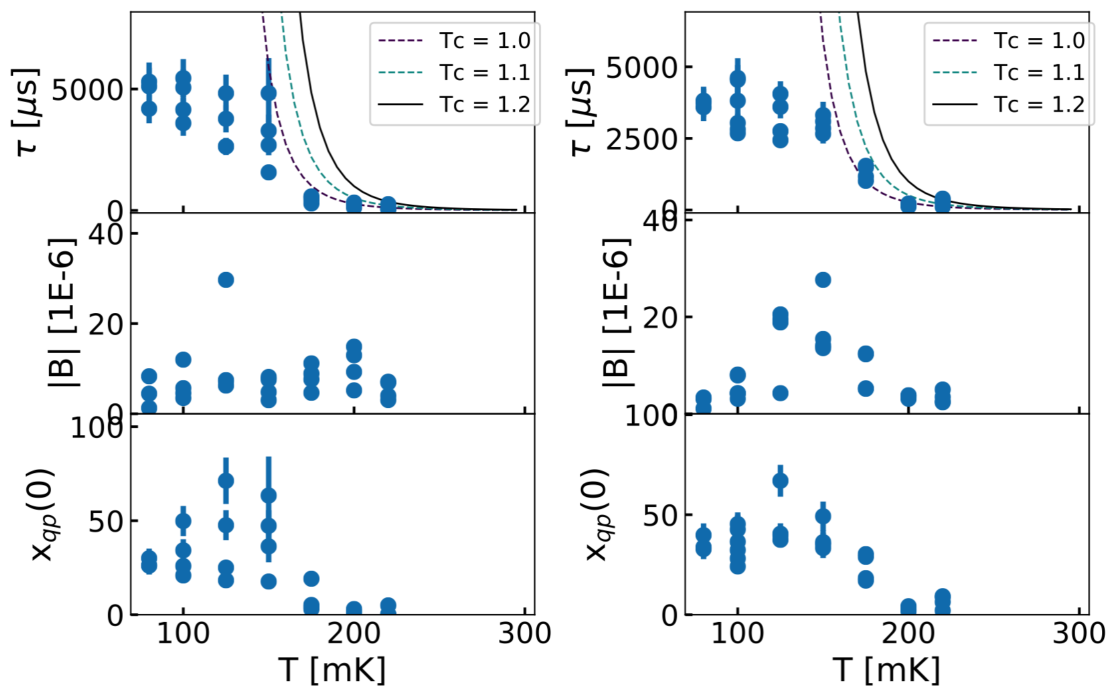

Figure 5.11: X(t) fit parameters for two 40 nm thick resonators. The *left* plot is for a resonator with a center conductor width of 0.6 μm, while the *right* plot is for a resonator with a center conductor width of 1.5 μm. Drive powers: 0.1 & 7.9 aW. Each data point represents a decay.

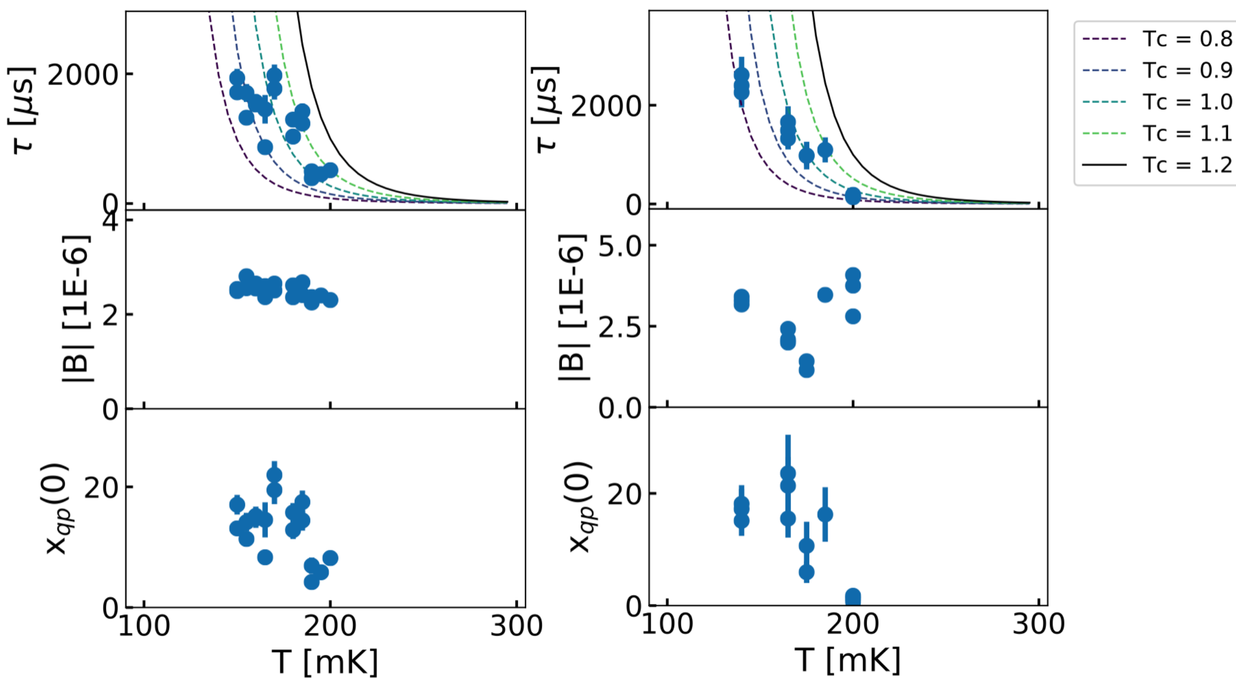

Figure 5.12: X(t) fit parameters for a 50 nm thick resonators with a center conductor width of 1.5 μm. *Left:* drive power of 0.32 aW, *Right:* drive power of 2.5 aW. Each data point represents a decay.

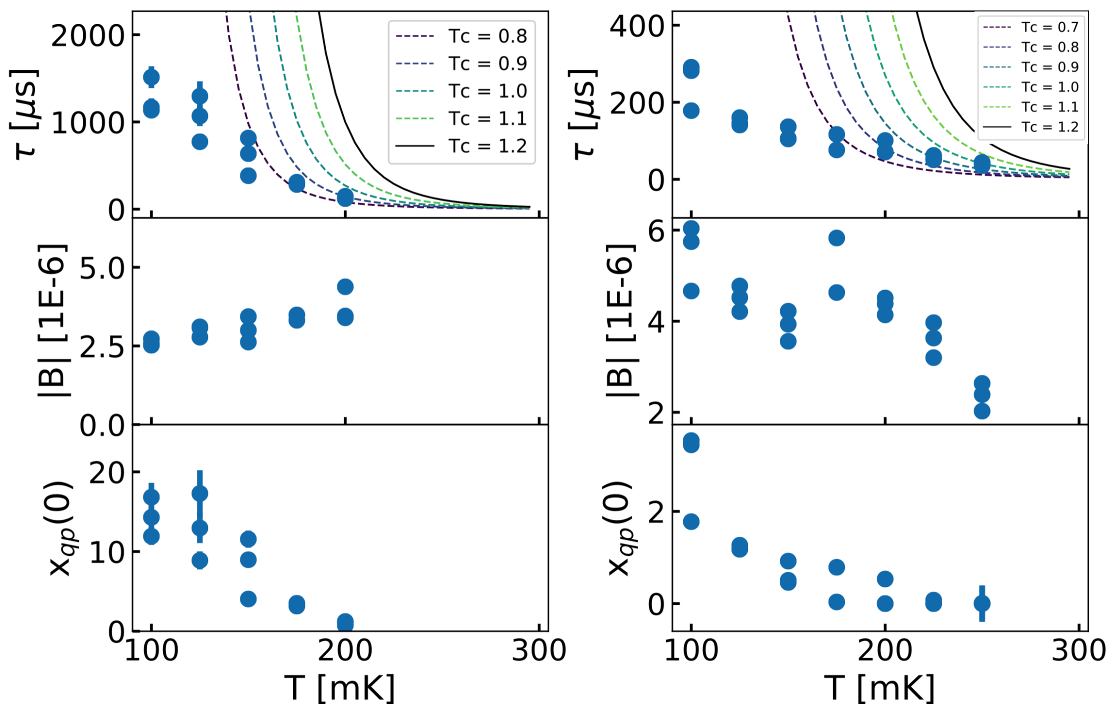

Figure 5.13: X(t) fit parameters for three 40 nm thick 1/2 λ resonators. The *left* plot is for a resonator with a center conductor width of 0.6 μm, while the *middle* and *right* plot are for resonators with center conductor widths of 1.5 μm. Each data point represents a decay.

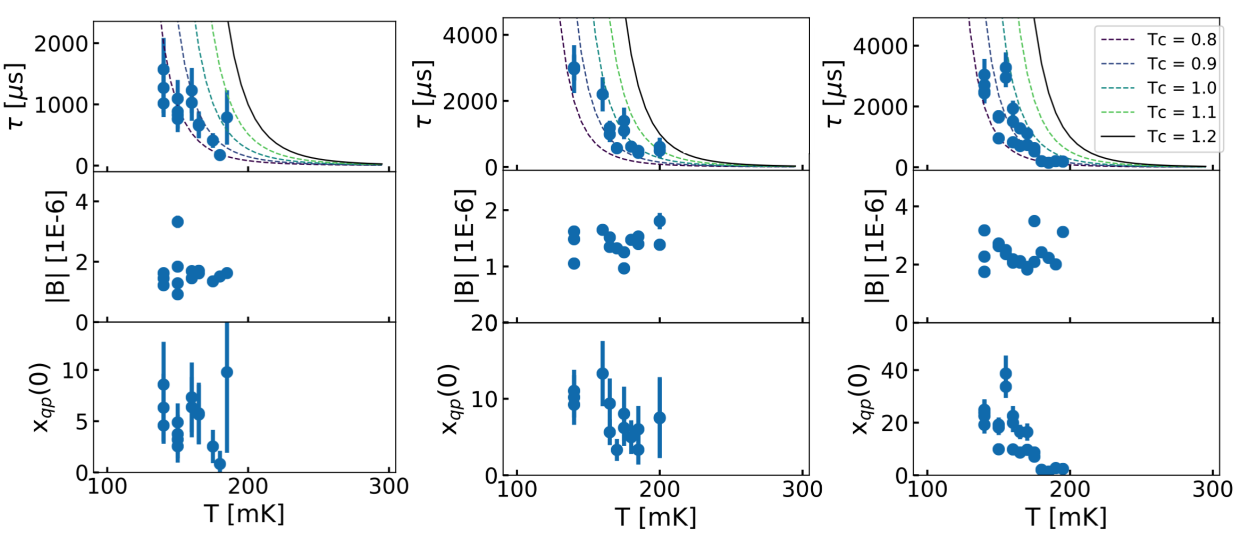

Figure 5.14: The $\Delta f/f_0$ decays at the lowest temperatures for the 20, 30, 40, and 50 nm thick resonators, whose X(t) fit parameters are plotted in Figures 5.9, 5.10, 5.11, and 5.12. This plot averages together the separate decays shown as individual data points in the aforementioned figures for simplified visualization.

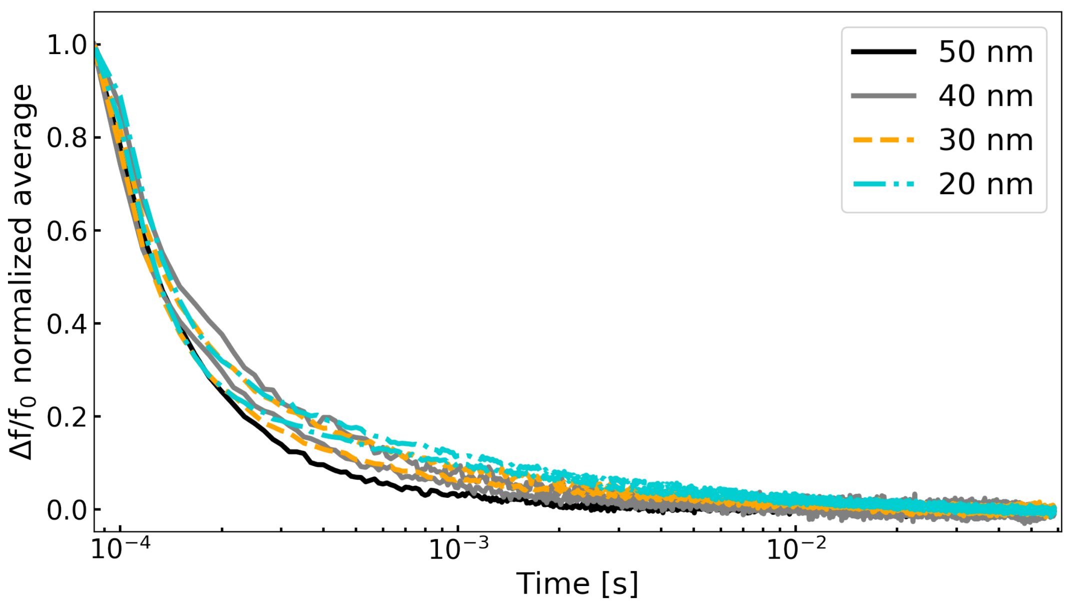

Figure 5.15: Examples showing the independence of the decay on the length of the pulse (*top*), and the LED type or LED voltage (*bottom row*). The *top* plot shows decays for three different duty cycles of the function generator corresponding to three different pulse durations. The *bottom left* plot shows a decay for LED type Thorlabs 1200L, with a peak wavelength at 1200 nm. The *bottom right* shows decays for LED type Thorlabs 1600L (peak wavelength of 1600 nm) powered at three different voltages corresponding to three different brightnesses. The decays do not appear to change within the bottom right plot, nor between the bottom left and bottom right plots.

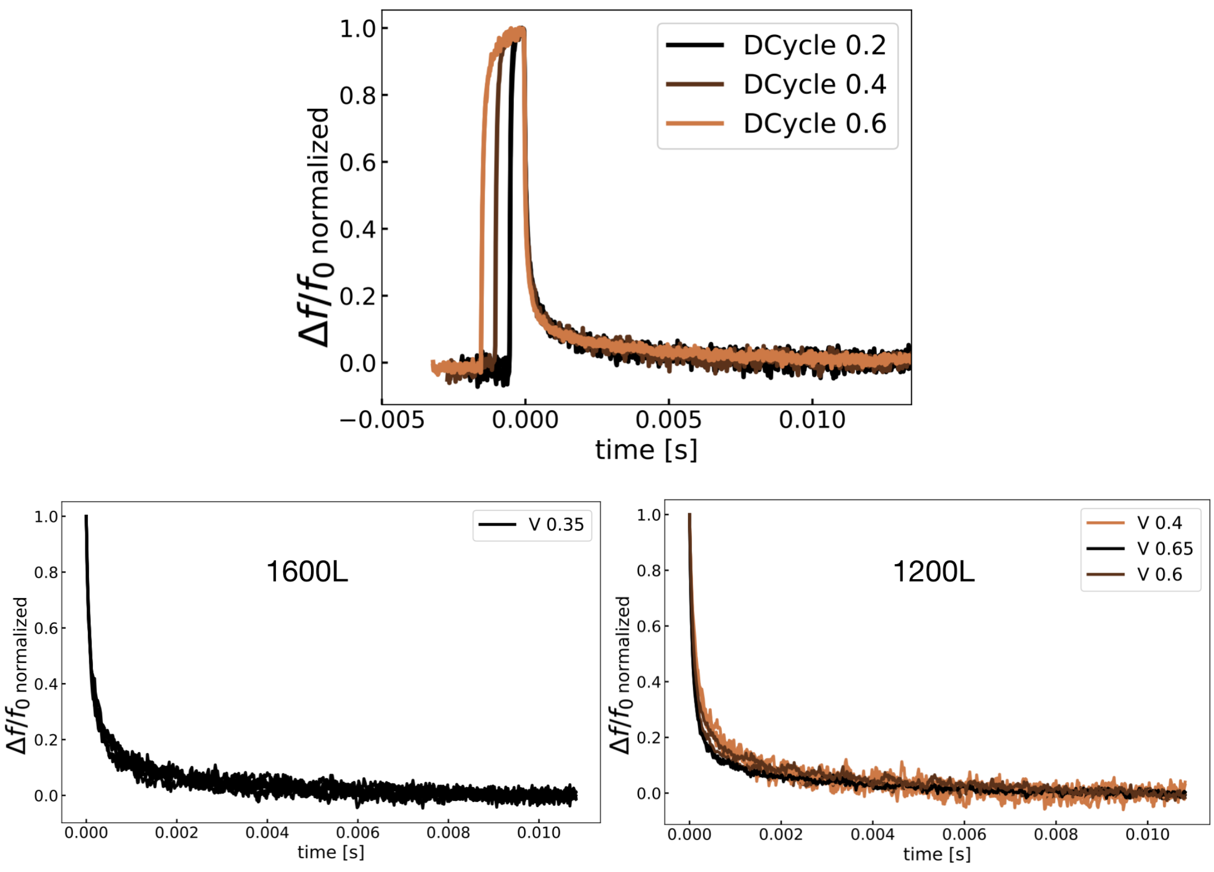

nm thick films mimics that seen in other works (see a summary in Zmuidzinas (2012)), with τ_{qp} flat with T_{det} below T $\sim$ 150-200 mK and decaying above that temperature. The 40 nm thick film was not measured at a low enough temperature to observe this behavior in these measurements. This shape is captured by Equation 5.2 – the flat portion of $\tau_{qp}(T_{det})$ is τ_{max}, and the decay is captured in the dependency of n_{qp} on T_{det}.

We tested the influence of LED voltage, LED type (peak wavelength at 1200 or 1600 nm), and LED pulse duration (up to 4 times longer) on the measured τ_{qp}. No statistically significant change in τ_{qp} was measured for any of these variables. Examples of these tests are shown in Figure 5.15.

The value of τ_{qp} is consistent for the 40 nm thick $\lambda/4$ grounded and the $\lambda/2$ ungrounded CPW. Additionally, we did not observe statistically significant differences in τ_{qp} for different center conductor lengths for a given metal film thickness, a signature that would evidence diffusion because longer resonators should result in longer diffusion times. This indicates that our measurements of τ_{qp} are not affected by diffusion of quasiparticles into the ground plane.

$x_{qp}(0)$ for all decays is $\gg 1$ reflecting our data's deviation from a single exponential due to the high illumination level required to overcome TLS noise. $x_{qp}(0)$ decreases with T_{det} due to its correlation with τ_{qp} (see Figure 5.2). Despite this correlation, the decrease in τ_{qp} remains when $x_{qp}(0)$ is forced to be constant with T_{det}.

5.4.1 Influence of Microwave Drive Power

We also measure τ_{qp} as a function of microwave drive power P_{drive} for 40 and 50 nm thick $\lambda/4$ resonators and 40 nm thick $\lambda/2$ resonators with results shown in Figure 5.16. We measured τ_{qp} across 26 dB of P_{drive} for the 50 nm thick array and 11 dB for the 40 nm array up to 3 dB below bifurcation, the power at which $S_{21}(f)$ develops a discontinuity. For the $\lambda/4$ resonators, τ_{qp} decreases with P_{drive}. This slow decay of τ_{qp} is not observed for the 40 nm thick $\lambda/2$ array we measured, rather, it suddenly drops within 2 dB of bifurcation.

It is not immediately obvious how microwave power of frequency ν_p with energy well be-

Figure 5.16: Power sweeps for our 40 nm (*squares*) and 50 nm (*circles*) thick $\lambda/4$ resonators, and the 40 nm thick $\lambda/2$ resonators (*black triangles*). Different resonators for a given metal film thickness are represented by different colors. The highest power for each resonator is 3 dB below bifurcation except for the $\lambda/2$ device that is 2 dB below bifurcation. Fits using $\tau_{qp} \propto P_{drive}^{\alpha}$ for the data are plotted in corresponding line colors, with *dotted lines* corresponding to 40 nm thick resonators and *dot dashed lines* to 50 nm thick resonators. No fit to the $\lambda/2$ resonator was made because no clear decay with power was seen apart from a sudden drop in τ_{qp} at the highest power, within 2 dB of bifurcation.

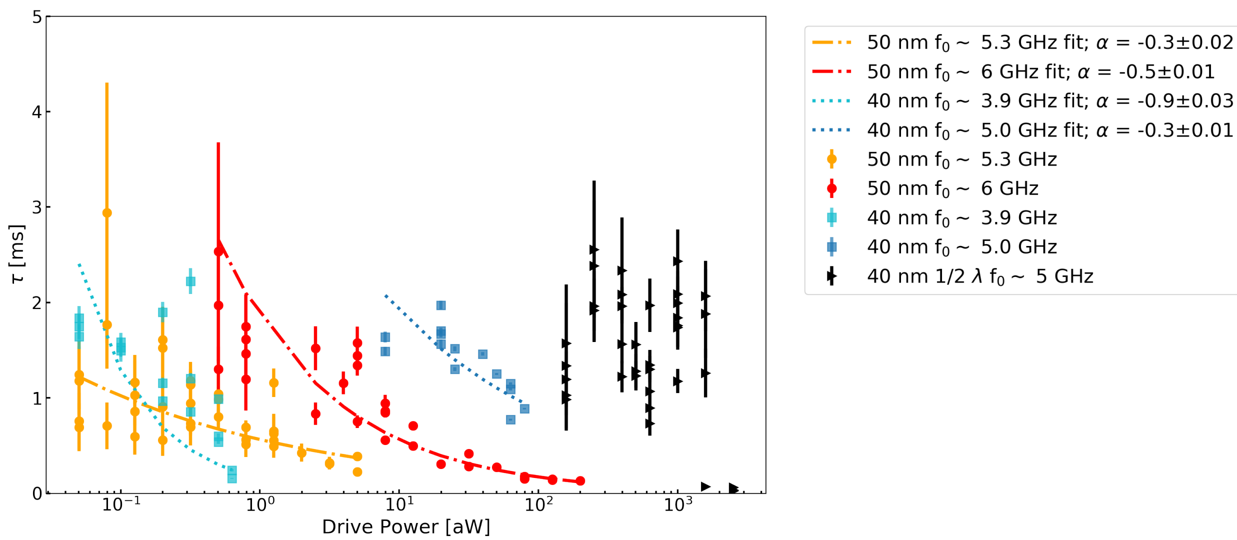

low the binding energy of Cooper pairs ($h\nu_p \ll 2\Delta_0$) suppresses the quasiparticle recombination time. However, this behavior has been observed by other groups in aluminum CPWs (e.g. in de Visser et al. (2014)). The mechanism causing this behavior is numerically modeled in Goldie and Withington (2014). They postulate that high drive powers create excess quasiparticles with $E{>}3\Delta_0$ that then suppress the time constant: quasiparticle populations already extant in the superconductor (primarily thermal quasiparticles) absorb energy from multiple RF photons leading to non-equilibrium populations, some of which have energy $E{>}3\Delta_0$. This excess energy means that the quasiparticles can emit a phonon with energy $>2\Delta_0$ upon scattering. That energetic phonon can then go on to break another Cooper pair and create excess quasiparticles above the levels of the original population.

No direct comparison to Goldie and Withington (2014) can be conducted because their calculations depend on the absorbed power per unit volume P_{abs} [aW/μm^3] whereas experimentally we plot τ_{qp} against P_{read}, the microwave drive power. de Visser 2014's Equation S10 relates P_{drive} to P_{abs} with (de Visser et al., 2014; Zmuidzinas, 2012)

$$P_{abs} = \frac{P_{drive}}{2}\frac{4Q^2}{Q_iQ_c}\frac{Q_i}{Q_{i,qp}}, \tag{5.10}$$

where Q is the total quality factor, Q_i is the internal quality factor, Q_c is the coupling quality factor, and $Q_{i,qp}$ is the quasiparticle quality factor. The value of $Q_{i,qp}$ is not known because Q_i is not limited purely by quasiparticle dissipation. Therefore, we cannot compare directly to Goldie and Withington's numerical predictions. Additionally, we do not *measure* P_{drive} with any equipment, rather we *estimate* it according to our understanding of the microwave losses of the readout electronics, attenuators, and coaxial cables used. It is doubtful that we are closer than 3 dB to the true value of P_{drive}. Instead of comparing directly to Goldie and Withington (2014), we compare to the findings of de Visser et al. (2014) and follow their methodology. They calculate a theoretical expectation of $\tau_{qp} \propto P_{drive}^{-0.5}$ and experimentally find $\tau_{qp} \propto P_{drive}^{-0.2\pm0.05}$ for a 50 nm thick $\lambda/2$ Al CPW resonator at 120 mK.

We fit and plot $\tau_{qp} \propto P_{drive}^{\alpha}$ in Figure 5.16, finding $<\alpha>$ = -0.4 for 50 nm and $<\alpha>$ = -

0.6 for 40 nm. The 40 nm thick resonators were tested at 140 mK due to a He compressor issue while the 50 nm thick were tested at 100 mK. It is possible that the steeper decay for the 40 nm thick resonators is caused by the higher temperature, which Goldie and Withington (2014) showed results in fewer non-equilibrium quasiparticles suppressing τ_{qp}. Therefore, lower temperatures will in turn require lower drive powers to remove the influence of the non-equilibrium quasiparticles and to observe the τ_{qp} intrinsic to the metal. At higher temperatures (the 40 nm thick resonators), the intrinsic, higher τ_{qp} of the metal will become dominant more quickly due to the lesser influence of the non-equilibrium quasiparticles.

Theoretically there should be a low drive power at which τ_{qp} ceases to be suppressed by the microwave-generated quasiparticles. Our measurements shown in Figure 5.16 do not show this low drive power limit where the time constant is constant with drive power. We could not reach the low power regime required to see this effect because at the lowest drive powers we measured, TLS noise significantly degraded the signal-to-noise ratio and overwhelmed the signal from the LED to the point that the software was no longer able to fit τ_{qp} from the decays. Practically, however, KIDs are usually operated within 3-6 dB of bifurcation (the 3-4 highest powers plotted for each resonator in Figure 5.16) due to their improved responsivity and noise performance at these powers. This means that although we do not measure the regime where drive power ceases to be important, we have characterized the influence of microwave-generated quasiparticles at the regime that is relevant to current KID characterization and development experiments.

5.4.2 Critical temperatures T_c

We compare our measured dependency of τ_{qp} on detector temperature T_{det} to the theoretical dependency in Kaplan et al. (1976)

$$\tau_0/\tau_r \approx \pi^{1/2}\left(\frac{2\Delta_0}{kT_c}\right)^{5/2}\left(\frac{T_{det}}{T_c}\right)^{1/2} e^{-\Delta_0/kT_{det}}, \tag{5.11}$$

Figure 5.17: *Left:* τ_{qp} is elevated for the 20 nm thick film. Plotted are τ_{qp} (T_{det}< 160 mK) corresponding to τ_{max} in Equation 5.2 (*black markers*), τ_{qp} for the lowest three P_{drive}'s for the 40 nm thick resonators, and the lowest four P_{drive}'s for the 50 nm thick resonators (*orange & red markers* for $\lambda/4$ and $\lambda/2$ resonators). Data are gathered from Figures 5.9, 5.10, 5.11, and 5.12 and 5.16. The inset includes each data point, while the main plot shows the mean and standard deviation of the points in the inset. The 40 nm thick resonators are shown as lower limits because the plateau of $\tau_{qp}(T_{det})$ was not definitively observed in Figure 5.11. The 50 nm thick resonator is shown as a lower limit because of the probable suppression of τ_{qp} due to excessively high microwave drive power (see text in Section 5.4.2 for elaboration). *Right:* T_c is elevated for the 20 nm thick film. The critical temperature T_c measured from three different possible tests described in Section 5.4.2 with fits to $T_c = A/t + T_{c,\infty}$ (Chubov et al., 1968).

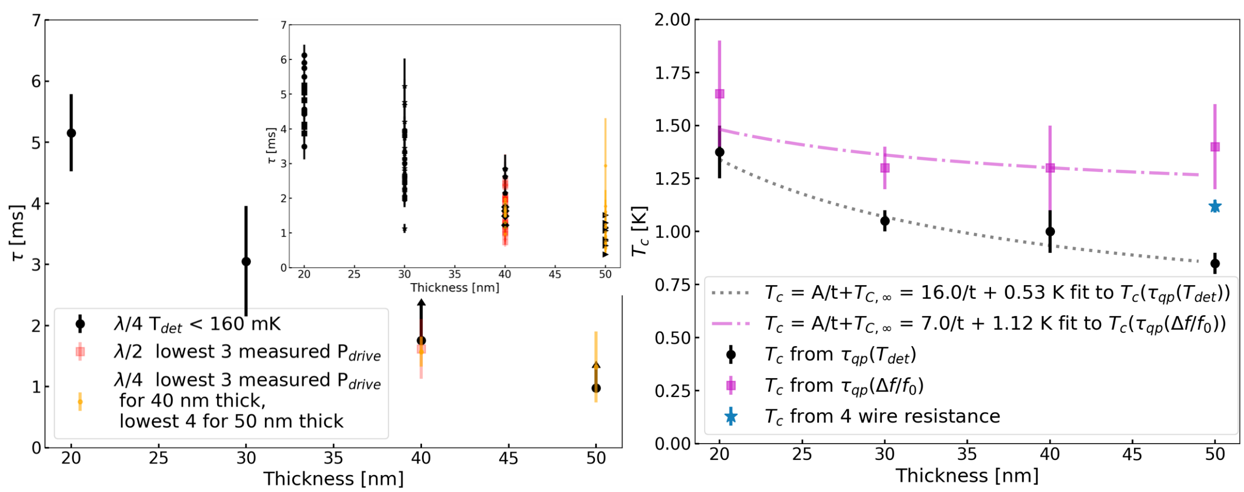

Figure 5.18: A four wire resistance measurement as a function of the temperature of the adiabatic dilution refrigerator (ADR) for the 50 nm thick array showing T_c = 1.12 K. The ADR temperature should correspond to T_{det} assuming sufficient thermal sinking.

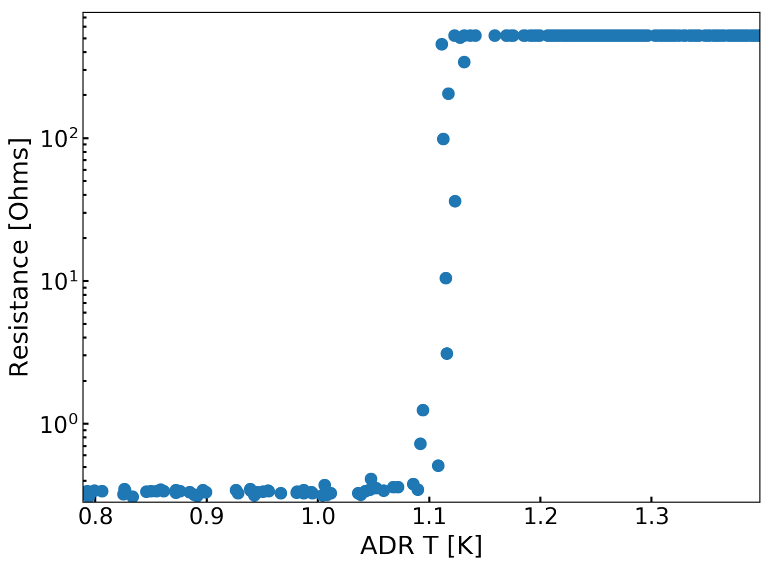

where τ_r in this reference is equivalent to our τ_{qp}. τ_0 is the characteristic decay time for the material (Kaplan et al., 1976)

$$\tau_0 = Z_1(0)\hbar/2\pi b(kT_c)^3, \tag{5.12}$$

where b is a constant that helps approximate the low-frequency interaction term between electrons and phonons, and $Z_1(0)$ is a renormalization parameter of the electron-phonon interaction. Both b and $Z_1(0)$ are properties of a given material and we use $10^3 b = 0.317$ meV $^{-2}$ and $Z_1(0) = 1.43$ from Table 1 of Kaplan et al. (1976). This equation originates from calculations of quasiparticle scattering and recombination times assuming the dominant quasiparticle energy relaxation processes are inelastic scattering with phonons and quasiparticle recombination. The recombination time is found to increase exponentially as $\exp(\Delta/kT_{det})$ at the lower temperatures because the quasiparticle population also decreases exponentially with temperature according to the Fermi-Dirac distribution.

We calculate Equation 5.11 for a range of values of T_c to compare the observed behavior of $\tau_{qp}(T_{det})$ to this theoretical expectation, plotted as dashed lines in Figures 5.9, 5.10, 5.11, and 5.12. The range of theoretical T_c's corresponding to the observed $\tau_{qp}(T_{det})$ for each thickness are shown as black circular points in the right panel of Figure 5.17.

For comparison we also measured T_c using $\Delta f/f_0$ as a function of T_{det} Zmuidzinas (2012)

$$\frac{\Delta f}{f} = 0.5\alpha\Big(\frac{\sigma_2 - max(\sigma_2)}{max(\sigma_2)}\Big), \tag{5.13}$$

Zmuidzinas (2012) where α is the kinetic inductance fraction $L_{kinetic}/L_{total}$ and σ_2 is defined using Mattis-Bardeen theory as

$$\sigma_2 = \frac{\pi\Delta_0}{hf}\Big(1 - \sqrt{\frac{2\pi k_B T_{det}}{\Delta_0}} e^{-\Delta_0/k_B T_{det}} - 2e^{-\Delta_0/k_B T_{det}} J_0(\xi)\Big); \quad \xi = \frac{hf}{2k_B T_{det}}. \tag{5.14}$$

Here, f is the microwave probe frequency and J_0 is a Bessel function of the first kind. Since $\Delta_0 = 1.76 k_B T_c$, we can use these equations to constrain both α and T_c. The 1σ bounds on T_c from these fits for each thickness are shown as magenta squares in Figure 5.17.

For the 50 nm thick array we also conducted a four-wire resistance measurement and found T_c = 1.12 K shown in Figure 5.18 and represented as a blue star in Figure 5.17.

The observed T_c from $\tau_{qp}(T_{det})$ for the 50 nm array is around 0.8 K, anomalously low for Al and much lower than the T_c from the four-wire resistance measurement (1.12 K). The discrepancy is likely caused by microwave readout power suppressing $\tau_{qp}(T_{det})$ and thereby reducing the corresponding T_c we measure. This mistake is probable; this measurement was taken at P_{drive} = 0.32 aW where $\tau_{qp} \lesssim 1$ ms, lower than the higher $\tau_{qp} \sim 1.5$ ms observed at the lowest powers for the 50 nm device with $f_0 \sim 6$ GHz.

We also fit $T_c = A/t + T_{c,\infty}$ to the data sets of T_c vs metal film thickness t according to the findings in Chubov et al. (1968), where A is a constant and $T_{c,\infty}$ is T_c for a perfect and infinitely thick metal (expected to be $\sim$ 1.2 for Al). The fits are included in the right hand panel of Figure 5.17. The increase of T_c in very thin superconducting films was postulated to be associated with electron energy quantization in the finite films and the corresponding restriction of the electron wave functions at the boundaries becoming significant (Paskin and Singh, 1965). This theory was challenged by Shapoval (1967) who instead calculated an elevated T_c arising purely from BCS superconductivity without the need for boundary and surface conditions for the electron wave functions, but without a proposed alternate mechanism for *why* the T_c increases for the thinnest films. Currently, the physics governing the behavior of T_c with film thickness is still not understood or widely agreed upon in the KID community (Zmuidzinas, 2012). Correspondingly, since τ_{qp} depends on T_c the understanding of why or how τ_{qp} changes with thickness is not understood. Regardless of the underlying mechanism, Shapoval (1967) found that $T_c = A/t + T_{c,\infty}$ and Chubov et al. (1968) found good correspondence between this formula and experiments at the time. We therefore compare our experiments to this formula as well.

These fits of $T_c(t)$ are shown in the right panel of Figure 5.17. We find $T_{c,\infty}$ from $\tau_{qp}(T_{det})$ is 0.53 $\pm$ 0.11 and $T_{c,\infty}$ from $\Delta f/f_0(T_{det})$ is 1.12 $\pm$ 0.32. The $\tau_{qp}(T_{det})$ method under-predicts T_c as compared to fits to $\Delta f/f_0(T_{det})$ and its fit $T_{c,\infty}$ is much less than theoretically expected. It is likely that measurements of T_c and thereby fits to $T_{c,\infty}$ from $\tau_{qp}(T_{det})$ are lower than expected due to 1/f suppression of the measured time constants as compared to their true value, as quantified in Table 5.1. Suppression of $\tau_{qp}(T_{det})$ also suppresses the derived value of T_c from the theoretical expectation of Kaplan et al. (1976). An example of the influence of extremely high drive power on the derived T_c is shown in Figure 5.12, which shows $\tau_{qp}(T_{det})$ for two drive powers. The higher drive power, well into the regime where τ_{qp} is suppressed according to Figure 5.16, has a T_c lower than the measurements out of this regime. We therefore assume that the "true" T_c lies between the T_c from the two methods. This assumption is corroborated by the four-wire resistance measurement for the 50 nm array that lies between the other two measured T_c's.

While the theoretical decay of T_c with thickness fits the data well for the T_c's extracted from the $\tau_{qp}(T_{det})$ measurements, the decay is less clear for the T_c extracted from $\Delta f/f_0(T_{det})$. Regardless of the overall trend of T_c(t), for both measures of T_c the 20 nm thick array has a measured T_c that is statistically elevated above the 30, 40, and 50 nm thick arrays at around $\sim$ 1.4-1.6 K. High critical temperatures of 1.4 K are also seen in the 20 nm thick Al LEKIDs developed for the STARFIRE project (Hailey-Dunsheath et al., 2018), corroborating our findings.

5.4.3 Influence of Metal Film Thickness

The black markers in the left panel of Figure 5.17 show the measured τ_{qp} for the $\lambda/4$ resonators at $T_{det}<$ 160 mK collected from Figures 5.9, 5.10, 5.11, and 5.12. These values of $\tau_{qp}(T_{det}<$ 160 mK) plotted in Figure 5.17 correspond to the pre-factor in Equation 5.2, τ_{max}. The plot shows that $\tau_{qp}(T_{det}<$ 160 mK) (== τ_{max}) appears to be elevated for the 20 nm thick film compared to the other film thicknesses measured.

The inset shows each measured value of $\tau_{qp}(T_{det}<160mK)$ collected from Figures 5.9, 5.10, 5.11, and 5.12 and the main plot shows the mean and standard deviations of these collections of values. The inset shows that the scatter between points is greater than the statistically calculated error bars – that is, that there are more systematic sources of error than are captured by just the noise. It is likely that the scatter is caused primarily by the noise affecting the fitter beyond the scope of the statistical error introduced by the noise that is calculated by Pythons curve_fit. We see evidence of this in Figure 5.4 that shows the scatter between points often exceeds the statistically created error bars. We found that increasing the level of the simulated noise increases the scatter between data points. While more simulated noise also increases the size of the error bars, it is not to the extent that it increases the scatter between data points. Therefore, the noise increases scatter above the statistical error introduced by the noise. The scatter between measurements is captured by the error bars in Figure 5.6. These error bars are comparably large to the scatter between data points in Figure 5.17 for all but the 30 nm thick array.

For the 40 nm thick film, the mean value of τ_{qp} in Figure 5.17 is considered a lower bound because we do not definitively observe the plateau of $\tau_{qp}(T_{det})$ and as such the values could be higher than observed at the lowest temperature measured. For the 50 nm thick film, it is probable that the measurement was taken at a drive power at which the microwave suppression of τ_{qp} was still a prevalent effect. Therefore, the value of τ_{qp} is considered a lower bound for the 50 nm thick film as well. We also include the measurements for the 40 nm thick un-grounded $\lambda/2$ CPW to test for diffusion of quasiparticles into the ground plane via comparison to the grounded $\lambda/4$ CPWs. The two CPW geometries appear equivalent in this plot, indicating our measurements are not affected by quasiparticle diffusion into the ground plane.

For the 40 and 50 nm thick resonators we also include measurements of τ_{qp} from the P_{drive} sweeps for the lowest values of P_{drive}, shown in orange in the left panel of Figure 5.17. These data points have a large scatter due to the influence of noise at the lower readout powers. Including these points significantly increases the mean value of τ_{qp} for the 50 nm thick array from ~ 0.9 ms to ~ 1.5 ms. This supports the argument made at the end of Section 5.4.2 that we conducted the $\tau_{qp}(T_{det})$ measurement for the 50 nm array at a P_{drive} that was high enough to suppress τ_{qp}.

As a check of the measurement made by the fits of X(t) to the decays, we also plot the averaged raw decays for the lowest temperatures from Figures 5.9, 5.10, 5.11, and 5.12 together in Figure 5.14. These correspond to the black markers in Figure 5.17. We then evaluate these decays by eye to check whether it is believable that the time constant is elevated for the 20 nm thick device. The initial portion of the data (time $\lesssim 1.5\times10^{-4}$ s) appears to decay at approximately the same rate for all thicknesses. After this time the decays diverge, with the 20 and 30 nm thick resonators showing a long and slow decay. The 50 nm thick resonator clearly completes its decay prior to any other thickness at around 2 ms. The 40 nm decays do appear to flatten out and reach ~ 0 at around 5 ms, prior to the 30 nm decays that appear to complete around 10 ms. We therefore conclude that the raw data reflect what the fit values to the data suggest, that τqp appears elevated for at least the 20 nm thick resonators compared to the arrays of 30, 40, and 50 nm thick.

5.5 Dark noise

The on and off resonance dark noise spectra in the frequency and dissipation directions S_{freq} and S_{diss} for a 50 nm thick resonator at 175 mK are shown in Figure 5.19. S_{freq} has excellent clearance over the off resonance system noise that itself appears rather flat. The readout system roll off is clearly visible in S_{freq} at $\sim 1\times10^4$ Hz. However, S_{freq} has a strong 1/f shape that dominates and blends out any obvious detector roll off.

The source of this 1/f noise is likely from two level system (TLS) noise discussed in Chapter 2, Section 2.3.3. Two signatures of TLS noise are shown to be present in our detectors in Figure 5.20 – the 1/f noise decreases at higher microwave drive powers as well as at higher temperatures.

To confirm the time constants derived from the optical decays, we measured the dark noise spectra and measure the time constant corresponding to the Lorentzian detector roll off as explained in Chapter 2, Section 2.2.7. The resonator ring time τ_{res} that is intrinsic to the LC circuit's quality factor Q_{tot} (see Equation 2.18) is a few tens of μs for our resonators. The other dominant timescale in the system is that which we want to measure: τ_{qp}. If the timescale corresponding to the Lorentzian roll off is longer than $\sim$10 μs then we can assume that the roll off is caused by the quasiparticle recombination time and not the resonator ring time.

For 50 nm we could only measure detector roll offs for the highest readout powers. For the lower readout powers there are two compounding issues that contribute to the inability to measure detector roll offs: first, the TLS noise increases for lower readout powers, and second, that the increased τ_{qp} at lower readout powers leads to a roll off at lower frequencies. In our case, if τ_{qp} = 1 ms (2 ms) we would expect a rolloff at $1/(2\pi\tau_{qp})$ = 160 Hz (80 Hz). These are well into the regime where the 1/f TLS noise dominates over any other noise contribution, inhibiting any possibility of corroborating the longer optical τ_{qp}'s with dark noise measurements.

We had the best success measuring the dark noise for the 50 nm resonators due to the high readout powers possible and the resulting shortened time constants that correspond to roll offs at higher frequencies. While some resonators for other metal thicknesses had drive powers comparable to the 50 nm resonators (e.g. the resonator with $f_0 \sim 3.9$ GHz on the 40 nm 1/4 λ array) their longer time constants remained hidden in the dominating 1/f TLS noise.

For dark noise spectra with detector roll offs present for the 50 nm thick array, we fit the roll off with single and double pole Lorentzians as well as 1/f + single pole Lorentzian to find the corresponding time constant. We find that the 1/f + single pole Lorentzian fits the data best, with an example fit shown in Figure 5.21. However, we found that the high amplitude TLS noise smoothed out the Lorentzian decay and therefore made it much more challenging to fit. Additionally, we found that the high 1/f noise increases the fit roll off frequency corresponding to a shorter τ_{qp} as compared to measurements with less 1/f influence.

For a variety of powers within 3 dB of bifurcation at 100, 150, and 175 mK for two resonators on the 50 nm thick array we achieved high enough S/N to fit the Lorentzian rolloff. We then measured τ_{qp} from corresponding LED decays for identical drive powers and temperatures. The values of τ_{qp} derived from the Lorentzian rolloff are plotted against those derived from the LED decay in Figure 5.22. Ideally, all points would lie on the $\tau_{qp,LED} = \tau_{qp,darknoise}$ line in this figure. However, the majority of points lie below this line indicating that the Lorentzian fits are measuring a shorter value for τ_{qp} than the LED decay. This could either be caused by the LED decay fits overestimating τ_{qp} (which we showed to be unlikely using the simulated data) or the dark noise Lorentzians underestimating the τ_{qp}. As stated earlier, the 1/f features of the noise tend to increase the roll off frequency fit by the Lorentzian corresponding to shorter τ_{qp}. This is consistent with the dark noise roll offs appearing to underestimate τ_{qp} with respect to the optical decays.

Figure 5.19: Frequency *black* and dissipation *orange* noise for on *solid lines* and off resonance *transparent, dashed lines* measurements for a 50 nm thick resonator at 175 mK. The on resonance noise shows clearance above the off resonance (system) noise, but no detector rolloff is seen in the spectrum.

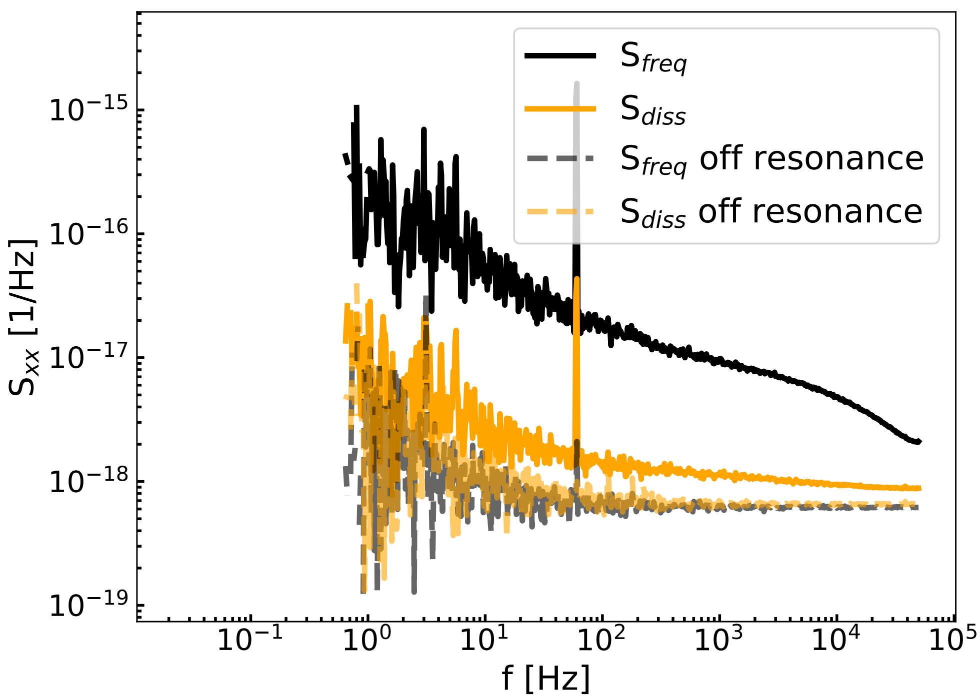

Figure 5.20: Evidence for TLS noise in a 50 nm thick resonator. **Top**: frequency (*darker lines*) and dissipation (*lighter lines*) noise spectra at microwave drive powers of -86 dBm (*purple*) and -89 dBm (*teal*). The higher microwave drive power results in lower noise levels in both S_{freq} and S_{diss}, as evidenced by the decrease by a factor of $\sim$2 times between the pair of purple (-86 dBm) and pair of teal lines (-83 dBm). **Bottom**: frequency and dissipation dark noise spectra at a microwave drive power of -86 dBm at T_{det} = 100 mK (*purple*) and 200 mK (*teal*). The higher detector temperature corresponds to a suppression of 1/f noise.

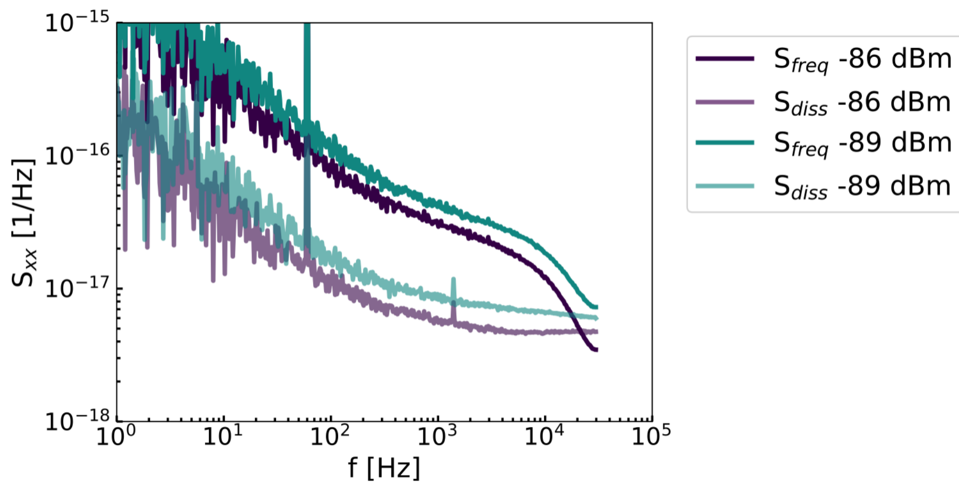

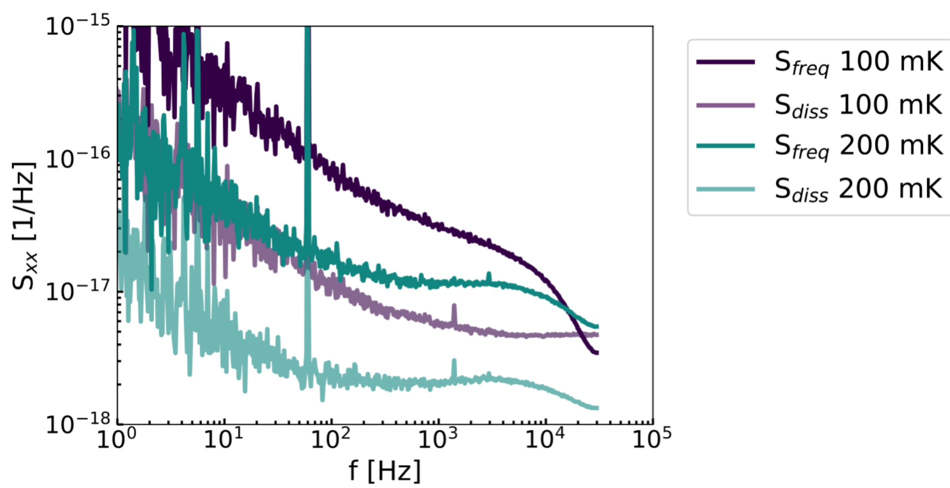

Figure 5.21: *Blue*: On resonance dark noise of a resonator at f_0 = 6.06 GHz at 100 mK. *Pink dashed line* shows the 1/f + Lorentz fit to the detector time constant-induced roll off at around 1000 Hz. The two components of this fit are shown separately as a *cyan dashed line* (the Lorentz component) and a *grey dashed line* (the 1/f component).

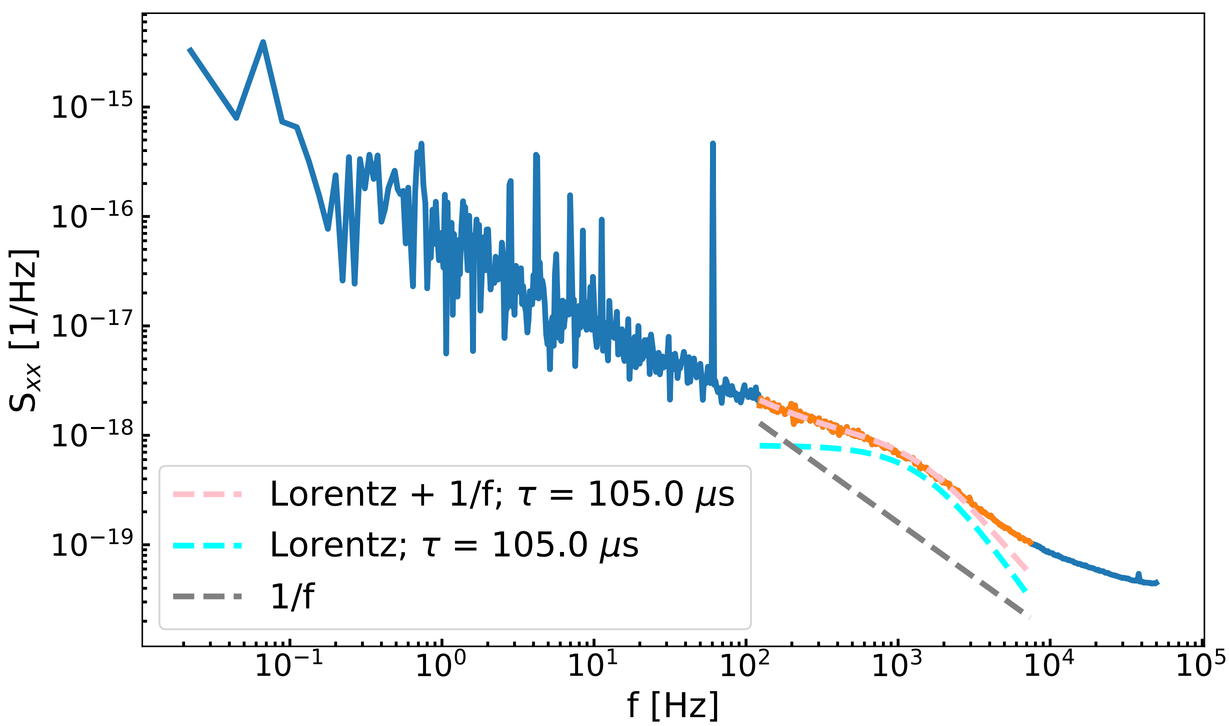

Figure 5.22: *Black points*: τ_{qp} derived from the Lorentzian fits to dark noise roll offs (like that in Figure 5.21) plotted against the τ_{qp} from a fit of X(t) to a LED-induced decay of Δf/f_0. *Grey dashed line*: the line along which $\tau_{qp,Lor} = \tau_{qp,LED}$. There is one outlier from the general group. Upon inspection of the data, no obvious difference is present between this data point and other data points included in this plot.

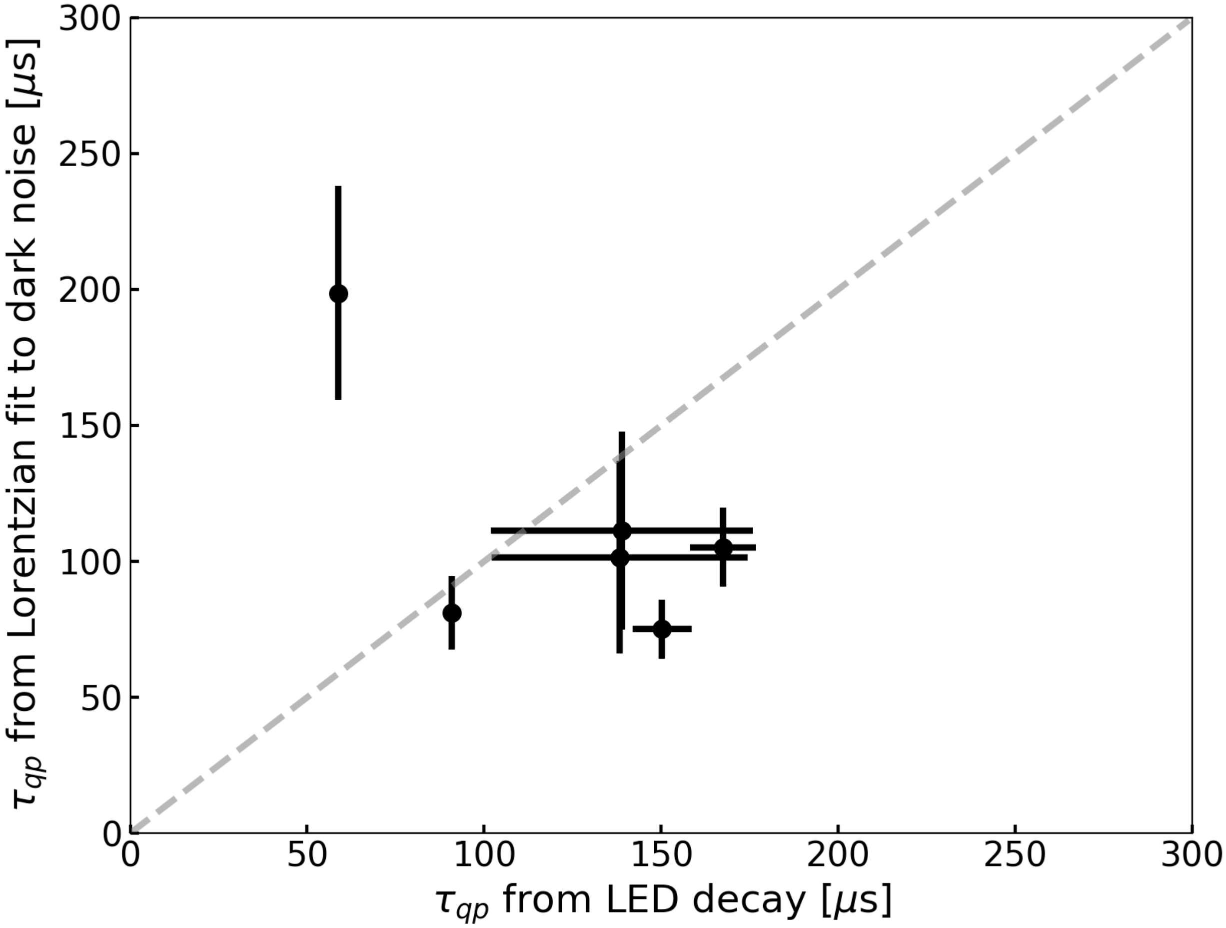

5.6 Discussion and Summary

We measured τ_{qp} for Al CPWs with metal film thicknesses of 20, 30, 40, and 50 nm by fitting Equation 5.9 to optical decays of $\Delta f/f_0$ resulting from flashes of infrared LED light. Using simulated data we determined the most accurate fitting technique that still under-predicts τ_{qp} due to 1/f contamination in the decays (Table 5.1 and Figure 5.8). The measured $\tau_{qp}(T_{det})$ matches the theoretical decay from Kaplan et al. (1976) (Figures 5.9, 5.10, 5.11, and 5.12). We find that τ_{qp} is suppressed by microwave drive power (Figure 5.16), consistent with the findings of de Visser et al. (2014). We find that τ_{qp} and T_c are elevated for the thinnest metal film of 20 nm (Figure 5.17), with the elevation of T_c for the thinnest Al films consistent with other works (Hailey-Dunsheath et al., 2018).

We found evidence that TLS noise dominates the noise spectra for these devices (Figure 5.20). We fit the Lorentzian roll offs to the highest driving powers of the 50 nm thick resonators (Figure 5.21) and compared them to the τ_{qp} measurements from the X(t) fits (Figure 5.22) and find decent correlation between the two.

The elevated τ_{qp} for the 20 nm thick device is unlikely to be caused by quasiparticle diffusion into the ground plane. We measured τ_{qp} for differing center conductor lengths, as well as for $\lambda/2$ grounded compared to $\lambda/4$ ungrounded 40 nm thick resonators, and saw no statistical difference. Therefore, we propose that τ_{qp} changes with thickness due to T_c's dependence on thickness. We measure an elevated T_c for the thinnest (20 nm) film and a correspondingly long τ_{qp} (3.8-5.8 ms). Evidence that the long τ_{qp} of 2.8-4.5 ms measured for the 30 nm thick film also corresponds to elevated T_c in the 30 nm thick film is tenuous (Figure 5.17 right panel). A four-wire resistance measurement of the 20, 30, and 40 nm thick arrays is needed to corroborate our findings.

Chapter 6

NGC6240

6.1 Abstract

NGC 6240 is a luminous infrared galaxy in the local Universe in the midst of a major merger. The merger has triggered star formation and AGN activity in each of the progenitor galaxies' nuclei. In the nuclear region of the merger the two nuclei appear separated by around 1 kpc, each surrounded by concentrations of dust and disks of stars. The molecular gas peaks in emission between the two nuclei as opposed to centered around the nuclei with the stars and gas.

We analyze high-resolution observations of CO J = $3-2$ and $6-5$ of the central few kpc of NGC 6240 taken at the Atacama Large Millimeter Array. Using these CO line observations, we model the density distribution and kinematics of the molecular gas located between the nuclei of the progenitor galaxies. Our models suggest the majority of this gas is a tidal bridge linking the two nuclei that could fall onto the nuclei prior to second pass and feed future starbursts. We also observe high velocity gas ($>$ 300 km/s) that is not captured by the model that could be accelerated by either gravitational forces from the merger or an AGN outflow. These findings shed light onto small-scale processes that can affect galaxy evolution and the corresponding star formation, with the tidal bridge depositing molecular gas onto the nuclei while other energetic forces accelerate molecular gas further out of the nuclear region.

6.2 Introduction

Ultra-luminous infrared galaxies (ULIRGs) are galaxies with infrared luminosities L_{IR} between 10^{12} and 10^{13} $L_{\odot}$, with IR contributions from dust heated by star formation and, often, active galactic nuclei (AGN) (Lonsdale et al., 2006). In ULIRGs, these processes appear to be triggered by major galaxy mergers and galaxy interactions the majority of the time, though there are exceptions. Their less luminous counterparts luminous infrared galaxies (LIRGs; L_{IR}= $10^{11}-10^{12}$ $L_{\odot}$) are not dominated by major mergers, though the most luminous LIRGs are primarily comprised of major mergers (Alonso-Herrero, 2013).

AGN are found to be more common in (U)LIRGs of higher IR luminosities (see Alonso-Herrero (2013) for a review). Earlier stages of galaxy interactions and mergers are more likely to classify as "composite", with IR luminosities sourced from comparable contributions from both star formation and AGN (Yuan et al., 2010). This means that (U)LIRGs of high IR luminosities in the early stages of a major merger are especially fruitful objects for studying the interaction between star formation and AGN activity.

ULIRGs are more numerous at redshifts $z>1$ than they are at z of 0 by orders of magnitude, indicating they represent one important piece of galaxy evolution. (U)LIRGs were discovered to be dominant contributors to the cosmic infrared background, further signifying their importance. Despite their importance, LIRGs are rare objects in the local Universe (Giovannoli et al., 2011). This means that any (U)LIRGs extant in the local Universe are important to study in detail, with the capability to study them in high resolution allowing an understanding of the detailed mechanics of this important galaxy class.

NGC 6240 (Wright et al. (1984); Thronson et al. (1990)) is a unique LIRG in the local universe ($z = 0.02448$) in the midst of a major merger event (Fried and Schulz, 1983) that is triggering high star formation rates (Genzel et al. (1998); Tecza et al. (2000)) and AGN activity in its two nuclei (Vignati et al., 1999). Its two AGN contribute ~20-24% of the bolometric luminosity (Armus et al., 2006), a rarity for local LIRGs for which only only $\sim$ 8% have an AGN bolometric contribution more than 25% of the infrared luminosity (Alonso-Herrero et al., 2011). It has been studied in detail using wavelengths ranging from the radio to X-ray, revealing it to be a complex and dynamic laboratory to inform the theory of galaxy evolution. Its far-infrared (FIR) luminosity $L_{FIR} \approx 10^{11.8}$ $L_{\odot}$ (Sanders et al. (1988); Thronson et al. (1990); Sanders and Mirabel (1996)) is just below that needed to classify it as a ULIRG, though it is expected to cross this threshold when a second starburst is triggered during final coalescence (Engel et al., 2010). As such, NGC 6240 presents an excellent opportunity to study in fine detail an example of a merger-driven ULIRG just before it passes this threshold.

The nuclear region ($<$1 kpc) of LIRGs and ULIRGs is a critical location to study as it is often the location of the highest star formation rates and the origin of many feedback processes, such as stellar winds and AGN outflows. Correspondingly, in the most luminous LIRGs ($L_{IR} > 6\times10^{11}$ $L_{\odot}$) the majority of the mid-IR emission originates from this central kpc region (Alonso-Herrero, 2013). The dynamics in this central region are complicated by the galaxy merger interactions that trigger the processes that lead to the high IR luminosity of most LIRGs (Lonsdale et al., 2006). The nuclear molecular gas dynamics of NGC 6240 are no exception, with a concentration of molecular gas *between* the two nuclei while the stars and dust are found to be concentrated *around* the two nuclei (Tacconi et al. (1999); Tecza et al. (2000); Engel et al. (2010) and others).

The molecular gas dynamics of NGC 6240's nuclear region have been studied and modeled for decades without consensus. Tacconi et al. (1999) observed a velocity gradient of CO $J = 2-1$ in the nuclear region and modeled it as a gravitationally stable, rotating disk. A similar disk model was used to describe the motions of HCN by Scoville et al. (2015). This inter-nuclear disk model has both been used to support other author's observations of molecular gas (e.g. Iono et al. (2007) in CO $J = 3-2$ and HCO^{+} $(4-3)$) and been claimed to be unphysical in the context of other observations (e.g. Gerssen et al. (2004) in $H\alpha$+[NII]). Alternate geometries have been proposed to explain this central molecular gas, including a tidal bridge connecting the two nuclei (Engel et al., 2010) and the origin site for a warm molecular outflow (Cicone et al., 2018).

With the lack of consensus on the geometry of the inter-nuclear molecular gas, updated detailed modeling of the central nuclear region using multiple new line transition observations in high resolution is needed. New telescopes have allowed the nuclear region of NGC 6240 to be observed in unprecedented detail and enable this detailed modeling. In this paper we use high resolution CO J = 6 − 5 and CO J = 3 − 2 observations from the Atacama Large Millimeter Array (ALMA) to model the velocity profile and density distribution of the central molecular gas. We find that while a disk-like concentration of gas can explain the majority of the line emission it cannot capture the kinematics of this central region. Our fiducial model finds the central molecular gas is flat and pancake-like, suggesting that it is not a self-gravitating, rotating disk but instead a transient tidal bridge connecting the two nuclei.

The residuals from the fiducial model definitively show high velocity gas that is not associated with the nuclear pancake-like concentration of gas. This high-velocity gas is potentially associated with a multi-phase AGN driven outflow studied by other groups (Cicone et al. (2018); Müller-Sánchez et al. (2018)), or could be gravitationally accelerated by the merging process. The presented observations also reveal extended emission features, some of which correspond to previously studied filamentary structures, and others of which are new to these observations.

Section 6.3 presents the ALMA observations of CO J = 3 − 2 and J = 6 − 5 and the continuum observations at 344 and 677 GHz. We analyze these observations in Section 6.4 and calculate the mass of the dust from the continuum emission, highlight observed extended molecular gas emission features, and explore the velocity structure of the gas. Sections 6.4.5 and 6.4.6 outline a non-local thermodynamic equilibrium disk model used to fit the nuclear molecular gas component as well as the fitting strategy for the model. The fiducial model is discussed in Section 6.4.7, finding that the molecular gas between the two nuclei is unlikely to be a self-gravitating disk. In Section 6.5 it is discussed whether the central molecular concentration could instead be a tidal bridge in light of the model findings. We explore possible sources of observed highly redshifted gas in Section 6.6. The anatomy of the merger and comparisons to simulations are discussed in Section 6.7. Finally, we conclude and summarize our findings and extend our findings to general galaxy evolution research in Section 6.8.

6.3 Observations

Observations of CO J = $3-2$ and J = $6-5$ were completed using ALMA in the extended configuration. These are the first observations of the nuclear region of NGC 6240 in CO J = $6-5$ and the highest resolution observations to date in CO J = $3-2$. The CO J = $3-2$ observations were completed during Cycle 2 for project number 2013.1.00813.S. The CO J = $6-5$ observation was completed for project number 2015.1.00658.S during Cycle 3. The longest baselines, full-width half-maxima (FWHM) of the beams, channel widths, reported channel root mean squared (RMS), central frequency of continuum observations, bandwidth (BW) of continuum observations, and RMS of continuum observations are reported in Table 6.1. For the integrated moment 0 maps, $\sigma_{integrated}$ is calculated using line channel RMS values $\sigma_{line\ channel}$ from Table 6.1. The RMS in the moment map is $\sigma_{integrated} = \sqrt{n_{chan}\sigma_{line\ channel}}$ where n_{chan} is the number of channels included in the integration.

To check the reliability of the provided reduced data products, we re-imaged the data using the Common Astronomy Software Applications (CASA) software package and the National Radio Astronomy Observatory (NRAO) provided imaging script. For CO J = $3-2$ no new structure emerged in the molecular gas when re-imaging the data, nor could we improve upon the noise. Therefore, we deemed the NRAO provided data products sufficient for analyses of the CO J = $3-2$ molecular gas.

Re-imaging was required for the CO J = $6-5$ and both continuum observations due to contamination in the continuum maps from extremely high-redshift ($\sim 400-740$ km/s) molecular gas. This CO line contamination created a false continuum source between the two nuclei with the same peak flux density as the southern nucleus. Upon re-imaging, this source disappeared. We use Briggs weighting for the re-imaged maps to match the weighting schemes used for the NRAO-provided molecular gas data products. For all images, $1''$ corresponds to 500 pc of projected distance and AGN locations from Hagiwara et al. (2011) are included as crosses. Below, we introduce all observations with further details of all figures discussed in the following sections.

Moments 0, 1, and 2 for CO J = $3-2$ are plotted in Figure 6.1 with the 344.4 GHz continuum contours included. Similarly, Figure 6.2 shows the moments of CO J = $6-5$ with the 677 GHz continuum contours. Observations in both J = $3-2$ and J = $6-5$ show a concentration of molecular gas between the two nuclei, as observed previously in other observation of the molecular gas (e.g. Tacconi et al. (1999), Scoville et al. (2015)). Also similar to previous observations, the CO has a velocity gradient from highly redshifted ($<v> \sim$ 350 to 400 km/s) to blueshifted ($<v> \sim$ -100 to -150 km/s) along a position angle of approximately 34°. Extended structure including two dim features of molecular gas to the SE and SW of the southern nucleus are observed in the CO J = $3-2$ maps that have also been observed in other wavelength bands, such as H_2 (Max et al., 2005), Fe XXV, and Hα (Wang et al., 2014). We also observe new extended emission to the north in the CO J = $3-2$ observation.

Channel maps for the CO J = $3-2$ and CO J = $6-5$ observations are shown in Figure 6.3 and Figure 6.4, respectively. The velocity structure of the nuclear region is well resolved in the higher sensitivity CO J = $3-2$ channel map and shows emission at 5σ above the noise at velocities greatly exceeding the average velocities in Figure 6.1: highly blueshifted emission (-415 km/s) near the southern nucleus and highly redshifted emission (665 km/s) located between the two nuclei. The CO J = $6-5$ channel map shows a similar velocity structure to that of the CO J = $3-2$, with a gradient of velocities between -358 and 723 km/s in the molecular gas between the two nuclei. This velocity structure is best resolved in the nuclear region between the two AGN, where the gas emission is brightest.

The continuum maps for all observations are plotted in Figure 6.5, showing a concentration of continuum emission around each of the nuclei. The concentration around the southern nucleus is much brighter than the northern concentration, with flux densities reported in Table 6.2. The location of the peak flux density is farther south in the 344 GHz observation than the 677 GHz observation by approximately 0.4″. This difference indicates a likely temperature or optical depth gradient along the axis between the two nuclei. The distribution of continuum emission, along with the offset in peak flux density locations, is also observed at comparable frequencies, beam sizes, and sensitivities in Scoville et al. (2015).

Targeted line	Configuration (Baseline)	Beam FWHM	Beam PA	Line channel width
CO J = $3-2$	Extended (1.6 km)	0.34 x 0.15"	-67^o	20 km/s
CO J = $6-5$	Extended (460 m)	0.32 x 0.23"	-83^o	20 km/s
	Line channel RMS	Continuum ν	Continuum BW	Continuum RMS
	2.52 mJy/beam	344.4 GHz	7.5 GHz	0.21 mJy/beam
	14.5 mJy/beam	677 GHz	6 GHz	2.1 mJy/beam

Table 6.1: Parameters of observations received from ALMA.

6.4 Data Analysis and a Test Model

6.4.1 Mass from Continuum

To calculate the mass of the dust M_d we use

$$M_d = \frac{S_\nu D_L^2}{\kappa_\nu B_\nu(T)} \text{ kg} \tag{6.1}$$

(Casey, 2012), where S_ν is the flux density in the continuum frequency band, D_L is the luminosity distance of 108 Mpc, κ_ν is the dust mass opacity coefficient, and $B_\nu(T)$ is the blackbody emission in the continuum frequency band for a dust temperature T. For this calculation, we use the 344 GHz (870 μm) as the measure of S_ν because it should have a lower optical depth than the 677 GHz continuum observation. The total S_{344GHz} for this observation is 27 mJy, 18% of the galaxy-integrated flux density measured at 850 μm by SCUBA of 150 mJy (Klaas et al., 2001). This flux density recovery is comparable to Scoville et al. (2015) who measured 18-24 mJy at 340 GHz for a comparable beam size and sensitivity with ALMA in Cycle 0. The SCUBA beam is as large as our entire ALMA observation with a diameter of 15$''$ (Klaas et al., 2001), and as such we can expect that their measurement includes extended emission that is lost in our high-resolution observations. Therefore we expect our value of S_ν to be lower than that measured with SCUBA.

For $B_\nu(T)$ we choose a dust temperature of 56 K, the dust temperature fit in Kamenetzky et al. (2014) who used a greybody fit to *Herschel*-SPIRE, IRAS, *Planck*, SCUBA, and ISO pho-

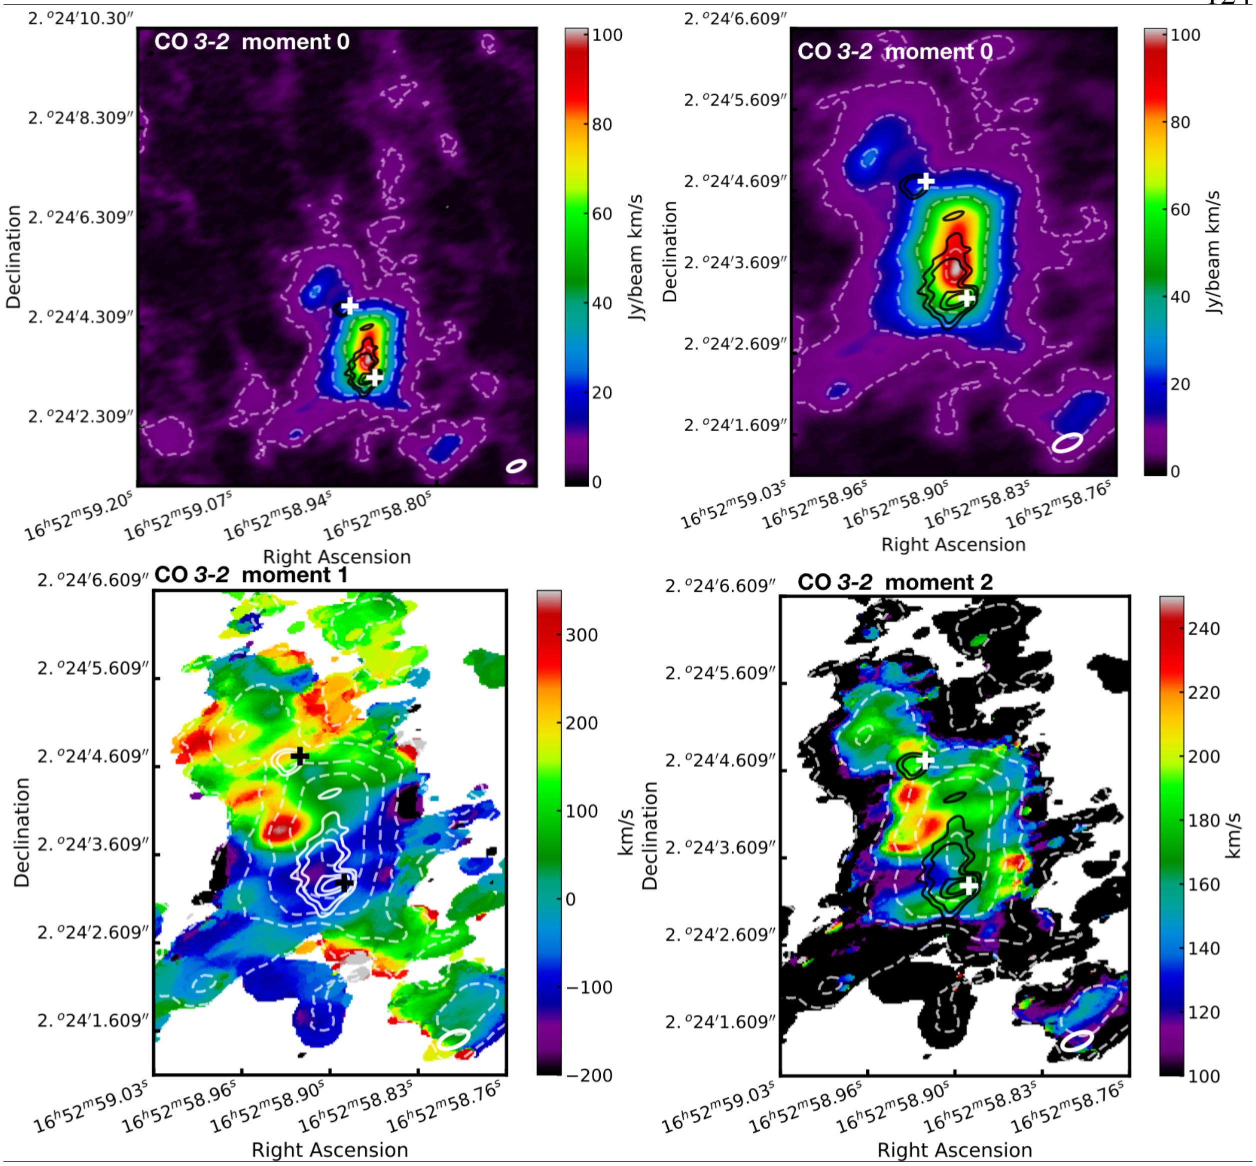

Figure 6.1: **Upper Left:** Full map of moment 0 of CO J = 3 − 2 with *dashed white contours* corresponding to 10, 25, 50, 100, 200, 300, 400, 500, and 600 $\sigma_{integrated}$ of the integrated line emission, *solid contours* show the 344.4 GHz continuum at 4, 6, 10, and 14 σ, and *crosses* show the location of the two AGN (Hagiwara et al., 2011). **Upper Right:** The same, but zoomed to show only the nuclear region. **Lower Left:** Moment 1 in the nuclear region with channels below 5 σ masked out prior to moment calculation. **Lower Right:** Moment 2 in the nuclear region with channels below 5 σ masked out prior to moment calculation. The beam FWHM contour is plotted in the lower right of each image, 1″ corresponds to 500 pc, east is left, and north is up. Velocities are calculated relative to the average velocity of CO J = 3 − 2 for the entire observation of 7200 km/s.

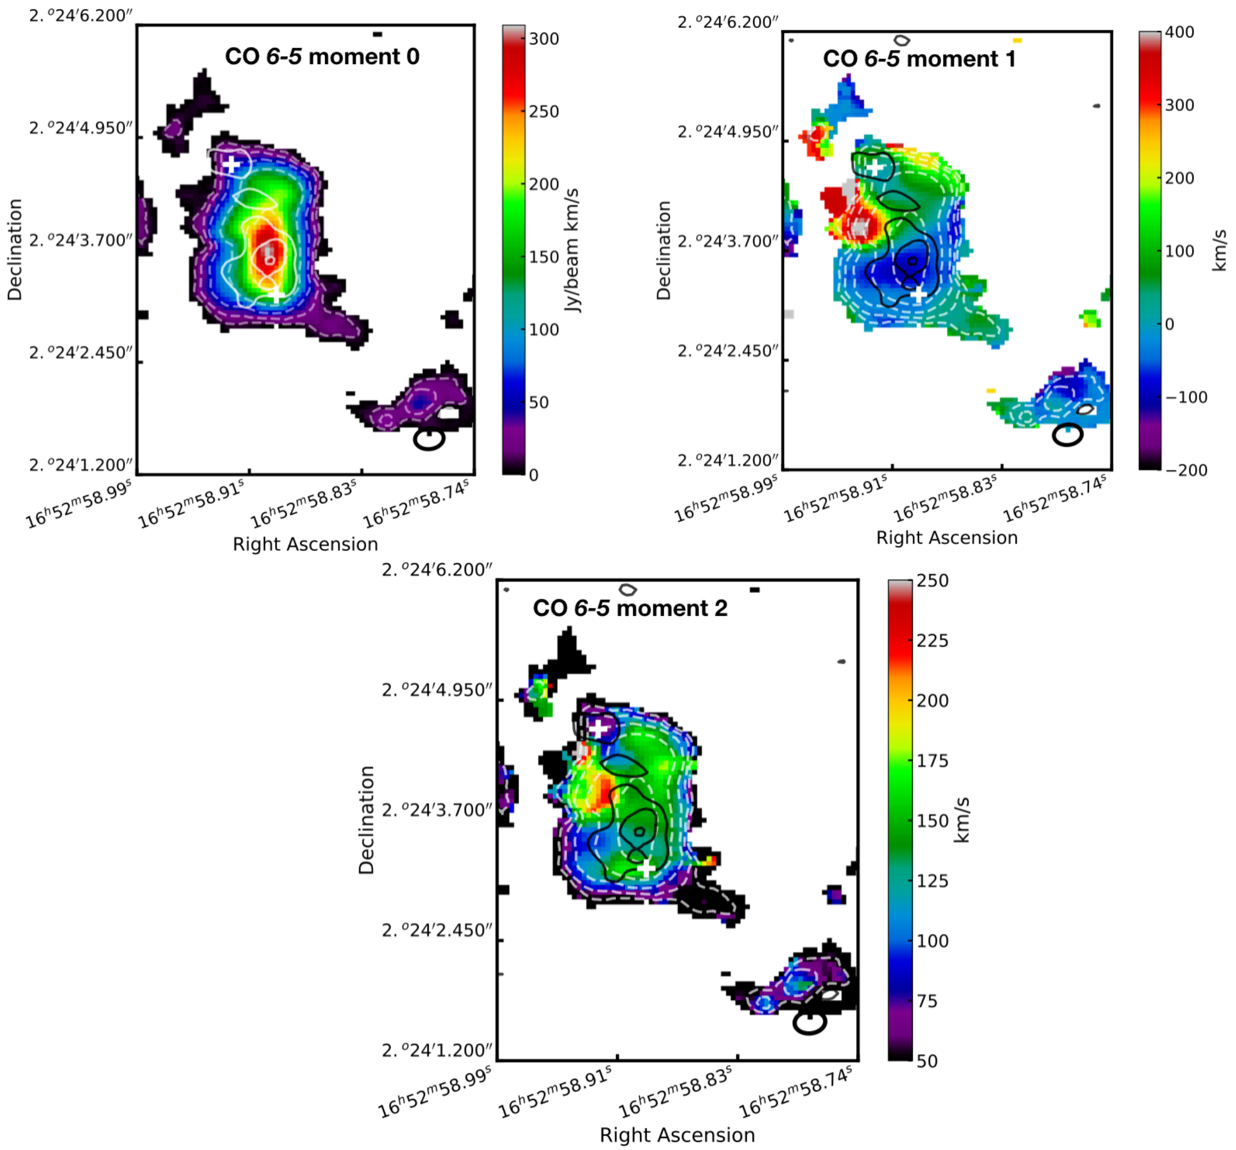

Figure 6.2: **Upper Left:** Moment 0 of CO J = 6 – 5, with *dashed contours* corresponding to 10, 25, 50, 100, 200 $\sigma_{integrated}$ of the integrated line emission and *solid contours* show the 677 GHz continuum at 5, 10, and 15 σ. Channels below 5σ are masked out of the integration. **Upper right:** Moment 1 of CO J = 6 – 5, with pixels below 5 σ masked out prior to moment calculation. **Lower panel:** Moment 2 of CO J = 6 – 5, with pixels below 5 σ masked out prior to moment calculation. The beam FWHM contour is plotted in the lower right of each image. Channels below 5σ are masked out of the integration for all maps, 1″ corresponds to 500 pc, east is left, and north is up. Velocities are calculated relative to the average velocity of CO J = 6 – 5 for the entire observation of 7465 km/s.

Figure 6.3: The channel map of primary beam corrected, continuum subtracted CO J = 3 − 2 observations show every third line channel (each panel is separated by 60 km/s). *White contours:* 10, 25, 50, 75, and 100 times the RMS of the CO J = 3 − 2 observations. *White crosses:* the locations of the AGN. The beam FWHM contour is plotted in the bottom right of each figure, 1″ corresponds to 500 pc, east is left, and north is up. Velocities are calculated relative to the average velocity of CO J = 3 − 2 for the entire observation of 7200 km/s.

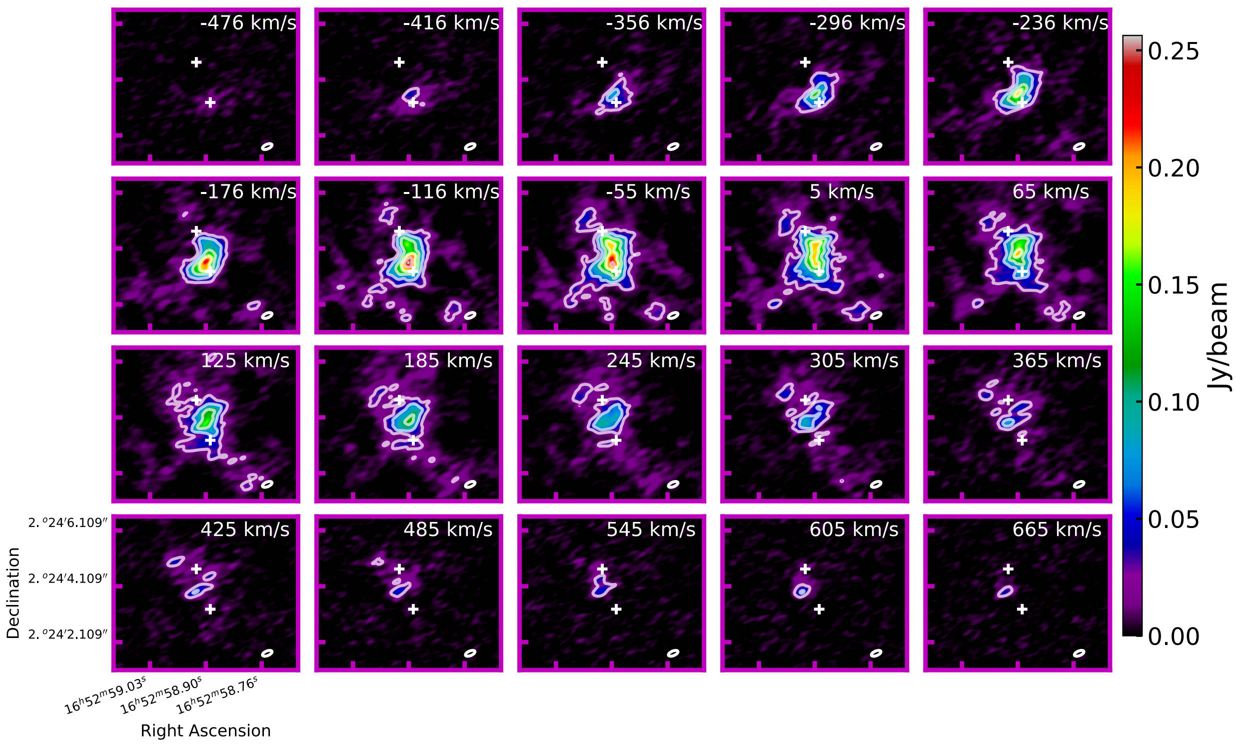

tometry to find the dust temperature and mass. This dust temperature is a galaxy-averaged property since observations used in Kamenetzky et al. (2014) have beam FWHM at least 100 times larger than the ALMA observations. It is likely that the dust temperature in NGC 6240's nuclear region is higher than the galaxy-averaged temperature due to the influences of concentrated star formation and AGN luminosity. However, we do not have flux density measurements across sufficient wavelengths in this central region to independently calculate the dust temperature. From James et al. (2002) we use $\kappa_{850} = 0.07$ m^{2}kg^{-1} and $\kappa_\nu \propto \nu^2$ to find $\kappa_{870} = 0.067$ m^{2}kg^{-1}.

The calculated masses for the regions outlined in Figure 6.6 are tabulated in Table 6.2. The total mass of the dust is calculated to be 1.2×10^7 $M_\odot$. As a check on this total dust mass we

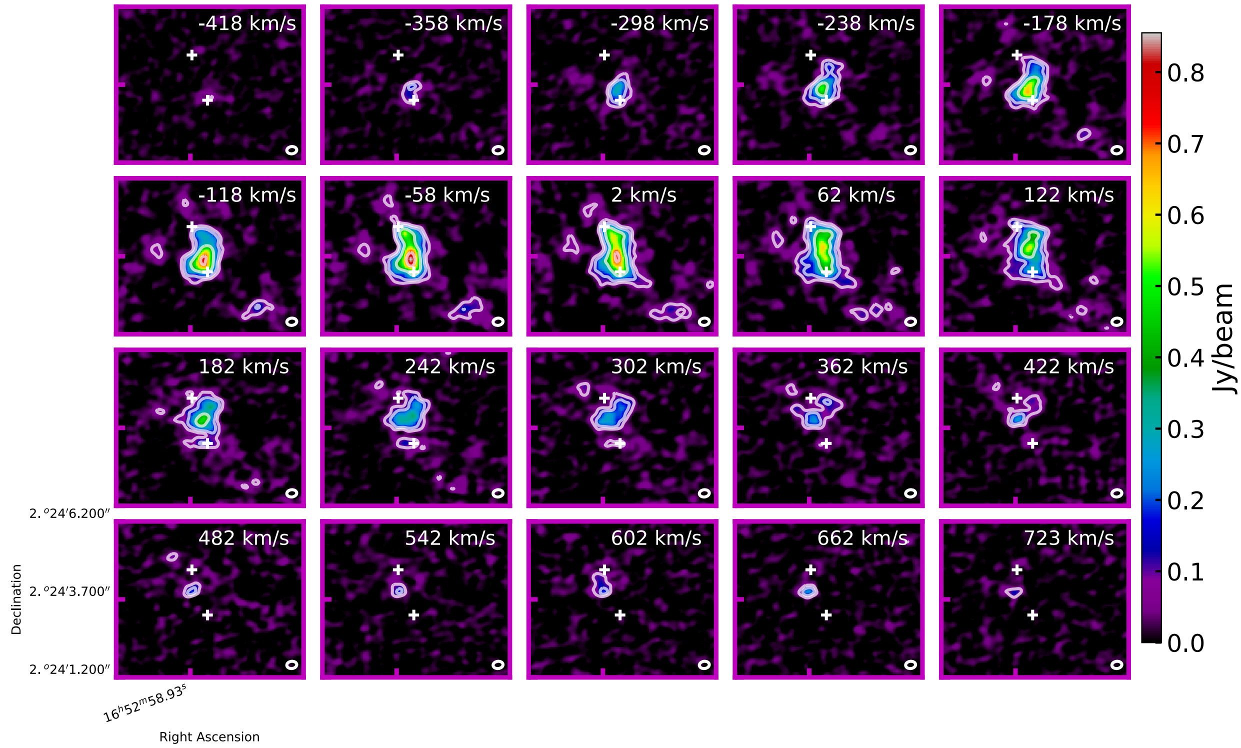

Figure 6.4: The channel map of primary-beam-corrected, continuum-subtracted CO J = 6 − 5 showing every third line channel. *White contours:* 5, 10, 25, 50, and 75 times the RMS of the CO J = 6 − 5 observations. *White crosses:* The locations of the AGN. The beam FWHM contour is plotted in the bottom right of each figure, 1″ corresponds to 500 pc, east is left, and north is up. Velocities are calculated relative to the average velocity of CO J = 6 − 5 for the entire observation of 7465 km/s.

compare to the galaxy-integrated dust mass of 5×10^7 $M_\odot$ found in Kamenetzky et al. (2014) from their model described earlier in this section. Our calculated dust mass is 24% of this value, consistent with the relative integrated flux density we calculate for the 344 GHz observations as compared to the SCUBA observations of 18%.

Using a gas-to-dust ratio of 100 we can convert the dust masses from Table 6.2 to a molecular gas mass. This is converted to a column density N_{H_2} by dividing the gas mass by the mass of molecular hydrogen and the projected size of the region on the sky. The average column density for the entire nuclear region (region 1 in Figure 6.6) is 4×10^{22} cm^{-2}, while the concentrations around the two nuclei both have slightly higher column densities of 2×10^{23} cm^{-2}. This value is

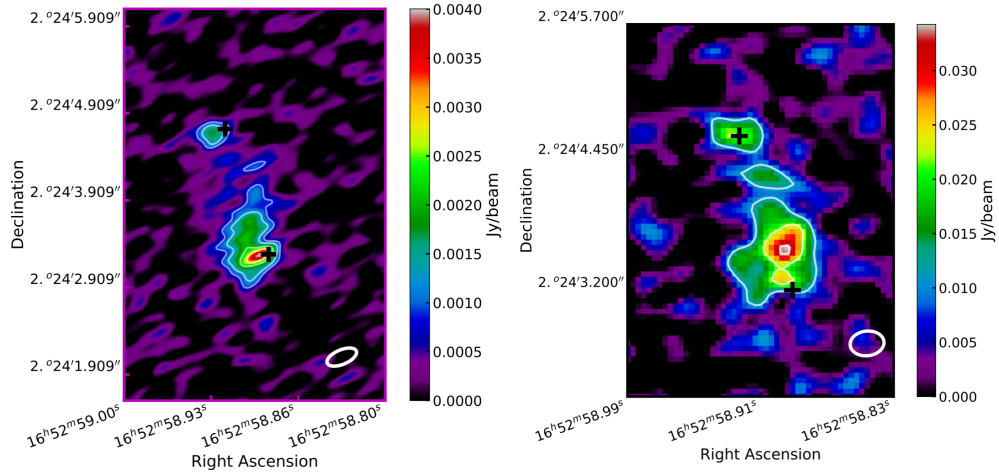

Figure 6.5: **Left:** *Colors*: The 344.4 GHz continuum observation. *Solid White Contours*: 4, 6, 10, and 14 RMS level contours of the extended configuration observation of the continuum at 344.4 GHz. *Black Crosses*: The locations of the two known AGN. The beam FWHM contour is plotted in the lower right of the image. **Right:** *Colors*: The continuum observation at 677 GHz. *Black Contours*: 5, 10, and 15 RMS level contours of the 677 GHz continuum. *Black Crosses*: The locations of the two known AGN. The beam FWHM contour is plotted in the lower right of the image, 1″ corresponds to 500 pc, east is left, and north is up.

consistent with the findings of Tacconi et al. (1999) of $N(H_2) \sim 1$–2×10^{23} cm^{-2}.

Region	677 GHz S_ν [Jy]	344.4 GHz S_ν [Jy]	M_d [$M_\odot$]	M_{gas}[M$_\odot$]	N_{H_2} [cm^{-2}]
1	0.27	0.027	1.2×10^7	1.2×10^9	4.0×10^{22}
2	0.03	0.002	9.0×10^5	9.0×10^7	2.1×10^{23}
3	0.116	0.013	5.8×10^6	5.8×10^8	2.4×10^{23}

Table 6.2: Integrated continuum flux densities S_ν, derived dust masses from the 344.4 GHz continuum calculated using Equation 6.1, gas masses assuming a gas-to-dust ratio of 100, and associated column densities of the regions outlined in Figure 6.6

6.4.2 Extended Molecular Gas Emission

A dim trail of gas extending directly north of the central concentration, which we call the "Northern Finger", is observed in the higher-sensitivity CO J = 3 − 2 observation but not in CO J = 6 − 5 or either continuum observation. It is not spatially coincident with observed H_2 V = 1 − 0

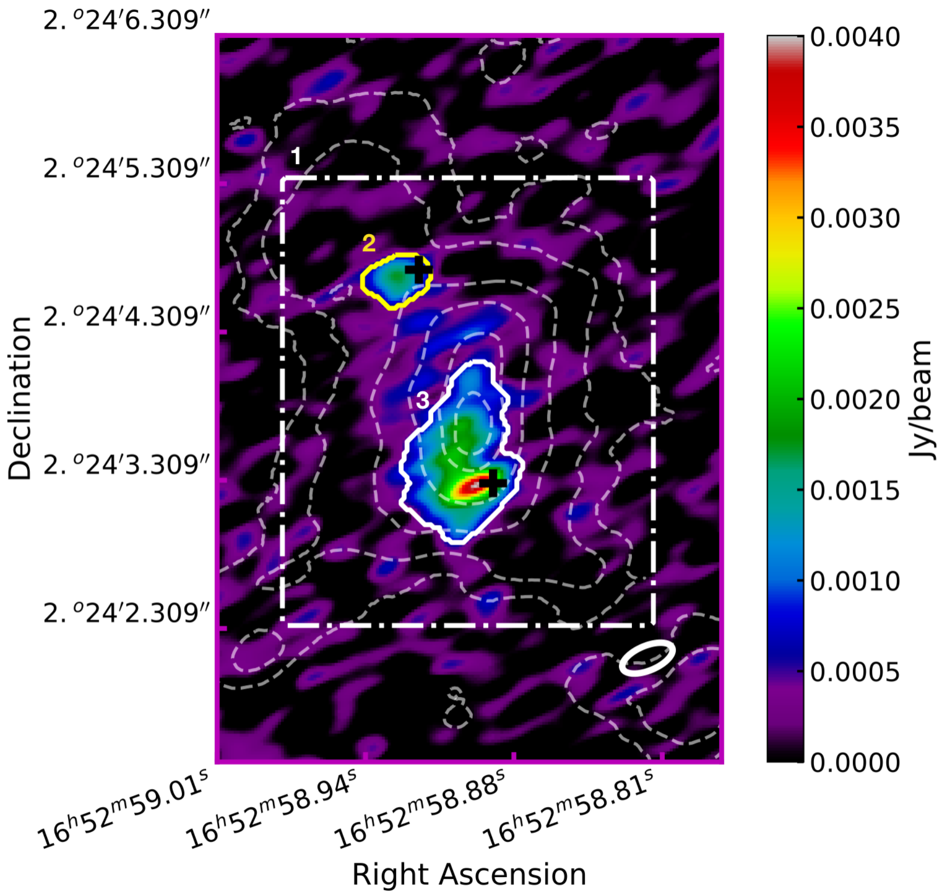

Figure 6.6: Regions for which we calculate the dust mass. *Color map:* Extended configuration 344 GHz continuum map. *Dashed contours*: CO J = 3 − 2 integrated line emission. Masses for these regions are tabulated in Table 6.2. The beam FWHM contour for the 344 GHz continuum is plotted in the bottom right.

S(1) and S(5) emission (Max et al. (2005), their Figure 12a). It dimly appears in their Figure 2a and 2c, false color images including Keck K' band, H band, J band, F814W filter from WFPC2 on *Hubble Space Telescope* (HST), and the F450W filter from WFPC2 on *HST*. Its southern half is spatially coincident with [O III] $\lambda = 5007$ and Hα emission (Müller-Sánchez et al. (2018), their Figure 1).

The Finger's average velocity is 44 km/s with a FWHM of $\sim$ 300 km/s. The low average velocity suggests this gas is unlikely to be an outflow. A nuclear inflow would cause σ_v to be elevated near the nuclei, for which there is no conclusive evidence. The low velocity is consistent

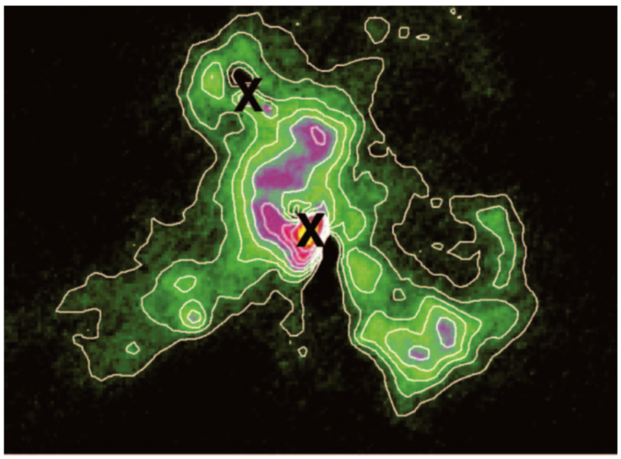

Max, Fig. 10 : integrated emission of H_2 1-0 S(1)

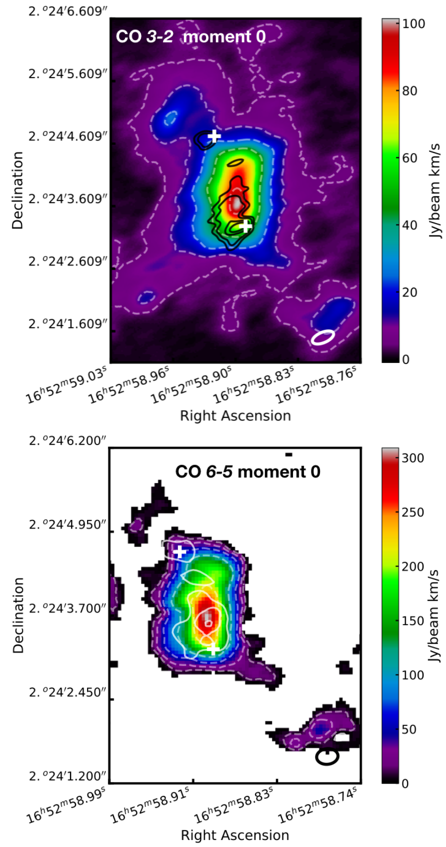

Figure 6.7: *Left*: integrated emission of H_2 $1-0$ S(1) from Max et al. (2005) showing a ribbon of hot H_2 with a reverse S shape and a peak of emission around the southern nucleus (nuclei marked with x's) *Right*: the integrated emission of the observations presented in this work showing a lack of correspondence between the their observations and ours. This difference is noted in their paper when comparing to the CO observations in Tacconi et al. (1999).

with a tidal tail or bridge but the high FWHM does not conform to this idea. The finger extends from the northern nucleus towards the northern dust lane of the galaxy, which is interpreted as a tidal tail by other authors (Gerssen et al. (2004); Yun and Hibbard (2001)). The top edge of the finger is spatially coincident with the eastern edge of the dust lane, suggesting a possible correspondence. The dust lane is much wider and extends much farther beyond the extent of the finger – approximately 7$''$ wide (roughly East/West) and over 20$''$ long (roughly North/South), beyond the extent of our observations.

There is also extended structure surrounding the two nuclei that align with H_2 V = $1-0$ S(1) and S(5) concentrations presented in Max et al. (2005) (their Figures 10 and 11, recreated in our Figure 6.7): a molecular concentration to the NE of the northern nucleus, a concentration to the SW of the southern nucleus, and dim, diffuse emission extending to the SE of the southern nucleus. The faint arms of H_2 extending SE and SW observed in Max et al. (2005) align with these faint arms extending SE and SW from the central concentration and are also bright in Fe XXV and Hα (Max et al. (2005); Wang et al. (2014)). Wang et al. (2014) postulates that these faint molecular arms are molecular gas entrapped and shocked by the superwind caused by a vigorous starburst in the southern nucleus. Similarly, Max et al. (2005) argues that these arms are a thin layer of gas at the edges of soft X-ray bubbles observed in Komossa et al. (2003), where a starburst driven wind is "driving shocks or ionization fronts into the interstellar medium and surrounding molecular clouds". We therefore conclude that the CO arms to the SE and SW of the southern nucleus are likely associated with the same starburst-driven superwind shocked gas. This is consistent with the weak dust emission at these locations.

6.4.3 Molecular Gas Between the Nuclei: a Test Model Motivated by History and Observations

In Tacconi et al. (1999) and in works since (e.g. Iono et al. (2007); Scoville et al. (2015)), the velocity gradient observed in the central molecular concentration was interpreted and modeled as a disk of gas between the two nuclei. The presence or absence of a molecular disk in the nuclear region of NGC 6240 is the basis for other author's arguments regarding important feedback processes, such as the outflow studied in Cicone et al. (2018). The disk model has been claimed to be unphysical in other studies of molecular gas in NGC 6240 (e.g. Gerssen et al. (2004)), but no further modeling has definitively proven or disproven the internuclear molecular disk model. More generally, the nuclear regions of (U)LIRGs are critical to understand as they are often the location of the highest star formation rates and the origin of many feedback processes such as stellar winds and AGN outflows (Alonso-Herrero, 2013). For these reasons, we model this nuclear region as a concentration of gas with a disk-like velocity profile using observations in CO J = $3-2$ and J =

$6-5$ to explore in detail the validity of this disk model. Our fiducial model finds the molecular gas concentration between the two nuclei is flat and pancake-like, with a high velocity dispersion that could cause it to dissipate within ~ 0.3 Myr. Its geometry suggests this transient structure could be a tidal bridge connecting the two nuclei.

For CO J = $6-5$ and J = $3-2$, the moment 1 observations in Figures 6.2 and 6.1 do suggest a potential disk with a clearly delineated redshifted to blueshifted structure. This structure was the motivation for disk models in previous observations. The velocity dispersion maps (moment 2) are less reminiscent of a simple disk structure, with both lines showing a pocket of high dispersion gas at the location of the highest velocity redshifted gas reflecting the wide line profiles present at this location. Ignoring this pocket of high dispersion, the CO J = $6-5$ moment 2 map is reminiscent of a disk with lower dispersion (80 km/s) towards the edge of the nuclear gas and higher dispersion (150 km/s) towards the center. The moment 2 map of CO J = $3-2$ is more complex, with more structured pockets of low and high dispersion gas. This difference is unlikely to be due to unresolved gas in CO J = $6-5$ because the resolutions of the CO J = $6-5$ and J = $3-2$ observations are nearly identical. Rather, it is more likely that the lower sensitivity and larger pixel sizes of the CO J = $6-5$ observation causes the difference in the observations.

A self-gravitating disk between the two nuclei is an unexpected phenomenon in a merging galaxy system like NGC 6240. One would expect that if the gas had time enough to settle into a self-gravitating disk it would have settled around one of the nuclei. The stellar population of NGC 6240 has followed this expectation and settled into two populations rotating around each of the nuclei, with the majority of the luminosity dominated by stars predating the merger(Tecza et al. (2000), Engel et al. (2010)). Therefore, a disk of gas between the two nuclei would indicate a separation of molecular gas and stars. A few processes could cause this separation, for example tidal forces channeling gas to the area between the nuclei (Ohyama et al., 2003), by an outflow powered by an AGN (Cicone et al., 2018), or via a head-on collision similar to the Bullet Cluster as proposed by Nakanishi et al. (2005). It is also possible the molecular gas disk settled around the center of mass (COM) of the system. We check if the gas is settled around the COM using dynamical masses from Engel et al. (2010) to locate the COM and compare to the location of the maximal line emission, the expected center of a self-gravitating disk. Engel et al. (2010) used stellar kinematic data to calculate the enclosed masses of the northern and southern nuclei at 250 pc and 320 pc, respectively, finding $M_{enclosed,N}$ = 2.5$\times10^9$ M$_\odot$ and $M_{enclosed,S}$ = 1.3$\times$ 10^{10} M$_\odot$. Using these masses, the center of mass is located at a projected distance of 132 pc (0.25$''$) from the southern nucleus along the axis between the two nuclei. This is at a projected distance of 80 pc (0.16$''$) from the maximum of the integrated CO J = 6 – 5 line emission, where we would expect the center of a gaseous disk to be. This is less than one beam FWHM of separation, indicating it is possible that this gas is coincident with the nuclei's COM.

The question remains of the proposed disk's position angle. Following the methodology presented in Cicone et al. (2018), we separate the observations into "quiescent" [-200, 250] km/s gas and high velocity gas. Figure 6.8 shows the contours of the observed highly redshifted (CO J = 3 – 2: [250, 700] km/s, CO J = 6 – 5: [250, 740] km/s) and highly blueshifted (CO J = 3 – 2: [-500, -200] km/s, CO J = 6 – 5: [-420, -200] km/s) gas in panels a and b. The contours of moderate velocity, or "quiescent" gases of velocity [-200, 250] km/s are shown in in panels c and d. The velocity gradient observed in moment 1 for CO J = 3 – 2 and J = 6 – 5 remains clearly present in the high velocity gas, with distinct centroids of emission at a PA of $\approx$ -34°. The lower velocity gas has much more overlap between the redshifted and blueshifted components, without as clear of a separation between centroids of emission. This lower velocity gas aligns approximately with the semi-major axis of the central gas concentration at a PA of $\approx$ 0°, along the north/south axis. Due to the distinct gradient present in the high velocity gas that more closely resembles a traditional disk, we choose the position angle based on the highest velocity gas. However, we find during modeling that the position angle does not influence the fiducial fit parameters outside of acceptable ranges already incorporated into the model.

We must also choose a location for the disk's center. If the gas were a simple disk, we would expect the maximum of the line emission to align closely with the turnover point of the velocity gradient. As noted in Tacconi and confirmed in these observations, this is not the case. This could potentially be explained by a disk with a highly asymmetric mass distribution leading to an offset in the maximum line emission. Optical depth effects could also displace the velocity gradient turnover point from the maximum of the line emission. Therefore, it is difficult to constrain the central location of this theoretical disk.

Taking into account all of these observations, we argue that CO J = 6 – 5 shows potential evidence of a disk at PA $\approx$ -34°. The center of the disk is likely to be location at the maximum of the line emission following the methodology outlined in Tacconi et al. (1999). Our modeled gas concentration is therefore centered at the maximum of the line emission and inclined to a PA of -34° to enable comparison to the proposed disk model.

6.4.4 Code for Generating the Test Model: the Line Modeling Engine (LIME)

To generate the model that tests the validity of the disk interpretation of the observations, we use the Line Modeling Engine (LIME) (Brinch and Hogerheijde, 2010). LIME uses non-local thermodynamic equilibrium radiative transfer and molecule rotational energy level populations calculations to predict the line and continuum emission from astronomical sources. The user defines a 3-D model describing the density of hydrogen molecules of the astronomical source, then assigns a 3-D temperature and velocity distribution to those molecules. The gas-to-dust ratio, the abundance of CO, and the 3-D distribution of the dust temperature are also set by the user. Finally, the user defines the distance to the source to obtain the correct distance scale per pixel and brightness observed.

The code does not require that the sources are in local thermodynamic equilibrium, and instead solves for population levels iteratively until the model populations have converged at all grid points. After convergence, LIME ray-traces photons to obtain an image of the modeled source at a user-defined observing angle. This simulation methodology allows great flexibility in geometries and kinematics of the simulated source and minimizes the assumptions made about the source to generate line profiles.

To compare the observed data to the simulated model images, we use the CASA package to

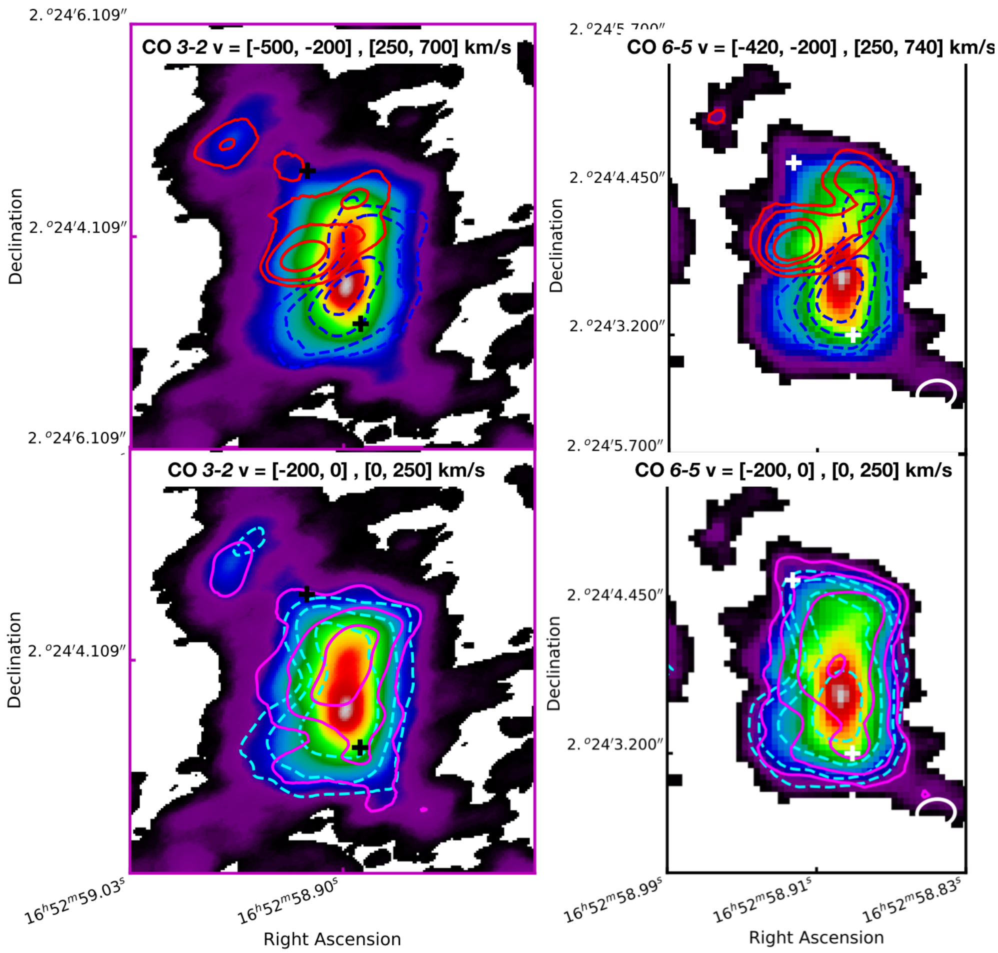

Figure 6.8: Contours of high and moderate velocity redshifted and blueshifted gas for CO J = 3 − 2 and J = 6 − 5, plotted over corresponding integrated line emission (moment 0). Contours are 0.2, 0.4, 0.6, and 0.8 of the maximum in each map. Beam FWHM contours are plotted in the bottom right of panels b and d. *Crosses* denote locations of the AGN. **a**: CO J = 3 − 2 high velocity gas; *red contours:* [250, 720] km/s, *dashed blue contours:* [-880,-200] km/s **b:** CO J = 6 − 5 high velocity gas; *red contours:* [250, 780] km/s, *dashed blue contours:* [-720,-200] km/s **c:** CO J = 3 − 2 moderate velocity gas; *red contours:* [0, 250] km/s, *dashed blue contours:* [-200,0] km/s **d:** CO J = 6 − 5 moderate velocity gas; *red contours:* [0, 250] km/s, *dashed blue contours:* [-200,0].

smooth and continuum subtract the simulated images. The resulting images have the same beam size as the observational images. We must also account for the redshift of the galaxy, which is not included in the LIME simulations. To do this, we find an average line profile for the entire observation in CO J = 3 − 2 and CO J = 6 − 5. We find the central velocity of these line profiles to be 7200 ± 20 km/s for CO J = 3 − 2 and 7465 ± 20 km/s for CO J = 6 − 5. We subtract these velocities from the NRAO-provided maps and can then compare the simulated maps to the observed maps directly.

6.4.5 Parameters of the Test Model

Using LIME, we model the 3-D central gas concentration with density and velocity definitions that are general to any concentration with exponentially decaying density and velocity profiles. Our 3-D model for the disk density follows the function

$$n_{H2} = n_{H_2,0} \exp(-r/r_{sh}) \exp(-|z|/z_{sh}), \tag{6.2}$$

where $r = \sqrt{x^2 + y^2}$, r_{sh} is the scale height of density in the disk radius, z_{sh} is the disk scale height, and $n_{H_2,0}$ is a characteristic H_2 number density of the disk. $n_{H_2,0}$ is calculated such that the total mass of the disk within the extent of our ALMA observations is g_{mass}. We also allow an overall density asymmetry by multiplying the density on one half of the disk (x>0) by a model parameter a_n. For the CO J = $3-2$ model we also allow the extent of the gas to differ for x>0 and x<0 by varying the scale height on either side of the disk.

We model the velocity with

$$v(r) = v_{circ}(1 - \exp(-r/r_{circ})), \tag{6.3}$$

where v_{circ} is the circular velocity and r_{circ} is the radial scale height of the circular velocity.

In Equations 6.2 and 6.3, parameters except r and z are fit parameters that can be tuned to match the resultant modeled line profiles to the data. In addition to those parameters, the LIME model requires a disk inclination, position angle, dust temperature T_{dust}, turbulent velocity v_{turb}, gas temperature T, and gas-to-dust ratio (set to 100). All available parameters, fiducial fit results, and acceptable ranges of fiducial fit parameters can be found in Table 6.3 and are discussed in Section 6.4.6.

The position angle of the gas concentration was modeled to be -34°, the angle between the maximum of the highly blueshifted gas and that of the highly redshifted gas from Figure 6.8. This is close to the position angle used by Tacconi of -40°. Changing the position angle of the gas concentration between these two values does not change the resultant fit parameters beyond their already established acceptable ranges.

6.4.6 Test Model Fitting Strategy and Fiducial Fits

In NGC 6240 the CO J = 6 – 5 emission has been found to be dominated by hot gas while the CO J = 3 – 2 emission is dominated by by cool/warm gas (Kamenetzky et al., 2014), meaning observations of the two transitions trace different gas temperatures. By extension, any models generated for the two transitions trace gas populations of different temperatures. To capture these physics we generate models to match the observed line profiles of CO J = 6 – 5 and J = 3 – 2 separately. Between the models of the two transitions we allow gas mass, density distribution, and gas temperature to vary. However, we maintain the same velocity structure, gas inclination and position angle, and dust temperature between the two models. This strategy requires that the cool/warm and hot gas occupy the same gravitational structure, but allows their extent and density within that structure to vary.

Due to the simpler nature of the CO J = 6 – 5 line profiles compared to those of CO J = 3 – 2, we first use CO J = 6 – 5 observations to constrain the velocity structure, inclination and position angle, and dust temperature. We then applied the fiducial fit parameters from the CO J = 6 – 5 model to find the gas density distribution and gas temperature of CO J = 3 – 2.

The fiducial model parameters are found by eye due to the large parameter space and number of fit parameters available to the model. The parameters are chosen to minimize the apparent differences between the modeled line profiles and the observed line profiles extracted at points separated by a beam FWHM along the major and minor axes of the modeled gas. Once a fiducial fit is found, ranges on acceptable fit parameters are found by varying one parameter at a time until the model's line profiles are no longer acceptably close to the observed line profiles. "Acceptable" fits are determined by a combination of factors: first, the modeled line profiles on average cannot appear to differ from the observed line profiles by more than $\sim 25\%$ in height, width, or central velocity. Second, the shape of the modeled line profiles should approximately match that of the observed, e.g. highly doubly-peaked profiles are unacceptable if the observed profiles are largely singly-peaked.

Many of the parameters can be well-constrained by the observed line profiles, for example r_{sh} by the rate of decay of the line profile heights with radius and v_{turb} by their width. Other parameters are more difficult to constrain. Due to the face-on nature of the gas concentration constrained by the singly-peaked profiles and the lack of gradient in their central velocities, v_{circ} is difficult to constrain leading to uncertainty in r_{circ}. However, v_{circ} must be high enough to create the singly-peaked profiles as higher circular velocities and low optical depths blur line profiles and cause them to appear singly-peaked. An upper bound on v_{circ} is determined because exceedingly high v_{circ} cause the line profiles' centers to deviate from the observed profiles.

All available parameters, fiducial fit values, and acceptable ranges of model parameters can be found in Table 6.3. The fiducial fit model is a nearly face-on pancake-like distribution of gas, with a scale height z_{sh} of 20 pc and a radial scale height of 400 pc at an inclination of 22.5° from face-on. The line profiles extracted from the fiducial models along the semi-major axis of the gas concentration for CO J = 6 – 5 and J = 3 – 2 are compared to those from the data in Figure 6.9.

6.4.7 Are the Fiducial Model Parameters Physical?

To explore the reasonableness of the modeled parameters, we compare the fiducial model values to those modeled by Kamenetzky et al. (2014). Similar to the dust temperature modeling described above, Kamenetzky et al. (2014) use *Herschel Space Observatory* observations of CO J = 1 – 0 to 13 – 12 to model gas temperatures, column densities, gas masses, and number densities for NGC 6240. The *Herschel* beam finds galaxy-averaged properties with a beam FWHM of $43''.5$ (Kamenetzky et al., 2014), leading to important differences between our model parameters and theirs. However, these differences can be attributed to probable physical differences between the nuclear and integrated gas properties in NGC 6240. We find that within the acceptable range of model parameters, the masses and temperatures of the gas within both models are consistent

Variable Name	Fit Value	Range
Shared Fiducial Model Values		
Inclination	22.5°	[15, 30]°
Position Angle	33.7°	[30, 40]°
T_{dust}	30 K	[20, 100] K
v_{turb}	140 km/s	[120, 160] km/s
v_{circ}	100 km/s	[50, 300] km/s
r_{circ}	200 pc	[0, 500] pc
CO J = 6 – 5 Fiducial Model Values		
Gas Mass	6.0×10^{8} $M_{\odot}$	$[4.0, 7.8]\times10^{8}$ $M_{\odot}$
r_{sh}	390 pc	[290, 500] pc
z_{sh}	20 pc	[18, 28] pc
a_n	1.0	[0.6, 1.0]
T	2,000 K	[1,500, 3,000] K
CO J = 3 – 2 Fiducial Model Values		
Gas Mass	7.6×10^{8} $M_{\odot}$	$[2.9, 9.2]\times10^{8}$ $M_{\odot}$
r_{sh}	350 pc	[250, 400] pc
$r_{sh,x<0}$	250 pc	[200, 400] pc
z_{sh}	13 pc	[10, 16] pc
a_n	1.0	[0.8, 1.8]
T	600 K	[400, 1500] K

Table 6.3: Fiducial LIME model values for CO J = 3 – 2 and J = 6 – 5 (*column 2*), as well as the ranges in parameter space that resulted in acceptable models (*column 3*). Gas masses are calculated by integrating the modeled density over the observed extent of the gas (the high signal-to-noise areas plotted in Figures 6.1 and 6.2).

with previous observations. However, the ratio of the modeled circular velocity to the velocity dispersion is high for such a thin, extended self-gravitating disk. Additionally, the large modeled velocity dispersion and the thin vertical extent indicate a transient structure that could dissipate within 0.3 Myr. Finally, the modeled velocity dispersion is comparable to the escape velocity of the gas concentration. These discrepancies suggest the gas is unlikely to be a self-gravitating disk.

Our fiducial models show that the CO J = 6 – 5 emission is dominated by hot (T = 2,000 K) gas, and the CO J = 3 – 2 emission by warm (T = 500 K) gas. In Kamenetzky et al. (2014) the CO emission from NGC 6240 is modeled as a combination of cold and hot gas. In their Figure 8, the CO J = 3 – 2 emission is shown to be dominated by a mixture of cold (16 K with 1σ range

[5, 50] K) and hot (1,260 K with 1σ range [790, 2,000] K) gas, while CO J = $6-5$ emission is produced primarily by the hot gas component. Our CO J = $3-2$ fiducial model temperature of 600 K lies between their modeled temperatures of these two components, supporting the argument that the CO J = $3-2$ emission is created by a combination of gas temperatures. Our CO J = $6-5$ observation is modeled to be dominated by hot gas, like their model, with a fiducial fit temperature of 2,000 K at the upper bound of their 1σ range. This is potentially due to the smaller ALMA beam sizes probing the nuclear region while Kamenetzky et al. (2014) used *Herschel* data that combines emission from both nuclear and extended regions. The nuclear region is likely to be much warmer than the extended gas due to star formation and AGN activity concentrated in the central kpcs. These high temperatures are supported by the presence in this region of $H_2 1-0S(1)$ that requires temperatures > 1000 K to be excited (Max et al. (2005); Meijerink et al. (2013)).

Kamenetzky et al. (2014) also models the masses of these cool and warm components, finding a cool gas mass of $2\times10^9 \pm 5\times10^8$ $M_\odot$ and a warm gas mass of $4\times10^8 \pm 10^8$ $M_\odot$. Our total modeled mass for CO J = $3-2$ is 7.6×10^8 $M_\odot$, higher than Kamenetzky et al. (2014)'s modeled warm gas mass but lower than their cool gas mass. This is consistent with the theory that CO J = $3-2$ is a mixture of cool and warm gas components. Our total modeled hot gas mass (CO J = $6-5$) is 6×10^8 $M_\odot$, higher than their best-fit warm gas mass by around 50%. Accounting also for our large range of acceptable masses, the range of masses for our warm gas model contains Kamenetzky et al. (2014)'s masses.

The total modeled mass of this test model (CO J = $3-2$ plus CO J = $6-5$), 1.4×10^9 $M_\odot$, is approximately equal to the nuclear region's (region 1, Figure 6.6) mass calculated in Section 6.4.1 from the continuum emission of 1.2×10^9 $M_\odot$. If this central region is indeed a disk, we would expect the disk to comprise of the majority of its mass. The hot gas modeled by CO J = $6-5$ (6×10^8 $M_\odot$) is about 50% of the mass of this same region.

The fiducial fit dust temperature of 30 K, though not tightly constrained, contains Kamenetzky et al. (2014)'s best fit dust temperature of 56 K that we used in our continuum mass calculations.

The ratio of our average rotational velocity to our modeled velocity dispersion is $<v>/\ \sigma \sim 1$. According to Tacconi et al. (1999) this is in the range that indicates a disk that must be geometrically thick. However, our fiducial model is quite thin, only tens of pc thick in comparison to a radial extend of hundreds of parsecs. Thicker fiducial models did not fit the observed data because their thickness causes more self-absorption, resulting in asymmetric and doubly peaked profiles. This aspect ratio is unlikely for a self-gravitating disk with such high velocity dispersion since the dispersion would have the effect of "puffing up" the disk. However, it is important to note that the acceptable range of values on the modeled parameters are large and there are ranges of v_{circ} and v_{turb} that allow this ratio to be < 1. Despite this, the acceptable range on v_{circ} mostly covers a range $>v_{turb}$, and therefore this argument holds for the majority of the acceptable range of parameters. It is also possible that with such a thin concentration of gas, the modeled high velocity dispersion could partially be accounted for by shear. That is, the gas concentration is being tidally disrupted in the plane of the sky.

We can also conduct a back-of-the envelope calculation to determine the lifetime of the gas concentration given its modeled velocity dispersion of 140 km/s (assuming no contribution from shear). If the vertical extent of the concentration is ~ 40 pc (two vertical scale heights), gas moving at 140 km/s would travel this distance after 300,000 years. This is an extremely short timescale in the context of galaxy mergers that occur over timescales of Gyr, meaning the model is suggesting a highly transient structure, an unlikely scenario for a self-gravitating disk.

As a further test of the fiducial model in the context of a self-gravitating disk, we can compare the velocity dispersion and circular velocities to the escape velocity from the modeled gas concentration. The escape velocity $\sqrt{2GM/r}$ for the entire gas concentration of 13.6×10^{8} $M_{\odot}$ at a distance of r_{sh} = 390 pc from the center is 170 km/s. This is comparable to the modeled velocity dispersion of 140 km/s, and is less than the circular velocity (100 km/s) added to the modeled dispersion. This means the gas' velocity can exceed the escape velocity of the gas concentration, indicating it is extremely unlikely this gas is a self-gravitating disk, especially in the context of the velocity dispersion indicating a transient structure.

Given the fiducial LIME model, we can also check the stability of the gas concentration by calculating the Toomre parameter $Q = \sigma_r \kappa \ / \ \pi \ G \ \Sigma_{gas}$. Here $\kappa = \sqrt{3} \ v_{max} \ / \ R$ is the epicyclic frequency, σ_r is the line-of-sight velocity dispersion, and Σ_{gas} is the mass surface density of the gas (Toomre, 1964). Q for the fiducial LIME model is 2.3 for CO J = $6-5$ and 2.5 for the CO J = $3-2$ model. A $Q > 1$ means the gas is stable against collapse at this time, with the majority of models within the acceptable parameter ranges fitting into this category. These high values of Q are consistent with this highly turbulent system. The majority of the luminosity being generated in this central mass concentration is then unlikely to be due to star formation, consistent with other papers that argue for heating from shocks and superwinds originating in starbursts around the nuclei, not in the central region between the two nuclei (e.g. Tecza et al. (2000); Max et al. (2005); Engel et al. (2010)).

6.4.8 Model Residuals: What is Captured by the Test Model?

Figures 6.10 and 6.11 show the normalized residuals in moments 0 and 1 between ALMA data and the fiducial LIME models for CO J = $6-5$ and J = $3-2$. The residuals for each moment take a similar shape for both emission lines. The low values of the moment 0 fractional residual in the central brightest region indicate that the test model accurately describes the majority of the line emission in this central region. However, the model misses the northernmost portion of the observed gas concentration, especially apparent in the CO J = $6-5$ residual. The test model does not accurately describe the $<v>$ data, as evidenced by the large values and amount of structure that remains in the fractional residuals of moment 1 for both CO J = $6-5$ and J = $3-2$. The non-normalized residual of $<v>$, shown in the bottom panel of each figure, remains largely unchanged from the observed values.

The residuals shed light on what components of NGC 6240's molecular gas are captured by the test model, if any. The low values for the integrated line emission residuals in the central/southern portion of the central gas concentration indicate that the models capture the majority of the CO emission in this, the brightest part of the ALMA observations. Capturing the majority of the integrated line emission indicates that we in turn have a decent constraint on the mass and temperature of this gas. This is important because good constraints on mass and temperature indicate we also have an accurate modeling of the *physical distribution* of the gas. Our models found that this central gas region is pancake-like: quite thin (tens of pc) with respect to its extent (hundreds of pc).

While the fractional residuals of the line emission appear to accurately capture the *distribution* of the gas, the velocity residuals tell a different story. The velocity residuals in this central area show a structural residual gradient across the modeled area, suggesting the simple exponentially decaying velocity profile (like that of a disk) does not accurately capture the *kinematics* of the gas. When taken with the arguments in Section 6.4.7 that the modeled velocity dispersion indicates a transient structure whose velocities can exceed the escape velocity of the gas concentration, we suggest that the nuclear molecular gas emission is dominated by a thin pancake-like gas concentration that is not rotating like a disk. This conclusion is also consistent with the findings of Section 6.4.7, finding our masses and temperatures fit within current observations but the high velocity dispersion relative to the circular velocity is unlikely for such a thin geometry if the gas were to be a self-gravitating disk.

6.5 Is the Central Molecular Gas a Tidal Bridge?

At the location of the brightest molecular gas emission the fractional residuals are quite small, indicating we accurately model the mass and temperature of the central molecular gas concentration. Gas mass and T are degenerate in the LIME modeler as increases in either parameter also increases the emission. Exceedingly high masses or low temperatures result in doubly-peaked profiles, providing constraints on both parameters. Accurate models of mass and temperature do not directly imply accurate models of gas distribution. The other parameters important for this are r_{sh}, well constrained by the visible extent of the gas, and z_{sh}. z_{sh} is somewhat difficult to constrain because the gas is nearly face-on, but thick gas result in doubly peaked profiles constraining z_{sh}

to small values (tens of parsecs). Density is directly calculated from the gas mass, r_{sh}, and z_{sh}. Our constraints on these values and the gas temperature indicate we have successfully modeled the *distribution* and density of the gas. According to the model this indicates that the gas is rather thin (tens of parsecs) with respect to its extent (hundreds of parsecs). However, as argued in Section 6.4.8 the central molecular gas kinematics are not well captured by the model as demonstrated by the high fractional residual values for $<v>$. This means we must consider other possibilities to explain the kinematics of this gas. The extended distribution and extreme aspect ratio suggests a tidal bridge between the two nuclei, stretched thin by tidal forces along the semimajor axis of the observed CO emission at PA = 0°.

If the gas is indeed a thin tidal bridge, its low density means it is unlikely to be in virial equilibrium and could result in sub-thermal excitation. To check for the possibility of sub-thermal emission we compare the observed density to the critical densities n_{crit} for the observed lines. The modeled average H_2 density is $\sim 1\times10^3$ cm^{-3}, to be compared to n_{crit} for CO J = $3-2$ emission of 3.6×10^4 cm^{-3} and 2.9×10^5 cm^{-3} for CO J = $6-5$. n_{crit} for both transitions is 1-2 orders of magnitude higher than the modeled density, indicating that sub-thermal emission is likely in this central region. Sub-thermal emission from non-virialized clouds could "easily boost the emission from gas in a drawn-out, thin and diffuse gas bridge compared to denser material around the nuclei" due to lower self absorption in the thinner gas (Engel et al. (2010); (Tacconi et al., 2008)). Essentially they argue the X-factor (the relationship between $N(H_2)$ and the CO luminosity) only applies for virialized, denser clouds, and can be radically modified for non-virialized, thin clouds. This would result in the observed phenomenon that the peak of the CO emission does not correspond to the peaks in the dust emission around the nuclei, which in turn supports the idea that this is a thin tidal bridge. This sub-thermal emission could help to explain the anomalously high L_{CO}/L_{FIR} observed in NGC 6240, around 10 times higher than other nearby galaxies (Kamenetzky et al., 2014).

The presence of a tidal bridge or ribbon between the two nuclei is also supported by observations of H_2 $1-0$ S(1) and S(5) presented in Max et al. (2005) and recreated in our Figure 6.7. They observe a ribbon of H_2 with a reverse S shape (their Figure 10) extending between the northern and southern nuclei and postulate it could be material flowing along a bridge connecting the two progenitor galaxy nuclei. They compare this geometry to the tidal bridges predicted by computer simulations (e.g. Barnes and Hernquist (1991); see also Figs. 9 and 10 in Barnes and Hernquist (1996); Figs. 1 and 4 in Barnes (2002)). One important note is that these simulations show tidal bridges on scale of $\sim$ 5-40 kpc, while the projected separation between the nuclei in NGC 6240 is only $\sim$ 1 kpc. Nonetheless, smaller scale simulations (e.g. Hopkins et al. (2013)) also show significant mass flowing between galaxy nuclei on scales $\lesssim$ 5 kpc. While the irregular H_2 morphology does not follow the morphology of CO, this could be explained by an altered CO-to-H2 conversion factor in this nuclear region (Tacconi et al., 2008) causing CO to be brightened with respect to H_2 and dust along the N/S axis. The differences in the two observations are unlikely to be caused by differences in sensitivities or beam sizes, as the H_2 observations show offsets in peak brightness located at the southern nucleus, farther than 1.5 beam FWHM away from the peak brightness observed in the CO observations.

Tidal bridges and filaments have been shown to have relatively small line-widths ($\sim$ 50 – 100 km/s) in tidal dwarf galaxies (Braine et al., 2001). In larger galaxies the linewidths and velocities are larger, for example the CO FWHM is $\sim$ 200 km/s with $<v> \sim$ 130 km/s in the bridge between the Taffy galaxies (Braine et al., 2003). In Arp 194 (total dynamical mass $> 10^{11}$ $M_\odot$), the dispersion is $\sim$ 25-125 km/s with $|<v>| <$ 150 km/s in the bridge connecting the two galaxies (Zasov et al. (2016), their Figure 3). In Figure 6.12 we plot the velocity and dispersion along the proposed axis of the tidal bridge (PA = 0°) to compare to these observed line widths and velocities. Following the methodology in Cicone et al. (2018), we also separate the quiescent (-200 to 250 km/s) from the high velocity gas as we did for Figure 6.8 and extract the moments for moderate velocity (quiescent) gas alone as well as for the entire observation. For both the quiescent gas as well as the full moment extraction, the velocities all lie below approximately $|100|$ km/s, with a structure that moves smoothly from redshifted to blueshifted, as expected for a tidal bridge. The total velocity dispersion shows values of 150-180 km/s, slightly higher than the bridge in Arp 194 but within range of the Taffy galaxy bridge. When we separate out the quiescent gas the velocity dispersions fall to approximately 110 km/s for this quiescent gas, the upper limit for the expected velocity dispersion of a tidal bridge or tail for even the small tidal dwarf galaxies.

6.6 Highly Redshifted Gas Origins

The modeled line profiles in Figure 6.9 distinctly show excess redshifted emission not captured by the model in the northeast (the dark brown/black line profiles). This is the same highly redshifted gas seen in the average velocity maps ($\sim$ 300 km/s in Figure 6.1 and $\sim$ 400 km/s in Figure 6.2) and in the channel maps (up to 665 km/s in Figure 6.3 and 724 km/s in Figure 6.4). It is also the highly redshifted gas plotted in Figure 6.8 with an average velocity of $\sim$400 km/s when isolated from the quiescent gas, or $\sim$250 km/s when including the quiescent gas. This gas is dim compared to the majority of the line emission, and as such does not augment the normalized fractional residual moment 0 maps above $\sim$0 despite not being described by the models.

This gas' extremely high velocity, especially as compared to the rest of the galaxy, implicates it as worthy of discussion. Given the current observations, we cannot definitively distinguish between the different energetic processes that could result in the high velocities observed. The gravitational energy of the merger is sufficient to have accelerated or slingshotted this gas to its current velocity. There is no definitive evidence for or against this gas representing a dwarf galaxy disturbed and accelerated by the merger. The gas' morphological and velocity correspondence to outflows studied by other authors (Cicone et al. (2018), Müller-Sánchez et al. (2018)) is a compelling argument, but other signatures of an AGN accelerating this gas are not present. Details of these three arguments are presented below.

6.6.1 Gravitational Energy from the Merger

It is possible that the gas' kinetic energy comes purely from the gravitational forces at work during the merging process.

As a first test, we conduct an order-of-magnitude calculation and compare the kinetic energy of the high-velocity gas to the gravitational potential energy of the nuclei. There is no observed continuum emission at this location, so we instead calculate the mass according to Solomon and Vanden Bout (2005), their equations 3 and 4. They find the hydrogen mass is

$$M(H_2) = \alpha 3.25 \times 10^7 S_{CO} \Delta v \nu_{rest}^{-2} D_L^2 (1+z)^{-1} M_\odot, \tag{6.4}$$

where α is known as the CO to H_2 conversion factor ($\sim$ 0.3 for NGC 6240 Kamenetzky et al. (2014)), $S_{CO}\Delta v$ is the velocity-integrated flux density in Jy km s^{-1}, ν_{rest} is the rest frequency of the observed line in GHz, D_L is the luminosity distance of the object in Mpc, and z is the redshift. If we integrate the flux density and average velocities over the high-velocity gas only, as in Figure 6.8 panel *a*, we find a mass of 9.3×10^7 $M_\odot$ and an average velocity of 400 km/s. The kinetic energy to accelerate this high-velocity gas is 1.5×10^{56} ergs. The gravitational potential energy Φ of the nuclei using the masses from Engel et al. (2010), $\Phi = GMm/r$, and the projected distance r is $\sim$ 1.5 kpc is $\sim 3\times10^{57}$ ergs. This amount of energy indicates that the merger contains enoughs energy to accelerate the gas gravitationally. This is unsurprising, especially given that the relative velocity of the merging galaxies was 150-300 km/s (Wang et al. (2014); Tecza et al. (2000)) and the average velocity of this gas is 400 km/s.

It is also possible that the gas participated in a gravitational slingshot with one of the nuclei, much like comets or spacecraft in our own Solar system. When participating in head-on interactions with a heavier body, an object can be accelerated to twice its initial velocity. Since the relative velocity of the merging galaxies was 150-300 km/s, this indicates velocities from 300-600 km/s are possible via this type of interaction.

6.6.2 Dwarf Galaxy

The distinct concentration of high-velocity gas and increased velocity dispersion suggests a possible dwarf galaxy at this location, accelerated and disturbed by the merger. Illustris simulations show that most galaxies undergoing mergers have at least one satellite galaxy (Brainerd and Yamamoto, 2019) and the velocity difference between the host and its satellites is most likely to be 0 km/s (their Figure 1f). We observe a velocity difference between the nuclei and the highly redshifted gas of 300-400 km/s, which is low probability according to their simulations, though not impossible. It is important to note, however, that the size scales studied in these simulations were 10's to 100's of kpc as opposed to the central kpc. Upon review of the literature, we found no evidence of simulations summarizing the impacts of satellite galaxies on mergers' central kpc.

6.6.3 AGN-driven Outflow

6.6.3.1 Kinematic and Morphological Correspondence to Previously Studied Outflows

It is possible that the gas' kinetic energy comes purely from an AGN-driven outflow. AGN outflows have been postulated and studied in NGC 6240 by other authors: Feruglio et al. (2013) proposed a molecular outflow that was then studied by Cicone et al. (2018), and an ionized outflow is observed and modeled in Müller-Sánchez et al. (2018).

Cicone et al. (2018) uses high-resolution nuclear observations of $[CI]^3P_1-^3P_0$ ([CI](1-0)) along with extended observations of CO J = $2-1$ and $1-0$ to study properties of an H_2 molecular outflow extending east/west over 10 kpc and originating in the central kpc. The high-resolution observations are reproduced in part in our Figure 6.13 and the extended observations are reproduced in part in our Figure 6.14. The high-resolution observations of $[CI](1-0)$ in Cicone et al. (2018) show a concentration of highly redshifted ($\sim$ 200 km/s) and high dispersion ($\sim$160 km/s) gas at the same location as our high velocity, high dispersion gas (see Figure 6.13). They postulate this location is the *initiation* point of a molecular outflow, with the elevated velocity dispersion indicating a possible widening of the outflow at that location due to a collision with quiescent gas. In the context of our observations, this high velocity dispersion could also be a geometric effect resulting from the superposition along our line of sight of emission from the AGN outflow and the more quiescent gas of the tidal bridge.

This redshifted gas is also observed in H_2 (Engel et al. (2010) H_2 V = $1-0$ S(1); see Figure 6.13). The average velocities of this highly redshifted gas from these, Cicone et al. (2018)'s, and our observations are comparable (H_2 V = $1-0$ S(1): 200 km/s, $[CI](1-0)$: 200 km/s, CO

J = $3-2$: 250 km/s, CO J = $6-5$: 300 km/s), though our velocities appear slightly higher than the other observations. Considering their comparable velocities and similar morphology, we argue that it is possible that all of these observations trace molecular gas subjected to the same energetic processes. Specifically, the correspondence in morphology and velocity between our CO observations and the [CI]($1-0$) indicate an association between the highly redshifted gas presented here and the east/west molecular outflow studied by Cicone et al. (2018).

The ionized outflow presented by Müller-Sánchez et al. (2018) was observed in [O III] and modeled as a truncated conical AGN outflow with a base that encompasses both nuclei. They interpreted this outflow as a cone produced by both AGN but with the majority of the energy stemming from the southern nucleus. A cartoon of their modeled outflow cone atop their [O III] observations is reproduced in our Figure 6.14. The average velocity of the [O III] outflow at the radius corresponding to the location of our highly redshifted gas is $\sim$200 km/s (Müller-Sánchez et al. (2018), their Figure 3d reproduced in part in our Figure 6.15). The average velocities at this location from our observations are 150 km/s (CO J = $3-2$) and 180 km/s (CO J = $6-5$), consistent with the [O III] velocity. These comparable velocities suggest that our highly redshifted gas is associated with the [O III] outflow. As a note, these average CO velocities were found using $0.5''$ by $1.5''$ boxes oriented along the slit (PA_2) used to extract [O III] velocities in their work, at a distance of $\sim1''$ from the galaxy's center. This is a different velocity extraction than we use in the rest of the paper and captures not only the highly redshifted gas but some of the more moderate velocity gas as well. Notably, the velocities of our molecular gas exceeds that of the H_2 gas extracted at that same location ($\sim$ 100 km/s; Figure 6.15), in spite of the morphological correspondence between the H_2 observation and ours shown in Figure 6.13. The velocities of the CO gas presented in this work do appear higher than the H_2 in Figure 6.13, as noted above. This difference can be potentially resolved in the context of the interpretation we present, with molecular gas emission resulting from outflowing gas superposed on emission from a tidal bridge with more quiescent velocities: the H_2 has a lower optical depth than CO and therefore could contain more emission from the more quiescent gas of the tidal bridge.

In summary, we argue that the redshifted gas has a probability to be associated with both the molecular outflow presented in Feruglio et al. (2013) and Cicone et al. (2018) as well as the ionized outflow modeled in Müller-Sánchez et al. (2018). The NE-oriented angle of the [O III] outflow and the large modeled opening angle (50°, Müller-Sánchez et al. (2018)) means the molecular and ionized outflows are spatially coincident (see Figure 6.14). The spatial coincidence along with the kinematic and morphological correspondence between the three observations suggests a possible multiphase outflow. The origin of this outflow is likely to be the southern AGN, the primary energy source for the ionized outflow presented in Müller-Sánchez et al. (2018). The highly redshifted gas (CO, [CI]$(1-0)$, and H_2 $(1-0)$ $S(1)$ in Figure 6.13) points along the axis of the ionized outflow towards the southern AGN. However, we cannot definitively point to it as the sole source of the energetics based solely on these arguments.

The idea of a multiphase outflow (an outflow containing cool gas as well as hot, ionized gas) is relatively new, with uncertain driving mechanisms from broad winds driven by radiation pressure to the mechanical action of radio plasma flowing from the AGN (Morganti, 2017). The examples of multiphase outflows that have been discovered find that molecular gas can often comprise the majority of the mass of an outflow. Only a few galaxies have the extensive multi-wavelength data needed to characterize these outflows, and as such NGC 6240 presents an exciting prospect if this is indeed a multiphase outflow. One example of a multiphase outflow is found in Seyfert 2 IC 5063, presented in Morganti et al. (2007) and Morganti et al. (2015) showing co-spatial ionized and CO emission reaching velocities of nearly 650 km/s. This example bears resemblance to the similarities between the [O III] emission presented by Müller-Sánchez et al. (2018), the [CI]$(1-0)$ from Cicone et al. (2018), and our CO emission.

6.6.3.2 AGN Energetics and Outflow Signatures

While the appearance of the gas kinematics indicate a possible outflow, we must check if the energetics of the AGN could accelerate this highly redshifted gas. The Bolometric luminosity of the AGN was found to be $\sim 5 \times 10^{45}$ erg s^{-1} (Lira et al., 2002). Accelerating the high-velocity gas to its current speed would require 1.5×10^{56} ergs according to the kinetic energy calculation presented above. This much energy could be deposited by the AGN in a very short $\sim$ 1,000 years (if the entire bolometric luminosity is deposited onto this gas, an unphysical assumption). Even if only 1% of the AGN energy is deposited on this gas, the acceleration could occur over 100,000 years, a short time in an astrophysical sense. If we assume the gas was initially close to the southern AGN, it would travel to its location $\sim$ 500 pc away in 10^6 years at its current speed. This timescale is much greater than the 1,000 (or 10,000) years it would take to accelerate the gas to its current velocity. Therefore, if the AGN injected the energy it is unlikely that the gas began proximal to the southern AGN. Rather, the jet would have to encounter the gas at a location closer to where the redshifted gas appears now.

When a jet encounters quiescent gas, shocks occur and velocity dispersion increases. Depending on the observed stage of the shocked gas this can lead to different observed properties. Initially, the shock will heat the gas, potentially dissociating dust grains as well as elevating the CO J = 6 – 5/J = 3 – 2 ratio.

Does the gas appear hotter than its surroundings? To check whether the gas is heated by shocks above the gas surrounding it, we look at the CO J = 6 – 5/ J = 3 – 2 ratio. An elevated ratio should indicate hotter gas. The mean ratio value for the entire observed region is 2.5, elevated to around 3.5 near the nuclei, 2.9 in the central brightest integrated line emission area, and 2.9 around the highly redshifted gas. This means the highly redshifted gas is not likely to be a different temperature than the majority of the molecular gas, which itself has been shown to be shock heated by star formation (Tecza et al., 2000). This does not support nor disprove the idea of AGN-accelerated, shocked gas.

As we do not see conclusive evidence for shock heating, we can look for signatures of post-shock cooling that can lead to the formation of neutral H and cause shock-induced star formation (Croft et al., 2006). Two examples of the post-shock cooling stage are Minkowski's object (Croft et al., 2006) and Centaurus A (Schiminovich et al., 1994). In both cases, star formation and neutral H are observed in the axis of the jet.

Does the highly redshifted gas contain star formation? CO J = $6-5$ and J = $3-2$ are primarily collisionally excited meaning star formation does not contribute much to these lines' overall luminosity, while star formation heats dust efficiently. No continuum emission is observed at the location of the highly redshifted gas, indicating that star formation is unlikely. Additionally, previous observations in the radio X band that traces thermal emission from star formation regions (Hagiwara et al. (2011); Gallimore and Beswick (2004)), ionized lines such as Fe XXV and Fe II that trace for young supernova remnants (Wang et al., 2014), K and H bands that observe current stars (Engel et al. (2010); Pollack et al. (2007)), and UV and Hα that detected knots of star formation in regions at and outside of the nuclei (Yoshida et al., 2016) see no notable star formation at this location. In addition, archival VLA data in the X-band (3.6 cm) that traces star formation sees no emission in this region.

Is there HI at this location? There are observations of NGC 6240 in absorption in HI (Beswick et al., 2001). There is broad H I absorption around the two nuclei, with some faint absorption between the two nuclei. They also observe narrow HI absorption extended around their entire observation. There was no notable HI absorption at the location of our highly redshifted gas, but the absorption was detected against a continuum band that has no observed emission at this location (Beswick et al. (2001), their Figure 1). Therefore we cannot conclusively say whether there is HI present.

There is a notable lack of conclusive evidence for localized shock heating or cooling at the location of the highly redshifted gas. It is also possible for AGN to entrain gas, leading to a more gentle acceleration and a lack of shocks. The gradual acceleration is observationally characterized by slowly increasing velocities along the axis of the jet or outflow (see e.g. Russell et al. (2017)). We do not see this signature, instead the velocity field is discontinuous, with high velocities appearing suddenly. Therefore, it is unlikely that the gas is entrained in the outflow.

6.6.4 Is the Gas Bound?

With such a high velocity, it is an interesting question to ask if the gas is bound to the nuclear region or whether it will eventually escape and deplete the nuclear region of star formation fuel.

The escape velocity v_{esc} from distance r to an object of mass $M(r)$ is $\sqrt{2GM(r)/r}$. The total mass of the nuclear region is approximately the mass of the gas calculated from our observations (Section 6.4.1) added to the mass of the nuclei from Engel et al. (2010) of $M_{north} = 2.5\times10^{9}$ and $M_{south} = 1.3\times10^{10}$ $M_{\odot}$. The highly redshifted gas is approximately 250 pc from the rough COM of the system, indicating an escape velocity of 760 km/s from the nuclear region. This exceeds even the highest velocity gas observed, indicating that although the velocities are extreme the gas is bound to the nuclear region.

However, when we consider nuclear components independently, we find that the gas is unbound from the tidal bridge (modeled mass of 13.6×10^{8} $M_{\odot}$; $v_{esc} \sim 220$ km/s) and from the northern nucleus ($v_{esc} \sim 290$ km/s). The southern nucleus is $\sim$ 500 kpc from the high velocity gas, corresponding to v_{esc} of 470 km/s. This is high enough to capture the majority of the gas, though the gas with the highest velocity could escape from the southern nucleus' gravitational well.

Taken as a whole, these arguments imply that the majority of the highly redshifted gas is bound to the nuclear region and to the southern nucleus specifically. This implies that this gas is likely to eventually feed back onto the southern nucleus to fuel future star formation. This viewpoint is consistent with other observations of outflowing molecular gas (either AGN driven or otherwise) having velocities low enough that they cannot escape (Morganti, 2017).

6.7 Anatomy of the Merger and Comparison to Simulations

The geometry of the merger that formed NGC 6240 is discussed in detail in Engel et al. (2010). They propose that the merger has a geometry that tends towards coplanar/prograde because of the extended tidal tails observed in NGC 6240. The stellar mass calculations suggest the merging galaxies had comparable masses. The stellar kinematics, which suggest the nuclei are counterrotating, support this merger geometry if one nucleus is “inclined behind the plane of the sky, and the other in front” (Engel et al., 2010).

Given an assumed coplanar prograde encounter of two nearly equal mass galaxies as postulated in Engel et al. (2010), we can compare the observed characteristics of NGC 6240 to that of current simulations. Lotz et al. (2008) presents simulations of equal-mass gas-rich mergers with a variety of geometries, including coplanar prograde encounters. All coplanar encounters of gas-rich Sbc-type galaxies simulated for their paper have two peaks of star formation, one around 1 Gyr and another stronger starburst around 2 Gyr after the merger begins. Tecza et al. (2000) argued that we are observing NGC 6240 shortly after the first encounter triggered an initial starburst, as did Engel et al. (2010) who argued NGC 6240 “has currently elevated levels of star formation compared to a quiescent galaxy; and will experience another, likely stronger, peak in star formation rate in the near future when the galaxies coalesce”. This viewpoint is consistent with the coplanar prograde merger results in Lotz et al. (2008) summarized above. Engel et al. (2010) also makes an important note that NGC 6240 is currently barely below the ULIRG classification of $L_{IR} \gtrsim 10^{12}$ $L_{\odot}$, and will likely breach that threshold once the second starburst is triggered, supported by the stronger second starburst predicted by the simulated merger models in Lotz et al. (2008).

The current rates of star formation measured by Engel et al. (2010) of approximately 25 $M_{\odot}$ yr^{-1} support this morphology and timescale according to the simulations in Lotz et al. (2008), which predict star formation rates of 15-30 $M_{\odot}$ yr^{-1} for these merger morphologies after the first pass. However, according to these simulations we cannot infer a tighter constraint than $\approx$ [0.5, 2] Gyr since the time the merger began because the star formation rates are elevated up to 0.25 Gyr before first pass and remain elevated until the time of the final merger according to Lotz et al. (2008). The existence of a tidal bridge in some cases helps to further constrain the time since the merger began (e.g. Hibbard and Mihos (1995)), but generally these require numerical simulations that do not constrain the uncertainty on the age since the beginning of the merger to less than a precision of 1 Gyr. In general, tidal features can be dispersed or fall back onto the galaxy nuclei anywhere between a few hundred Myr to a few Gyr after formation, which also does not help constrain the time since the merger began (Souchay et al., 2013). Therefore, the presence of a tidal bride between the two nuclei only allows us to support previous works (e.g. Tecza et al. (2000), Engel et al. (2010)) and contend that the merger is sometime between first pass and final coalescence, a range of merger ages between approximately 0.5-2 Gyr according to simulations (e.g. Lotz et al. (2008), Hopkins et al. (2013)).

It is possible that the tidal bridge will fall onto the nuclei prior to the second pass and final coalescence, possibly triggering a starburst or further AGN activity. We calculate the free-fall time of the gas of mass m onto a nucleus of mass M using

$$t_{ff} = \frac{\pi}{2} \frac{R^{3/2}}{\sqrt{2G(M+m)}}, \tag{6.5}$$

where R is the distance between the gas and the nucleus. As this is an order-of-magnitude calculation, we split the molecular gas mass in half and assume half will fall to the northern nucleus and half to the southern. We use the nuclei masses of 1.3×10^{10} $M_\odot$ for the southern nucleus and 2.5×10^{9} $M_\odot$ for the northern (Engel et al., 2010). The projected radius from the maximum of the integrated line emission to the northern nucleus is 590 pc and 145 pc to the southern. This means $t_{ff,northern} = 4.3\times10^{6}$ years and $t_{ff,southern} = 2.5\times10^{6}$ years. That is, the free-fall time of the gas onto the nuclei is a few Myr while the time between the first pass and the galaxies' maximum separation prior to the second pass is $\sim$ 400 Myr (Lotz et al., 2008). Therefore, it is likely that this gas will fall onto the nuclei prior to the next pass, possibly adding to the current nuclear starburst. One issue with this interpretation is the high modeled velocity dispersion and small vertical extent that indicate a transient structure that will dissipate after $\sim$0.3 Myr (see argument at the end of Section 6.4.7). Therefore, it is possible that the nuclear concentration of gas will dissipate prior to streaming onto the nuclei.

6.8 Summary

NGC 6240 presents an interesting test case of a galaxy merger between first pass and final coalescence, an intermediate and turbulent stage of galaxy evolution. It provides a detailed stage upon which we can study in detail what comes just before galaxies evolve into ULIRGs, as it is just below the classification threshold of a ULIRG but will evolve into a ULIRG when the second, stronger, merger-induced starburst is triggered. We presented high-resolution ALMA observations of CO J = 6 − 5 and J = 3 − 2, the first observations of the nuclear region of NGC 6240 in CO J = 6 − 5 and the highest resolution observations to date in CO J = 3 − 2. We observe similar morphology to previous CO observations, a concentration of gas between the two nuclei that is distinct from the continuum that is centered around the nuclei (Tacconi et al. (1999); Scoville et al. (2015), etc.). We model the central molecular gas concentration using LIME and find a thin, pancake-like distribution of gas whose velocities and velocity dispersions indicate a transient concentration that is not a self-gravitating disk. We instead argue that the nuclear region shows superposed emission from a tidal bridge and outflowing gas accelerated either by gravitational energy from the merger or an AGN. This work demonstrates the importance of high-resolution multi-line observations when trying to disentangle the effects of energetic gas acceleration mechanisms, star formation, and tidal forces in the central regions of major mergers.

We argue that the majority of the central molecular gas concentration is a tidal bridge connecting the two nuclei of the progenitor galaxies. Our fiducial models found that this central gas region is likely quite thin (tens of pc) with respect to its extent (hundreds of pc). That this central gas is a bridge connecting the nuclei is supported by Engel et al. (2010) who argue for an altered CO-to-H_2 conversion factor in the central region that would be exacerbated by a drawn-out, thin, and diffuse gas bridge. The H_2 observations presented in Max et al. (2005) also support the idea of a tidal bridge, which are in turn motivated by simulations from Barnes and Hernquist (1991) and Barnes and Hernquist (1996) that show mergers can result in material flowing along bridges connecting progenitor galaxies. Gas that is subject to gravitational torques, such as that in a tidal bridge in the nuclear region, will likely fall into the nuclear regions by the time of final coalescence (Souchay et al., 2013). Therefore, the molecular gas in the central region will likely fall into the nuclear regions of NGC 6240 and contribute to the second starburst that will turn NGC 6240 into a bona fide ULIRG. However, the high velocity dispersion with respect to the vertical extent of the gas means it is also possible that the structure will dissipate prior to this streaming. These observations and models shed light on one mechanism for star formation in major gas-rich mergers, that is, small-scale tidal bridges forming between progenitor galaxy nuclei that may ultimately feed into the nuclear regions.

There is highly redshifted gas that is not captured by the LIME model, and has velocities too high to be associated with the tidal bridge. We explore the possible acceleration mechanisms for this highly redshifted molecular gas with $<v>\sim$ 400 km/s when isolated from the more quiescent (v = [-250,200] km/s) gas of the galaxy. The energetics are such that it is difficult to determine the source of the gas' velocity, with both the gravitational energy of the merger and the southern AGN containing enough energy to accelerate the gas. The gas' velocity and morphology agrees closely with models and observations of the molecular outflow presented in Cicone et al. (2018) as well as the ionized outflow in Müller-Sánchez et al. (2018), meaning it may well be driven by the AGN outflow. Its velocity, though extreme, does not exceed the escape velocity of the nuclear region. Therefore it will likely fall back onto the nuclear region and fuel future star formation.

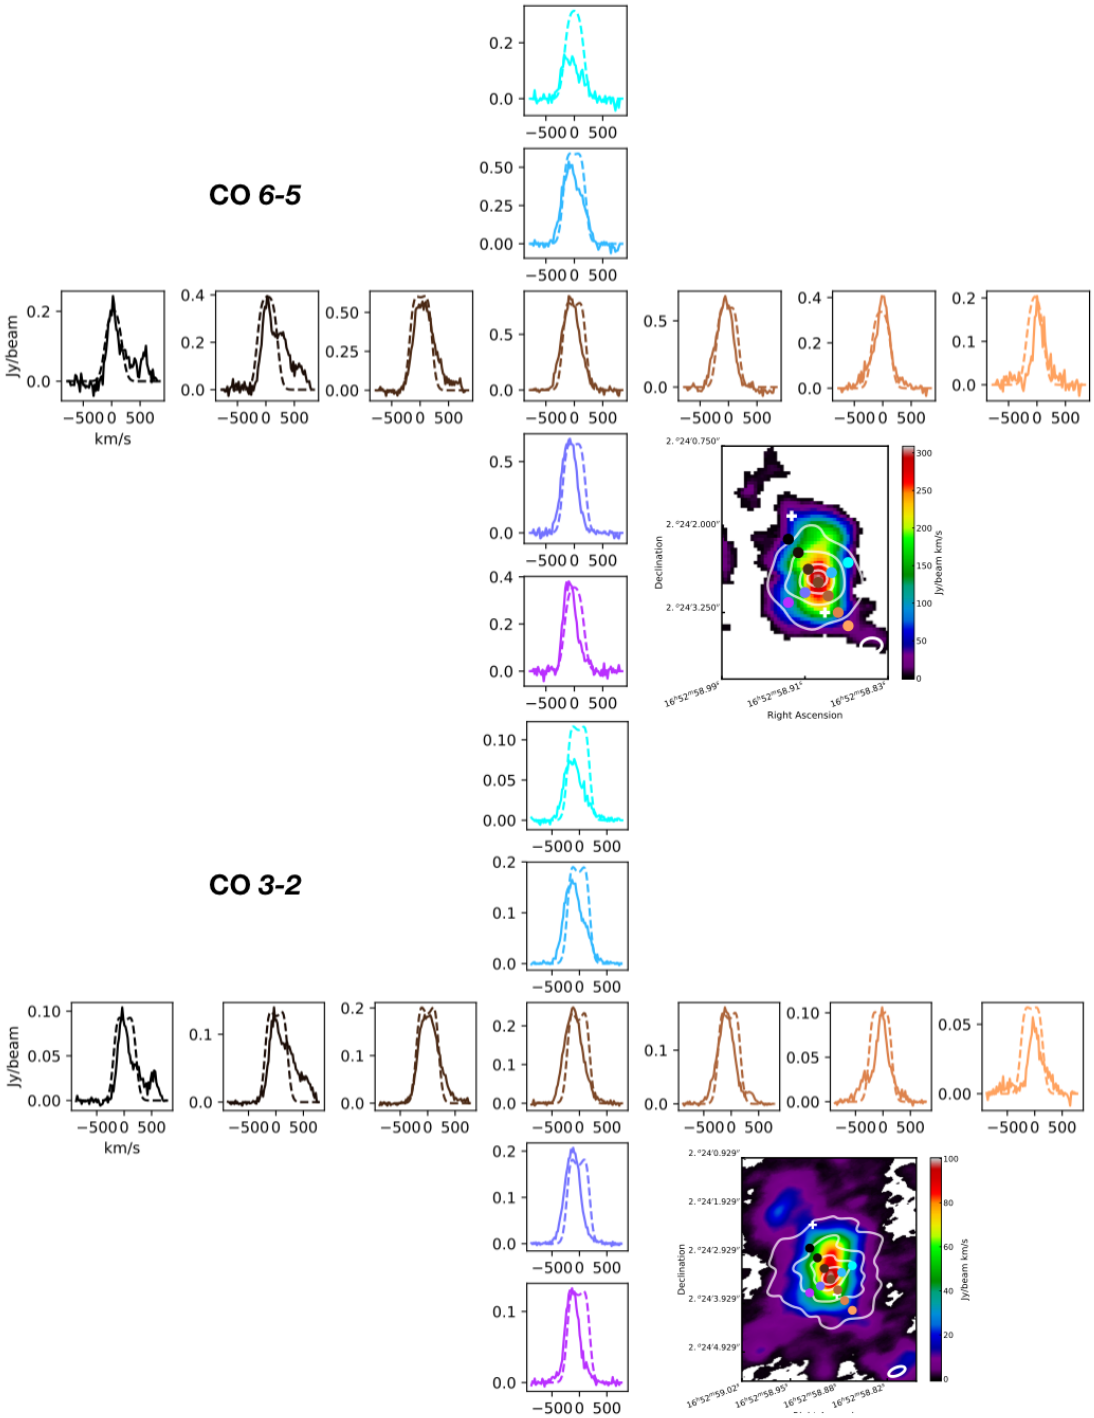

Figure 6.9: The line profiles for the fiducial fit CO J = 6 – 5 (**top**) and J = 3 – 2 (**bottom**) LIME models compared to the observed CO J = 6 – 5 and J = 3 – 2 line profiles extracted along -34°, the angle between the maximum of the highly blueshifted and redshifted gas. *Dashed Lines:* model line profiles at the location color-coded in the inset image. *Solid Lines:* Observed line profiles at the location color-coded in the inset image. **Inset Image:** The color map is the moment 0 map for the corresponding observation, and contours show the model's integrated line emission.

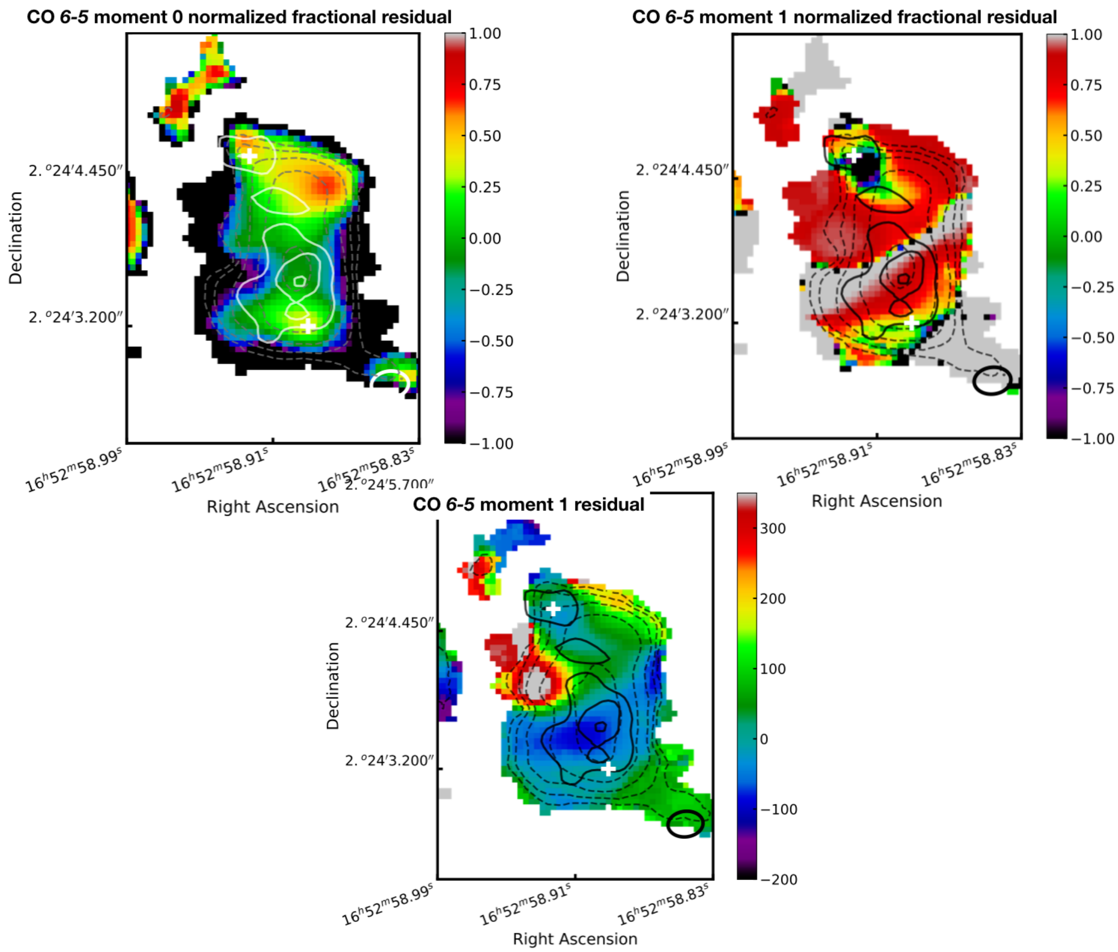

Figure 6.10: Normalized fractional residuals (*color*) between CO J = 6 − 5 data and simulation for moment 0 (*upper left*) and moment 1 (*upper right*). The *lower panel* shows the non-normalized residual in moment 1. *Dashed contours* correspond to the moment 0 data, *solid contours* correspond to the 677 GHz continuum emission, and *crosses* show the locations of the two AGN. Beam FWHM contours are shown in the bottom right of each panel.

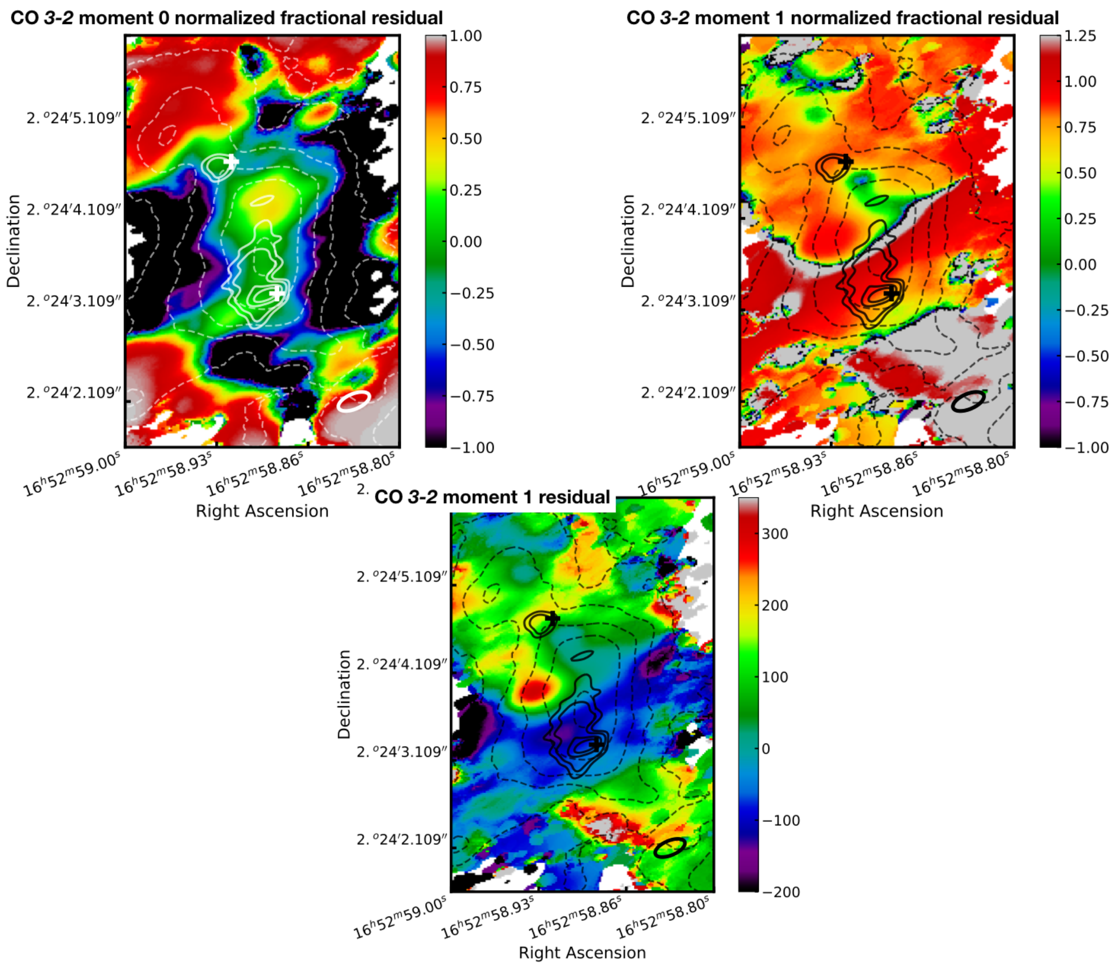

Figure 6.11: Normalized fractional residuals (*color*) between CO J = 3 − 2 data and simulation for moment 0 (*upper left*) and moment 1 (*upper right*). The *lower panel* shows the non-normalized residual in moment 1. *Dashed contours* correspond to the moment 0 data, *solid contours* correspond to the 344.4 GHz continuum emission, and *crosses* show the locations of the two AGN. Beam FWHM contours are shown in the bottom right of each panel.

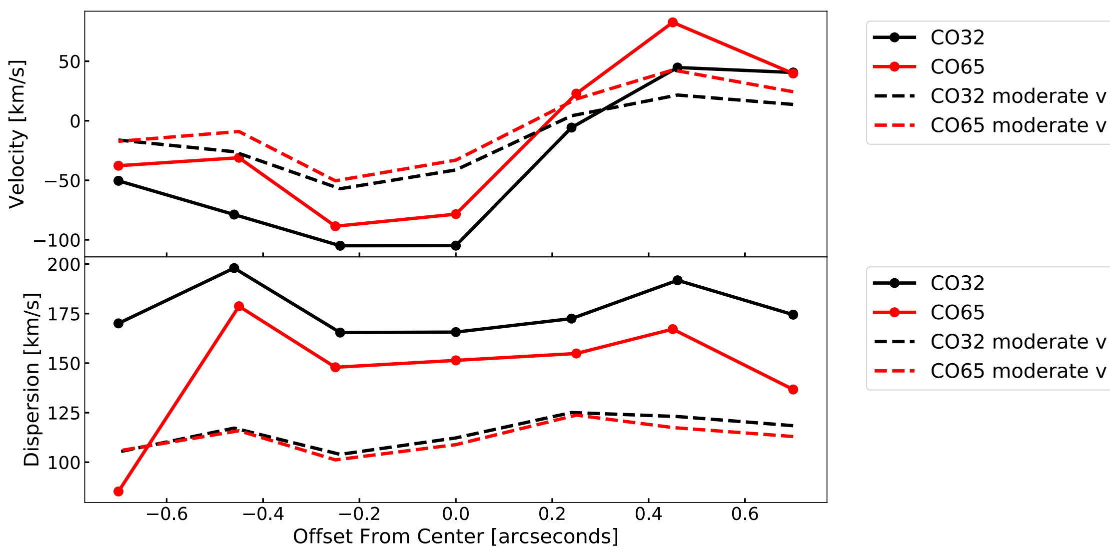

Figure 6.12: The average velocity and dispersion of CO 6-5 and CO 3-2 extracted along the semi-major axis of the central gas region (0^o), for all gas velocities (*solid*) and only moderate velocity [-200, 250] km/s gas (*dashed*).

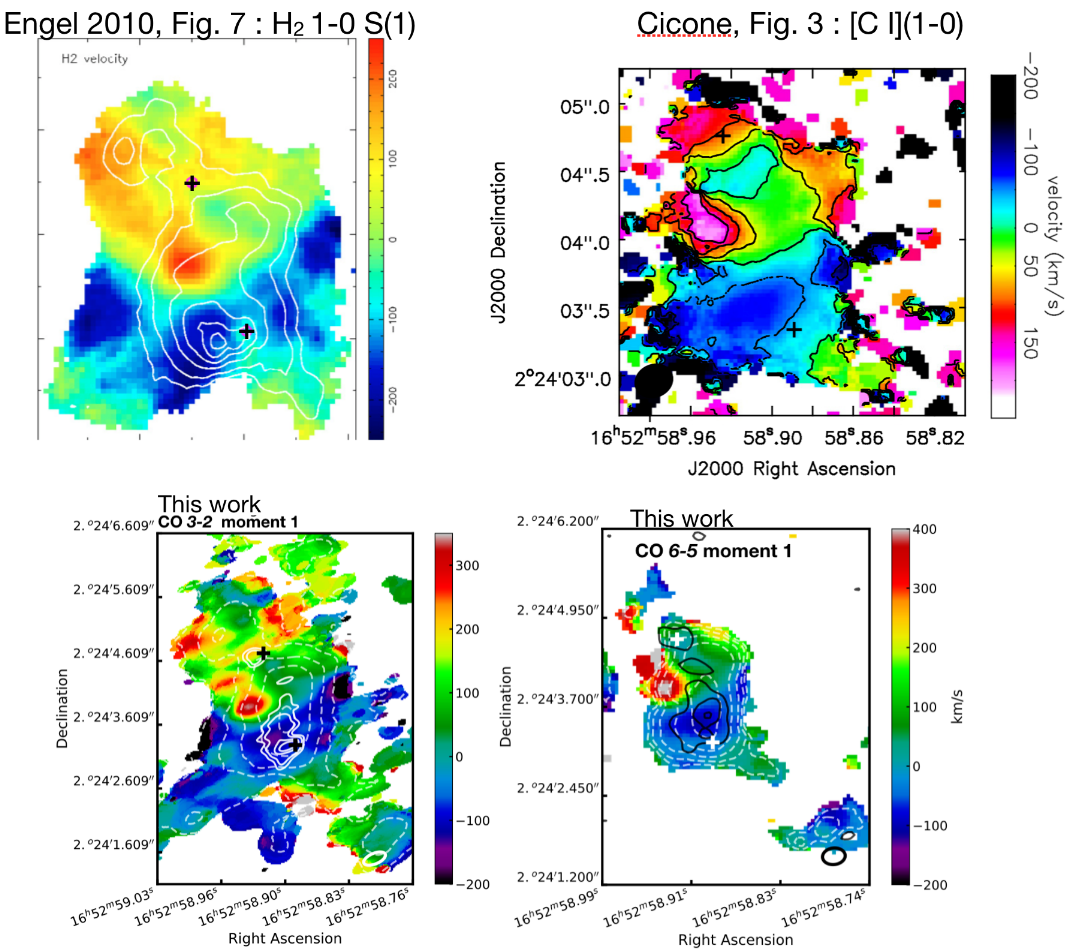

Figure 6.13: A comparison of the molecular gas observations from the outflow papers referenced (Müller-Sánchez et al. (2018), Cicone et al. (2018)) and this work, showing morphological and kinematic correspondence. **Upper left**: H_2 $(1-0)$ $S(1)$ velocity map from Engel et al. (2010) and reproduced in Müller-Sánchez et al. (2018). Müller-Sánchez et al. (2018) claims that this gas is a perturbed rotating disk produced by the tidal forces of the merger based on the arguments of Tacconi et al. (1999). **Upper right**: [CI]$(1-0)$ velocity map from Cicone et al. (2018). They claim that this gas shows the inner initiation region of an extended molecular outflow (Figure 6.14, upper right). **Bottom**: CO J = $3-2$ and J = $6-5$ velocity maps from this work. The crosses show the location of the nuclei in each map.

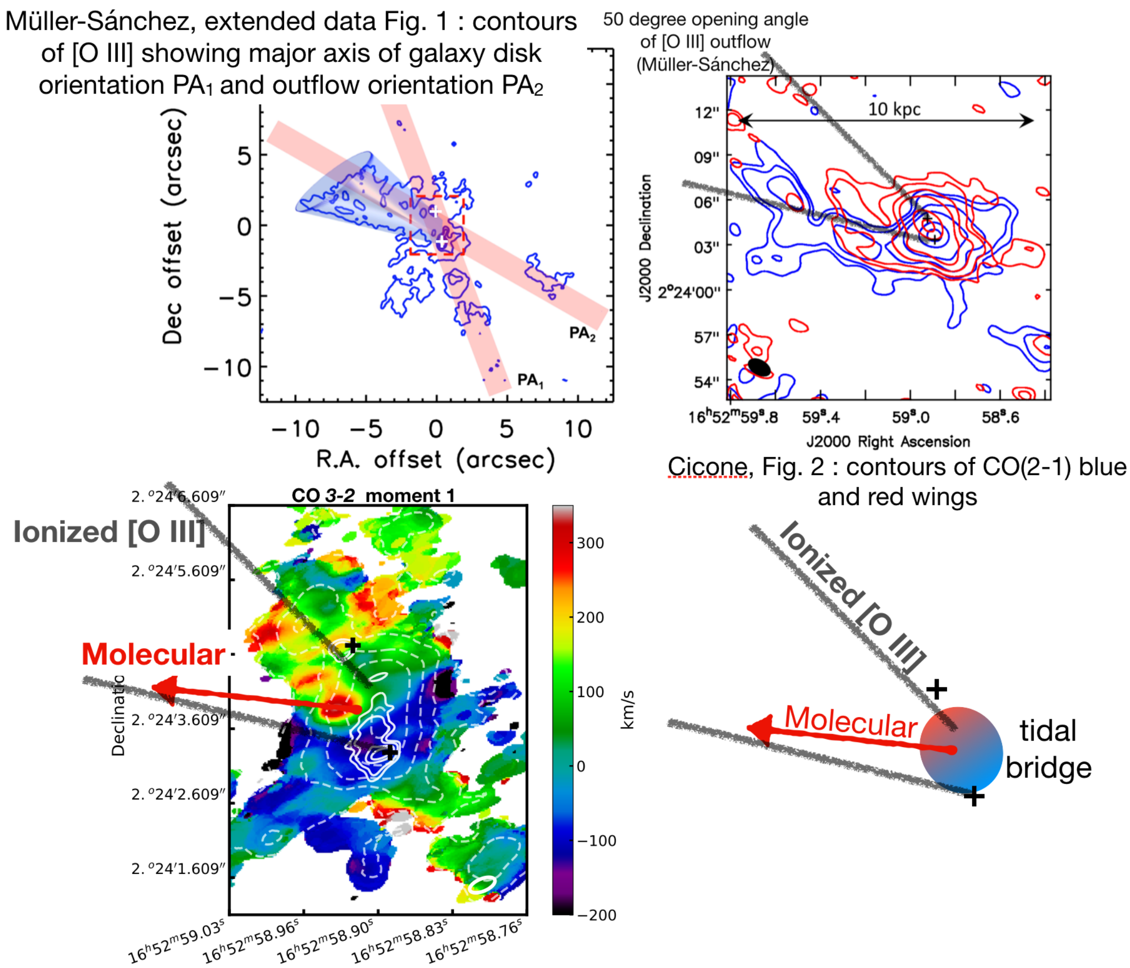

Figure 6.14: Cartoons and observations showing the overlap between the outflows of Müller-Sánchez et al. (2018) and Cicone et al. (2018) as well as the tidal bridge proposed in this observation. **Upper left**: the cartoon of the ionized [O III] outflow from Müller-Sánchez et al. (2018) (represented by the blue cone, though the gas is redshifted), their extended data Figure 1. Blue contours show the [O III] emission, white crosses the approximate location of the AGN. PA_1 is the position angle of the galactic disk. PA_2 is the position angle of the outflow and is the angle along which the velocities are extracted in Figure 6.15. **Upper right**: the contours of the CO$(2-1)$ v = [650, 200] km/s (blue wing) and v = [250, 800] km/s (red wing) showing the molecular outflow from Cicone et al. (2018), their Figure 2. The gray lines are a cartoon we added overlaying the orientation and angular extent of the ionized outflow modeled in Müller-Sánchez et al. (2018). **Lower left**: the angular extend of the ionized outflow (Müller-Sánchez et al., 2018) and the approximate axis of the redshifted portion of the molecular outflow (Cicone et al., 2018) drawn as cartoons on top of our CO J = $3-2$ velocity map. **Bottom right**: the cartoon from the lower left including the tidal bridge suggested by our LIME models.

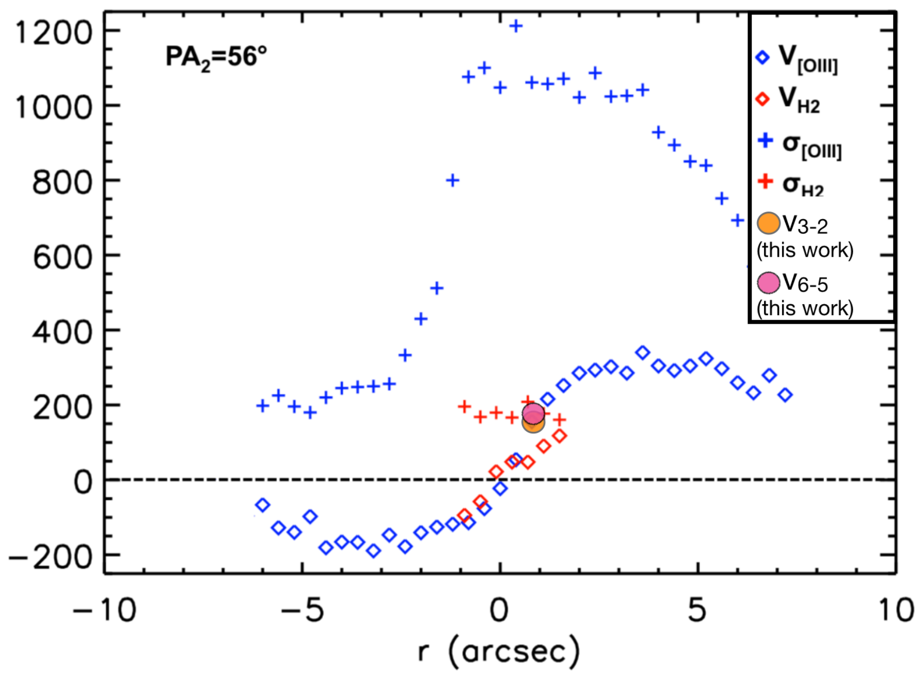

Müller-Sánchez, Fig. 3d : velocities of [O III] and H2 extracted along outflow axis

Figure 6.15: The velocity and its standard deviation of [O III] and H_2 $(1-0)$ $S(1)$ extracted along the position angle of the outflow (PA_2 in our Figure 6.14) from Müller-Sánchez et al. (2018). The H_2 is extracted from the map in the upper left of Figure 6.13. The velocity of CO J = $3-2$ from this work is plotted as an orange circle and that of CO J = $6-5$ is a pink circle. Both show close correspondence to the observations of [O III] from their work (which they associate with an outflow), but with velocities notably higher than the H_2 observation (which they associate with a disk).

Chapter 7

Future Work

7.1 Detectors

7.1.1 Low Volume Detectors (Chapter 3)

The low volume detector tests left our team with a number of questions that were unanswerable at the time of test, and remained unanswered for the duration of this thesis work. Given future funding, there are a few steps that can be taken to further characterize the noise and the physics that define the detector properties:

(1) The anomalously high T_c of 1.9 K shortens the quasiparticle lifetime and consequently lowers the responsivity. If bilayer devices are desired, finding the cause of the elevated T_c is critical. A sheet of TiN/Al bilayer with the Al on the top surface was fabricated in 2015 and its T_c was found to be lower than the bilayer with TiN on the top surface. This points to a possibility that depositing the TiN atop the Al during the fabrication process is altering the Al layer and heightening the T_c. A test of TiN/Al/TiN trilayer KIDs, with Al comprising the bulk of the detector volume, should be conducted.

(2) The source of the 1/f noise (possibly TLS noise or variable stray light absorbed by the detectors) must be identified as it is the major limiter at signal modulation frequencies $<$ 10 Hz.

(3) Our measurements for τ_{qp} at this time were found by fitting the noise rolloff in the NEP

spectra. An alternative measurement of τ_{qp} should be made by observing the response of the array when flashed by an IR LED.

(4) The optical coupling to the detectors by the lenslet should be better characterized with a simulation of the lenslet. This would help to constrain the conversion from blackbody temperature to optical power and increase the accuracy of our measurements of responsivity.

(5) Implementing a backshort would help to improve optical coupling to the detectors, which would in turn improve the responsivity at all optical loads.

To test the influence of stray light on the measurements taken previously, the Al/TiN bilayer devices needed to be measured in the CU testbed under dark conditions. In late 2017, nearly two years after the tests of the Al/TiN bilayer device were completed in the JPL testbed, we had completed the CU testbed to a point where we could test devices. We integrated the array into our testbed with dual intentions: primarily, to confirm the functionality of our testbed and secondarily, to evaluate differences between our testbed and that at JPL (e.g. a lower stray-light environment).

Initial tests of the arrays allowed our group to de-bug a number of issues extant in the system and are of no import to the scope of this thesis. Later tests of the devices included IR LED measurements of τ_{qp} in the same setup as the tests in Chapter 5. These measurements confirmed the short timeconstants ($\sim 100\ \mu$s) measured from the Lorentzian rolloffs of the noise PSDs measured in the JPL cryostat. However, these initial tests lacked the requisite wirebonds to properly thermally sink and electrically ground the array. Therefore, measurements of the quality factor and noise completed on this device could not be used for direct comparison with the JPL testbed and should be repeated.

7.1.2 Phonon Recycling KIDs (Chapter 4)

The back-thinned devices described in Chapter 4 were fabricated prior to simulations of the phonon recycled arrays began. The low simulated recycling factor of 1.49 for this geometry requires extremely careful testing to measure a 50% difference in time constants. The new testbed at CU would likely facilitate these careful tests, and we recommend measuring τ_{qp}/τ_{exp} using the IR LED fiber optic setup described in Chapter 5 combined with a light diffuser. The measurements for this array were also completed years before our group's understanding of the intricacies of the quasi-exponential decays described in Chapter 5. As such, another improvement on this test would be to fit the decays with X(t) rather than a single exponential. This would help to more carefully quantify the differences between back-thinned and standard devices.

To prove the efficacy of phonon recycling for elongating τ_{qp}, we also recommend further fabrications of phonon recycling arrays. Three possibilities of fabrication exist:

(1) The simulations show high recycling factors for low perimeter coverage fractions, with $P_{rf} \sim 20$ for a coverage fraction of 10^{-2}. Among the different parameters that can be changed (p, t_{thin}, and R_L), the perimeter coverage fraction gives the most τ_{qp} boosting within reasonable fabrication parameter spaces. Therefore, we recommend fabricating a standard KID array comprised of half standard detectors and half with p of 10^{-2}. The simulated boosting of 20 times τ_{qp} should be quite obvious in this instance and demonstrate the efficacy of this effect.

(2) Another simple fabrication test is to re-fabricate the array with half back-thinned and half standard inductors, but with a thinner membrane in order to boost the difference in τ_{qp} between the devices. The thinnest SOI wafers on the market have a top layer thickness of 100 nm (waferpro.com). Given an R_L of 90 μm, this corresponds to a t_{thin} / R_L of 10^{-3}. Our simulations did not extend to this extremely low t_{thin}, but the dependency was stronger than an exponential and suggest τ_{qp} increases of over ten times.

(3) The most complex re-fabrication would be of the fully released inductor array. We do not recommend this approach after the two unsuccessful fabrication attempts of this complex design without simplifications to the design. In the event that this is desired, we also recommend including more detectors on the array, with half of the detectors remaining

non-released for a direct comparison of phonon recycled and standard devices.

7.1.3 Time Constant Measurements of Al CPW Devices (Chapter 5)

There are three tests that should be completed in order to finalize our understanding of the dependency of τ_{qp} on metal film thickness:

(1) direct T_c measurement of all wafers (a four-wire resistance measurement of the 20, 30, and 40 nm thick arrays is needed to corroborate our findings.)

(2) $\tau_{qp}(T_{det})$ for the 50 nm thick device at lower power where time constants are no longer influenced by microwave-generated quasiparticles.

(3) We also fabricated a 100 nm thick CPW device, for which testing was not completed at low enough powers to have included in this work. τ_{qp} should be measured at a wide range of driving powers as well as for T_{det} from 80 mK to 300 mK. T_c should be measured directly for this device using a four-wire resistance measurement as well as from $\Delta f/f_0(T_{det})$.

There is also some work that could be done using the simulated decays to further characterize the influence of both the noise and our fitting technique on the recovered τ_{qp}. There is a plan in place to hopefully complete this work by the due date of the accompanying paper's referee response, November 14. Currently the simulated decays are matched to the measured data by eye, iterating on values of $\tau_{qp,simulated}$, $x_{qp}(0)_{simulated}$, and $B_{simulated}$ as well as the noise levels until both the PSD and the decay of the simulated dataset appear to match the measured data. This is by its nature a qualitative process that needs to be done for every decay separately. We cannot repeat this process for every decay (there are around 100 and each takes about 4 hours to do the full analysis). As such, a quantitative, computer-based process should be developed if further characterization of the influence of the decay's noise and the fitting technique on the measured τ_{qp} is desired.

One way to potentially mimic the decays quantitatively in an automated way is to fit the measured decay with X(t) (Equation 5.9, Chapter 5), use the fit values for $\tau_{qp,simulated}$, $x_{qp}(0)_{simulated}$, and $B_{simulated}$, then use the Fourier Transform of the measured decay to derive noise parameters.

7.1.4 Short-Wavelength devices and progress towards the GEP and GEP-B

In order to facilitate future FIR observatories like the GEP and GEP-B (see Introduction), KIDs will need to be able to observe at wavelengths as short as 10 μm. The devices presented in this thesis are designed to be sensitive to light at 350 μm. Peter Day, Joanna Perido, and Jason Glenn are leading the development of detectors sensitive to this wavelength range, briefly introduced in Chapter 1, Section 1.4.2. Optical testing of these devices is necessary to test the efficacy of the inductor design at absorbing 10 μm light. Neither the testbed at CU nor JPL possesses a method to deliver 10 μm light isolated from other wavelengths to the detectors. To enable this testing, the testbeds will require 10 μm bandpass filters. Joanna Perido is simulating metal mesh filters that, when combined with commercial filters, will allow isolation of 10 μm light and the optical testing of these short-wavelength devices. Once this work is completed, iterations on the design of the 10 μm absorber can be completed to optimize the absorbers for the wavelength ranges needed for the GEP and GEP-B.

7.2 NGC 6240

There is a wealth of observations of NGC 6240, and almost every molecule of interest to this work has been observed at some transition at some resolution. However, there is always more we can learn, and we suggest future observations of NGC 6240 below.

Understanding the variations in the dust spectral index β in the nuclear region of NGC 6240 could be informative on understanding the influence of shocks and outflows on dust. High-resolution observations of the nuclear region in CO J lines that have not yet been observed, e.g. $4-3$, $5-4$, $7-6$, and the continuum bands at these frequencies would allow the constraint of the dust spectral index. β and the dust temperature are degenerate, meaning these observations would also allow us to constrain the dust temperature in high-resolution.

Dust is destroyed in both continuous (C-type) and dissociative (J-type) shocks (Guillet et al. (2009), Guillet et al. (2011)). Even though only a few percent of the dust grains are destroyed, the shock process changes the distribution of the dust grain sizes significantly (J-type figure 7 in Guillet et al. (2009), C-type figure 4 in Guillet et al. (2011)). The result is a distribution of dust grain sizes that is more uniform, with many more smaller dust grains in a shocked region as compared to in an unshocked region. This would have the effect of decreasing the dust emissivity β as compared to the surroundings. Similarly, higher values of β ($\gtrsim$2) are associated with larger, icier dust grains Lis et al. (1998). The presence of AGN outflows impacting quiescent gas could therefore produce lower β's in the influenced gas, and potentially higher β's on the far side of the influenced gas that is more shielded from the AGN radiation. This geometry is seen in clouds studied in the Orion Nebula, with pockets of $\beta \simeq$2.5 on the sides of clouds shielded from the high-energy radiation from the Trapezium cluster Lis et al. (1998).

We discussed in Chapter 6 the possible causes of the highly redshifted gas observed in the nuclear region. We could not distinguish between three possibilities that we proposed: gravitational acceleration of molecular gas, gas accelerated by an AGN outflow, or the existence of a dwarf galaxy at that location.

Gravitational acceleration might be difficult to observationally disentangle from gas accelerated by an AGN outflow, as both would cause velocities high enough to form shocks at this location. An AGN outflow would cause a wider effect than a gravitationally slingshotted cloud of gas, and would likely be biconical. Cicone et al. (2018) saw extended, biconical signatures while Müller-Sánchez et al. (2018) saw extended, single-cone features. Our observations do not extend to large enough angular scales to observe this structure. We therefore recommend a re-observation of the compact array observations in CO J = $6-5$ and $3-2$ that were failed for poor sensitivity to look for the connection between this nuclear high-velocity gas and the extended structure around this nuclear region. This is similar to the analysis of Cicone et al. (2018), but the differences would be illuminating. Their extended observations were in CO J = $1-0$ and $2-1$ and their nuclear high-resolution observations are in [CI]($1-0$), all of which trace *cool* molecular gas. We wish to observe the extended configuration of CO J = $6-5$ especially to see how the hot molecular gas corresponds to the hot ionized outflow modeled in Müller-Sánchez et al. (2018).

Bibliography

P. A. R. Ade, R. W. Aikin, M. Amiri, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, J. A. Bonetti, J. A. Brevik, I. Buder, E. Bullock, G. Chattopadhyay, G. Davis, P. K. Day, C. D. Dowell, L. Duband, J. P. Filippini, S. Fliescher, S. R. Golwala, M. Halpern, M. Hasselfield, S. R. Hildebrandt, G. C. Hilton, V. Hristov, H. Hui, K. D. Irwin, W. C. Jones, K. S. Karkare, J. P. Kaufman, B. G. Keating, S. Kefeli, S. A. Kernasovskiy, J. M. Kovac, C. L. Kuo, H. G. LeDuc, E. M. Leitch, N. Llombart, M. Lueker, P. Mason, K. Megerian, L. Moncelsi, C. B. Netterfield, H. T. Nguyen, R. O'Brient, R. W. O. IV, A. Orlando, C. Pryke, A. S. Rahlin, C. D. Reintsema, S. Richter, M. C. Runyan, R. Schwarz, C. D. Sheehy, Z. K. Staniszewski, R. V. Sudiwala, G. P. Teply, J. E. Tolan, A. Trangsrud, R. S. Tucker, A. D. Turner, A. G. Vieregg, A. Weber, D. V. Wiebe, P. Wilson, C. L. Wong, K. W. Yoon, and J. Z. and. ANTENNA-COUPLED TES BOLOMETERS USED IN BICEP2,keck array, AND SPIDER. The Astrophysical Journal, 812(2):176, oct 2015. doi: 10.1088/0004-637x/812/2/176. URL `https://doi.org/10.1088%2F0004-637x%2F812%2F2%2F176`.

A. Alonso-Herrero. AGN and star formation activity in local luminous and ultraluminous infrared galaxies. arXiv e-prints, art. arXiv:1302.2033, Feb 2013.

A. Alonso-Herrero, M. Pereira-Santaella, G. H. Rieke, and D. Rigopoulou. LOCAL LUMINOUS INFRARED GALAXIES. II. ACTIVE GALACTIC NUCLEUS ACTIVITY FROMSPITZER/INFRARED SPECTROGRAPH SPECTRA. The Astrophysical Journal, 744 (1):2, dec 2011. doi: 10.1088/0004-637x/744/1/2. URL `https://doi.org/10.1088%2F0004-637x%2F744%2F1%2F2`.

L. Armus, J. Bernard-Salas, H. W. W. Spoon, J. A. Marshall, V. Charmandaris, S. J. U. Higdon, V. Desai, L. Hao, H. I. Teplitz, D. Devost, B. R. Brand l, B. T. Soifer, and J. R. Houck. Detection of the Buried Active Galactic Nucleus in NGC 6240 with the Infrared Spectrograph on the Spitzer Space Telescope. ApJ, 640(1):204–210, Mar 2006. doi: 10.1086/500040.

L. Armus, J. Bauer, D. Benford, E. Bergin, A. Bolatto, C. Bradford, C. Chen, A. Cooray, N. Evans, D. Farrah, J. Glenn, P. Goldsmith, A. Harris, G. Helou, D. Leisawitz, D. Lis, P. Marcum, G. Melnick, S. Milam, L. Mundy, D. Neufeld, K. Pontoppidan, A. Pope, M. Rizzo, K. Sandstrom, K. Sheth, and E. Wright. The Science Case and Concept for a Far-Infrared Surveyor. In The Science Case and Concept for a Far-Infrared Surveyor, 2015.

R. Barends, J. J. A. Baselmans, S. J. C. Yates, J. R. Gao, J. N. Hovenier, and T. M. Klapwijk. Quasiparticle relaxation in optically excited high-q superconducting resonators. Phys. Rev. Lett., 100:257002, Jun 2008. doi: 10.1103/PhysRevLett.100.257002. URL https://link.aps.org/doi/10.1103/PhysRevLett.100.257002.

J. E. Barnes. Formation of gas discs in merging galaxies. MNRAS, 333(3):481–494, Jul 2002. doi: 10.1046/j.1365-8711.2002.05335.x.

J. E. Barnes and L. Hernquist. Transformations of Galaxies. II. Gasdynamics in Merging Disk Galaxies. ApJ, 471:115, Nov. 1996. doi: 10.1086/177957.

J. E. Barnes and L. E. Hernquist. Fueling starburst galaxies with gas-rich mergers. ApJ, 370: L65–L68, Apr. 1991. doi: 10.1086/185978.

J. Baselmans, S. Yates, P. Diener, and P. de Visser. Ultra low background cryogenic test facility for far-infrared radiation detectors. Journal of Low Temperature Physics, 167(3):360–366, May 2012. ISSN 1573-7357. doi: 10.1007/s10909-012-0511-0. URL https://doi.org/10.1007/s10909-012-0511-0.

J. J. A. Baselmans and S. J. C. Yates. Long quasiparticle lifetime in Aluminum Microwave Kinetic Inductance Detectors using coaxial stray light filters. In B. Young, B. Cabrera, and A. Miller, editors, American Institute of Physics Conference Series, volume 1185 of American Institute of Physics Conference Series, pages 160–163, Dec. 2009. doi: 10.1063/1.3292305.

Baselmans, J. J. A., Bueno, J., Yates, S. J. C., Yurduseven, O., Llombart, N., Karatsu, K., Baryshev, A. M., Ferrari, L., Endo, A., Thoen, D. J., de Visser, P. J., Janssen, R. M. J., Murugesan, V., Driessen, E. F. C., Coiffard, G., Martin-Pintado, J., Hargrave, P., and Griffin, M. A kilo-pixel imaging system for future space based far-infrared observatories using microwave kinetic inductance detectors. A&A, 601:A89, 2017. doi: 10.1051/0004-6361/201629653. URL https://doi.org/10.1051/0004-6361/201629653.

R. J. Beswick, A. Pedlar, C. G. Mundell, and J. F. Gallimore. A MERLIN neutral hydrogen absorption study of the luminous infrared merger NGC 6240. MNRAS, 325:151–158, July 2001. doi: 10.1046/j.1365-8711.2001.04445.x.

J. L. Black and B. I. Halperin. Spectral diffusion, phonon echoes, and saturation recovery in glasses at low temperatures. Phys. Rev. B, 16:2879–2895, Sep 1977. doi: 10.1103/PhysRevB.16.2879. URL https://link.aps.org/doi/10.1103/PhysRevB.16.2879.

A. W. Blain, I. Smail, R. J. Ivison, J. P. Kneib, and D. T. Frayer. Submillimeter galaxies. Phys. Rep., 369(2):111–176, Oct 2002. doi: 10.1016/S0370-1573(02)00134-5.

R. W. Boyd. Photon bunching and the photon-noise-limited performance of infrared detectors. Infrared Physics, 22(3):157 – 162, 1982. ISSN 0020-0891. doi: https://doi.org/10.1016/0020-0891(82)90034-3. URL http://www.sciencedirect.com/science/article/pii/0020089182900343.

J. Braine, P. A. Duc, U. Lisenfeld, V. Charmand aris, O. Vallejo, S. Leon, and E. Brinks. Abundant molecular gas in tidal dwarf galaxies: On-going galaxy formation. A&A, 378:51–69, Oct 2001. doi: 10.1051/0004-6361:20011109.

J. Braine, E. Davoust, M. Zhu, U. Lisenfeld, C. Motch, and E. R. Seaquist. A molecular gas bridge between the Taffy galaxies. A&A, 408:L13–L16, Sept. 2003. doi: 10.1051/0004-6361:20031148.

T. G. Brainerd and M. Yamamoto. Satellite galaxies in the Illustris-1 simulation: anisotropic locations around relatively isolated hosts. MNRAS, 489:459–469, Oct. 2019. doi: 10.1093/mnras/stz2102.

C. Brinch and M. R. Hogerheijde. LIME - a flexible, non-LTE line excitation and radiation transfer method for millimeter and far-infrared wavelengths. A&A, 523:A25, Nov. 2010. doi: 10.1051/0004-6361/201015333.

J. Bueno, M. D. Shaw, P. K. Day, and P. M. Echternach. Proof of concept of the quantum capacitance detector. Applied Physics Letters, 96(10):103503, 2010. doi: 10.1063/1.3339163. URL `https://doi.org/10.1063/1.3339163`.

J. Bueno, V. Murugesan, K. Karatsu, D. J. Thoen, and J. J. A. Baselmans. Ultrasensitive kilo-pixel imaging array of photon noise-limited kinetic inductance detectors over an octave of bandwidth for thz astronomy. Journal of Low Temperature Physics, 193(3):96–102, Nov 2018. ISSN 1573-7357. doi: 10.1007/s10909-018-1962-8. URL `https://doi.org/10.1007/s10909-018-1962-8`.

B. L. Cardozo. Gaas blocked-impurity-band detectors for far-infrared astronomy. doi: 10.2172/840239.

C. M. Casey. Far-infrared spectral energy distribution fitting for galaxies near and far. MNRAS, 425:3094–3103, Oct. 2012. doi: 10.1111/j.1365-2966.2012.21455.x.

C. M. Casey, D. Narayanan, and A. Cooray. Dusty star-forming galaxies at high redshift. Physics Reports, 541(2):45 – 161, 2014. ISSN 0370-1573. doi: https://doi.org/10.1016/j.physrep.2014.02.009. URL `http://www.sciencedirect.com/science/article/pii/S0370157314000477`. Dusty star-forming galaxies at high-redshift.

A. Catalano, P. Ade, Y. Atik, A. Benoit, E. Bréele, J. J. Bock, P. Camus, M. Chabot, M. Charra, B. P. Crill, N. Coron, A. Coulais, F. X. Désert, L. Fauvet, Y. Giraud-Héraud, O. Guillaudin, W. Holmes, W. C. Jones, J. M. Lamarre, J. Macías-Pérez, M. Martinez, A. Miniussi, A. Monfardini, F. Pajot, G. Patanchon, A. Pelissier, M. Piat, J. L. Puget, C. Renault, C. Rosset, D. Santos, A. Sauvé, L. D. Spencer, and R. Sudiwala. Impact of particles on the Planck HFI detectors: Ground-based measurements and physical interpretation. A&A, 569:A88, Sep 2014. doi: 10.1051/0004-6361/201423868.

P. Chubov, V. Eremenko, and Y. Pilipenko. Dependence of the critical temperature and energy gap of superconducting aluminum films on their thickness. Zh. Eksp. Teor. Fiz., 55: 752-65(1968)., 1 1968.

C. Cicone, P. Severgnini, P. P. Papadopoulos, R. Maiolino, C. Feruglio, E. Treister, G. C. Privon, Z. yu Zhang, R. D. Ceca, F. Fiore, K. Schawinski, and J. Wagg. ALMA [c i]3 p 13 p 0 observations of NGC 6240: A puzzling molecular outflow, and the role of outflows in the global α CO factor of (u)LIRGs. The Astrophysical Journal, 863(2):143, aug 2018. doi: 10.3847/1538-4357/aad32a. URL https://doi.org/10.3847%2F1538-4357%2Faad32a.

F. Combes. Molecular gas in distant galaxies from ALMA studies. A&A Rev., 26(1):5, Aug 2018. doi: 10.1007/s00159-018-0110-4.

C. J. Conselice. The evolution of galaxy structure over cosmic time. Annual Review of Astronomy and Astrophysics, 52(1):291–337, 2014. doi: 10.1146/annurev-astro-081913-040037. URL https://doi.org/10.1146/annurev-astro-081913-040037.

S. Croft, W. van Breugel, W. de Vries, M. Dopita, C. Martin, R. Morganti, S. Neff, T. Oosterloo, D. Schiminovich, S. A. Stanford, and J. van Gorkom. Minkowski's Object: A Starburst Triggered by a Radio Jet, Revisited. ApJ, 647(2):1040–1055, Aug 2006. doi: 10.1086/505526.

P. Day, H. Leduc, B. Mazin, A. Vayonakis, and J. Zmuidzinas. A broadband superconducting detector suitable for use in large arrays. Nature, 425:817–21, 10 2003. doi: 10.1038/nature02037.

T. de Haan, A. Suzuki, S. T. P. Boyd, R. H. Cantor, A. Coerver, M. A. Dobbs, R. Hennings-Yeomans, W. L. Holzapfel, A. T. Lee, G. I. Noble, G. Smecher, and J. Zhou. Recent Advances in Frequency-Multiplexed TES Readout: Vastly Reduced Parasitics and an Increase in Multiplexing Factor with sub-Kelvin SQUIDs. arXiv e-prints, art. arXiv:1908.07642, Aug 2019.

P. J. de Visser, J. J. A. Baselmans, P. Diener, S. J. C. Yates, A. Endo, and T. M. Klapwijk. Number fluctuations of sparse quasiparticles in a superconductor. Phys. Rev. Lett., 106:167004, Apr 2011. doi: 10.1103/PhysRevLett.106.167004. URL https://link.aps.org/doi/10.1103/PhysRevLett.106.167004.

P. J. de Visser, J. J. A. Baselmans, S. J. C. Yates, P. Diener, A. Endo, and T. M. Klapwijk. Microwave-induced excess quasiparticles in superconducting resonators measured through correlated conductivity fluctuations. Applied Physics Letters, 100(16):162601, Apr 2012. doi: 10.1063/1.4704151.

P. J. de Visser, J. J. A. Baselmans, J. Bueno, N. Llombart, and T. M. Klapwijk. Fluctuations in the electron system of a superconductor exposed to a photon flux. Nature Communications, 5:3130, Feb 2014. doi: 10.1038/ncomms4130.

P. Diener, H. G. Leduc, S. J. C. Yates, Y. J. Y. Lankwarden, and J. J. A. Baselmans. Design and Testing of Kinetic Inductance Detectors Made of Titanium Nitride. Journal of Low Temperature Physics, 167(3-4):305–310, May 2012. doi: 10.1007/s10909-012-0484-z.

P. M. Echternach, K. J. Stone, C. M. Bradford, P. K. Day, D. W. Wilson, K. G. Megerian, N. Llombart, and J. Bueno. Photon shot noise limited detection of terahertz radiation using a quantum capacitance detector. Applied Physics Letters, 103(5):053510, 2013. doi: 10.1063/1.4817585. URL https://doi.org/10.1063/1.4817585.

H. Engel, R. I. Davies, R. Genzel, L. J. Tacconi, E. K. S. Hicks, E. Sturm, T. Naab, P. H. Johansson, S. J. Karl, C. E. Max, A. Medling, and P. P. van der Werf. NGC 6240: merger-induced star formation and gas dynamics. A&A, 524:A56, Dec. 2010. doi: 10.1051/0004-6361/201015338.

D. Farrah, K. E. Smith, D. Ardila, C. M. Bradford, M. J. DiPirro, C. Ferkinhoff, J. Glenn, P. F. Goldsmith, D. T. Leisawitz, T. Nikola, N. Rangwala, S. A. Rinehart, J. G. Staguhn, M. Zemcov, J. Zmuidzinas, J. Bartlett, S. J. Carey, W. J. Fischer, J. R. Kamenetzky, J. Kartaltepe, M. D. Lacy, D. C. Lis, L. S. Locke, E. Lopez-Rodriguez, M. MacGregor, E. Mills, S. H. Moseley, E. J. Murphy, A. Rhodes, M. J. Richter, D. Rigopoulou, D. B. Sanders, R. Sankrit, G. Savini, J.-D. Smith, and S. Stierwalt. Review: far-infrared instrumentation and technological development for the next decade. Journal of Astronomical Telescopes, Instruments, and Systems, 5(2):1 – 34, 2019. doi: 10.1117/1.JATIS.5.2.020901. URL `https://doi.org/10.1117/1.JATIS.5.2.020901`.

C. Feruglio, F. Fiore, R. Maiolino, E. Piconcelli, H. Aussel, D. Elbaz, E. Le Floc'h, E. Sturm, R. Davies, and C. Cicone. NGC 6240: extended CO structures and their association with shocked gas. A&A, 549:A51, Jan 2013. doi: 10.1051/0004-6361/201219746.

D. Flanigan. PhD thesis, 2018.

J. W. Fried and H. Schulz. NGC 6240 - A unique interacting galaxy. A&A, 118:166–170, Feb. 1983.

A. Fyhrie, J. Zmuidzinas, J. Glenn, P. Day, H. G. LeDuc, and C. McKenney. Progress towards ultra sensitive KIDs for future far-infrared missions: a focus on recombination times. In Proc. SPIE, volume 10708 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, page 107083A, Jul 2018. doi: 10.1117/12.2312867.

O. L. Fvre, C. Sanjuan, L. Tasca, Vvds Team, Massiv Team, and Vuds Team. The Galaxy Merger Rate History (GMRH) since z 3. In W.-H. Sun, C. K. Xu, N. Z. Scoville, and D. B. Sanders, editors, Galaxy Mergers in an Evolving Universe, volume 477 of Astronomical Society of the Pacific Conference Series, page 133, Oct. 2013.

J. F. Gallimore and R. Beswick. Parsec-Scale Radio Structure of the Double Active Nucleus of NGC 6240. AJ, 127(1):239–251, Jan 2004. doi: 10.1086/379959.

J. Gao, J. Zmuidzinas, B. A. Mazin, H. G. LeDuc, and P. K. Day. Noise properties of superconducting coplanar waveguide microwave resonators. Applied Physics Letters, 90(10):102507, Mar 2007. doi: 10.1063/1.2711770.

R. Genzel and C. J. Cesarsky. Extragalactic results from the infrared space observatory. Annual Review of Astronomy and Astrophysics, 38(1):761–814, 2000. doi: 10.1146/annurev.astro.38.1.761. URL `https://doi.org/10.1146/annurev.astro.38.1.761`.

R. Genzel, D. Lutz, E. Sturm, E. Egami, D. Kunze, A. F. M. Moorwood, D. Rigopoulou, H. W. W. Spoon, A. Sternberg, L. E. Tacconi-Garman, L. Tacconi, and N. Thatte. What Powers Ultraluminous IRAS Galaxies? ApJ, 498:579–605, May 1998. doi: 10.1086/305576.

R. Gereth and K. Hubner. Phonon mean free path in silicon between 77 and 250°k. Phys. Rev., 134:A235–A240, Apr 1964. doi: 10.1103/PhysRev.134.A235. URL `https://link.aps.org/doi/10.1103/PhysRev.134.A235`.

J. Gerssen, R. P. van der Marel, D. Axon, J. C. Mihos, L. Hernquist, and J. E. Barnes. Hubble space TelescopeObservations of NGC 6240: A case study of an ultraluminous infrared galaxy with obscured activity. The Astronomical Journal, 127(1):75–89, jan 2004. doi: 10.1086/380223. URL `https://doi.org/10.1086%2F380223`.

E. Giovannoli, V. Buat, S. Noll, D. Burgarella, and B. Magnelli. Population synthesis modelling of luminous infrared galaxies at intermediate redshift. A&A, 525:A150, Jan 2011. doi: 10.1051/0004-6361/201014898.

J. Glenn, C. M. Bradford, R. Amini, B. Moore, A. Benson, L. Armus, K. Alatalo, J. Darling, P. Day, J. Domber, D. Farrah, A. Fyhrie, B. Hensley, S. Lipscy, D. Redding, M. Rogers, M. Shannon, J. Steeves, C. Tucker, G. Wu, and J. Zmuidzinas. The Galaxy Evolution Probe: a concept for a mid and far-infrared space observatory. In Proc. SPIE, volume 10698 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, page 106980L, Jul 2018. doi: 10.1117/12.2314076.

J. Glenn, K. Alatalo, R. Amini, L. Armus, J. Bally, A. Benson, C. M. Bradford, J. Comerford, J. Darling, P. Day, J. Domber, D. Farrah, A. Fyhrie, M. Shannon, B. Hensley, S. Lipscy, B. Moore, E. Nelson, S. Oliver, B. Oppenheimer, D. Redding, M. Rodgers, E. Rosolowsky, R. Shirley, H. Smith, J. Steeves, A. Tielens, C. Tucker, G. Wu, and J. Zmuidzinas. The Galaxy Evolution Probe. In BAAS, volume 51, page 112, Sep 2019.

J. R. Goicoechea, M. Etxaluze, J. Cernicharo, M. Gerin, D. A. Neufeld, A. Contursi, T. A. Bell, M. De Luca, P. Encrenaz, N. Indriolo, D. C. Lis, E. T. Polehampton, and P. Sonnentrucker. Herschel Far-infrared Spectroscopy of the Galactic Center. Hot Molecular Gas: Shocks versus Radiation near Sgr A*. ApJ, 769(1):L13, May 2013. doi: 10.1088/2041-8205/769/1/L13.

D. J. Goldie and S. Withington. Non-Equilibrium Superconductivity in Kinetic Inductance Detectors for THz Photon Sensing. arXiv e-prints, art. arXiv:1401.2291, Jan 2014.

V. Guillet, A. P. Jones, and G. Pineau Des Forêts. Shocks in dense clouds. II. Dust destruction and SiO formation in J shocks. A&A, 497(1):145–153, Apr 2009. doi: 10.1051/0004-6361/200811115.

V. Guillet, G. Pineau Des Forêts, and A. P. Jones. Shocks in dense clouds. III. Dust processing and feedback effects in C-type shocks. A&A, 527:A123, Mar 2011. doi: 10.1051/0004-6361/201015973.

Y. Hagiwara, W. A. Baan, and H.-R. Klöckner. Very Long Baseline Interferometry Observations of NGC 6240: Resolving the Double Nuclei and Radio Supernovae. AJ, 142(1):17, Jul 2011. doi: 10.1088/0004-6256/142/1/17.

S. Hailey-Dunsheath, A. C. M. Barlis, J. E. Aguirre, C. M. Bradford, J. G. Redford, T. S. Billings, H. G. LeDuc, C. M. McKenney, and M. I. Hollister. Development of Aluminum LEKIDs for Balloon-Borne Far-IR Spectroscopy. Journal of Low Temperature Physics, 193(5-6):968–975, Dec 2018. doi: 10.1007/s10909-018-1927-y.

T. Hashimoto, A. K. Inoue, K. Mawatari, Y. Tamura, H. Matsuo, H. Furusawa, Y. Harikane, T. Shibuya, K. K. Knudsen, K. Kohno, Y. Ono, E. Zackrisson, T. Okamoto, N. Kashikawa, P. A. Oesch, M. Ouchi, K. Ota, I. Shimizu, Y. Taniguchi, H. Umehata, and D. Watson. Big Three Dragons: A z = 7.15 Lyman-break galaxy detected in [O III] 88 μm, [C II] 158 μm, and dust continuum with ALMA. PASJ, 71(4):71, Aug 2019. doi: 10.1093/pasj/psz049.

S. W. Henderson, J. R. Stevens, M. Amiri, J. Austermann, J. A. Beall, S. Chaudhuri, H.-M. Cho, S. K. Choi, N. F. Cothard, K. T. Crowley, S. M. Duff, C. P. Fitzgerald, P. A. Gallardo, M. Halpern, M. Hasselfield, G. Hilton, S.-P. P. Ho, J. Hubmayr, K. D. Irwin, B. J. Koopman, D. Li, Y. Li, J. McMahon, F. Nati, M. Niemack, C. D. Reintsema, M. Salatino, A. Schillaci, B. L. Schmitt, S. M. Simon, S. T. Staggs, E. M. Vavagiakis, and J. T. Ward. Readout of two-kilopixel transition-edge sensor arrays for Advanced ACTPol. In W. S. Holland and J. Zmuidzinas, editors, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII, volume 9914, pages 346 – 362. International Society for Optics and Photonics, SPIE, 2016. doi: 10.1117/12.2233895. URL `https://doi.org/10.1117/12.2233895`.

T. Herter. Ibc arrays: Present and future prospects. Experimental Astronomy, 3(1):255–263, Mar 1994. ISSN 1572-9508. doi: 10.1007/BF00430173. URL `https://doi.org/10.1007/BF00430173`.

J. E. Hibbard and J. C. Mihos. Dynamical Modeling of NGC 7252 and the Return of Tidal Material. AJ, 110:140, July 1995. doi: 10.1086/117502.

P. F. Hopkins, T. J. Cox, L. Hernquist, D. Narayanan, C. C. Hayward, and N. Murray. Star formation in galaxy mergers with realistic models of stellar feedback and the interstellar medium. MNRAS, 430(3):1901–1927, Apr 2013. doi: 10.1093/mnras/stt017.

J. Hubmayr, J. E. Austermann, J. A. Beall, D. T. Becker, S. J. Benton, A. S. Bergman, J. R. Bond, S. Bryan, S. M. Duff, A. J. Duivenvoorden, H. K. Eriksen, J. P. Filippini, A. Fraisse, M. Galloway, A. E. Gambrel, K. Ganga, A. L. Grigorian, R. Gualtieri, J. E. Gudmundsson, J. W. Hartley, M. Halpern, G. C. Hilton, W. C. Jones, J. J. McMahon, L. Moncelsi, J. M. Nagy, C. B. Netterfield, B. Osherson, I. Padilla, A. S. Rahlin, B. Racine, J. Ruhl, T. M. Rudd, J. A. Shariff, J. D. Soler, X. Song, J. N. Ullom, J. V. Lanen, M. R. Vissers, I. K. Wehus, S. Wen, D. V. Wiebe, and E. Young. Design of 280 GHz feedhorn-coupled TES arrays for the balloon-borne polarimeter SPIDER. In W. S. Holland and J. Zmuidzinas, editors, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII, volume 9914, pages 185 – 198. International Society for Optics and Photonics, SPIE, 2016. doi: 10.1117/12.2231896. URL `https://doi.org/10.1117/12.2231896`.

D. Iono, C. D. Wilson, S. Takakuwa, M. S. Yun, G. R. Petitpas, A. B. Peck, P. T. P. Ho, S. Matsushita, Y. M. Pihlstrom, and Z. Wang. High-Resolution Imaging of Warm and Dense Molecular

Gas in the Nuclear Region of the Luminous Infrared Galaxy NGC 6240. ApJ, 659(1):283–295, Apr 2007. doi: 10.1086/512362.

A. James, L. Dunne, S. Eales, and M. G. Edmunds. SCUBA observations of galaxies with metallicity measurements: a new method for determining the relation between submillimetre luminosity and dust mass. MNRAS, 335:753–761, Sept. 2002. doi: 10.1046/j.1365-8711.2002.05660.x.

R. M. J. Janssen, J. J. A. Baselmans, A. Endo, L. Ferrari, S. J. C. Yates, A. M. Baryshev, and T. M. Klapwijk. Performance of hybrid NbTiN-Al microwave kinetic inductance detectors as direct detectors for sub-millimeter astronomy. In Proc. SPIE, volume 9153 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, page 91530T, Jul 2014. doi: 10.1117/12.2055537.

J. Kamenetzky, N. Rangwala, J. Glenn, P. R. Maloney, and A. Conley. A Survey of the Molecular ISM Properties of Nearby Galaxies Using the Herschel FTS. ApJ, 795:174, Nov. 2014. doi: 10.1088/0004-637X/795/2/174.

S. B. Kaplan, C. C. Chi, D. N. Langenberg, J. J. Chang, S. Jafarey, and D. J. Scalapino. Quasiparticle and phonon lifetimes in superconductors. Phys. Rev. B, 14:4854–4873, Dec 1976. doi: 10.1103/PhysRevB.14.4854. URL `https://link.aps.org/doi/10.1103/PhysRevB.14.4854`.

M. Kenyon, P. K. Day, C. M. Bradford, J. J. Bock, and H. G. Leduc. Ultrasensitive transitionedge sensors (tess) for farir/submm spaceborne spectroscopy. AIP Conference Proceedings, 1185(1):56–59, 2009. doi: 10.1063/1.3292406. URL `https://aip.scitation.org/doi/abs/10.1063/1.3292406`.

U. Klaas, M. Haas, S. A. H. Müller, R. Chini, B. Schulz, I. Coulson, H. Hippelein, K. Wilke, M. Albrecht, and D. Lemke. Infrared to millimetre photometry of ultra-luminous IR galaxies: New evidence favouring a 3-stage dust model. A&A, 379:823–844, Dec 2001. doi: 10.1051/0004-6361:20011377.

S. Komossa, V. Burwitz, G. Hasinger, P. Predehl, J. S. Kaastra, and Y. Ikebe. Discovery of a Binary Active Galactic Nucleus in the Ultraluminous Infrared Galaxy NGC 6240 Using Chandra. ApJ, 582:L15–L19, Jan. 2003. doi: 10.1086/346145.

K. Kukita and Y. Kamakura. Monte carlo simulation of phonon transport in silicon including a realistic dispersion relation. Journal of Applied Physics, 114(15):154312, 2013. doi: 10.1063/1.4826367. URL `https://doi.org/10.1063/1.4826367`.

J. H. Lacy, C. Sneden, H. Kim, and D. T. Jaffe. H2, CO, and dust absorption through cold molecular clouds. The Astrophysical Journal, 838(1):66, mar 2017. doi: 10.3847/1538-4357/aa6247. URL `https://doi.org/10.3847%2F1538-4357%2Faa6247`.

S. J. Lilly, O. L. Fèvre, F. Hammer, and D. Crampton. The canada-france redshift survey: The luminosity density and star formation history of the universe to [ITAL]z[/ITAL] $\sim$ 1. The Astrophysical Journal, 460(1), mar 1996. doi: 10.1086/309975. URL `https://doi.org/10.1086%2F309975`.

P. Lira, M. J. Ward, A. Zezas, and S. S. Murray. Chandra HRC and HST observations of NGC 6240: resolving the active galactic nucleus and starburst. MNRAS, 333:709–714, July 2002. doi: 10.1046/j.1365-8711.2002.05456.x.

D. C. Lis, E. Serabyn, J. Keene, C. D. Dowell, D. J. Benford, T. G. Phillips, T. R. Hunter, and N. Wang. 350 micron continuum imaging of the orion a molecular cloud with the submillimeter high angular resolution camera. The Astrophysical Journal, 509(1):299–308, dec 1998. doi: 10.1086/306500. URL `https://doi.org/10.1086%2F306500`.

E. K. Lofthouse, S. Kaviraj, C. J. Conselice, A. Mortlock, and W. Hartley. Major mergers are not significant drivers of star formation or morphological transformation around the epoch of peak cosmic star formation. Monthly Notices of the Royal Astronomical Society, 465(3):2895–2900, 11 2016. ISSN 0035-8711. doi: 10.1093/mnras/stw2895. URL `https://doi.org/10.1093/mnras/stw2895`.

C. Lonsdale, D. Farrah, and H. Smith. Ultraluminous Infrared Galaxies, volume 2, pages 285–336. 08 2006. doi: 10.1007/3-540-30313-8_9.

J. M. Lotz, P. Jonsson, T. J. Cox, and J. R. Primack. Galaxy merger morphologies and time-scales from simulations of equal-mass gas-rich disc mergers. MNRAS, 391(3):1137–1162, Dec 2008. doi: 10.1111/j.1365-2966.2008.14004.x.

P. J. Love, A. W. Hoffman, N. A. Lum, K. J. Ando, J. Rosbeck, W. D. Ritchie, N. J. Therrien, R. S. Holcombe, and E. Corrales. 1024 x 1024 Si:As IBC detector arrays for JWST MIRI. In T. J. Grycewicz and C. J. Marshall, editors, Focal Plane Arrays for Space Telescopes II, volume 5902, pages 58 – 66. International Society for Optics and Photonics, SPIE, 2005. doi: 10.1117/12.623473. URL `https://doi.org/10.1117/12.623473`.

D. Lutz. Far-infrared surveys of galaxy evolution. Annual Review of Astronomy and Astrophysics, 52(1):373–414, 2014. doi: 10.1146/annurev-astro-081913-035953. URL `https://doi.org/10.1146/annurev-astro-081913-035953`.

D. D. E. Martin and P. Verhoeve. Superconducting tunnel junctions. ISSI Scientific Reports Series, 9:441–457, 2010.

D. Marx and W. Eisenmenger. Phonon scattering at siliconcrystal surfaces. Zeitschrift für Physik B Condensed Matter, 48(4):277–291, Dec 1982. ISSN 1431-584X. doi: 10.1007/BF01305187. URL `https://doi.org/10.1007/BF01305187`.

D. C. Mattis and J. Bardeen. Theory of the anomalous skin effect in normal and superconducting metals. Phys. Rev., 111:412–417, Jul 1958. doi: 10.1103/PhysRev.111.412. URL `https://link.aps.org/doi/10.1103/PhysRev.111.412`.

C. E. Max, G. Canalizo, B. A. Macintosh, L. Raschke, D. Whysong, R. Antonucci, and G. Schneider. The core of NGC 6240 from keck adaptive optics andHubble space TelescopeNICMOS observations. The Astrophysical Journal, 621(2):738–749, mar 2005. doi: 10.1086/427546. URL `https://doi.org/10.1086%2F427546`.

R. Meijerink, L. E. Kristensen, A. Weiß, P. P. van der Werf, F. Walter, M. Spaans, A. F. Loenen, J. Fischer, F. P. Israel, K. Isaak, P. P. Papadopoulos, S. Aalto, L. Armus, V. Charmandaris, K. M. Dasyra, T. Diaz-Santos, A. Evans, Y. Gao, E. González-Alfonso, R. Güsten, C. Henkel, C. Kramer, S. Lord, J. Martín-Pintado, D. Naylor, D. B. Sanders, H. Smith, L. Spinoglio, G. Stacey, S. Veilleux, and M. C. Wiedner. Evidence for CO Shock Excitation in NGC 6240 from Herschel SPIRE Spectroscopy. ApJ, 762:L16, Jan. 2013. doi: 10.1088/2041-8205/762/2/L16.

R. Morganti. The many routes to AGN feedback. Frontiers in Astronomy and Space Sciences, 4: 42, Nov 2017. doi: 10.3389/fspas.2017.00042.

R. Morganti, J. Holt, L. Saripalli, T. A. Oosterloo, and C. N. Tadhunter. IC 5063: AGN driven outflow of warm and cold gas. A&A, 476(2):735–743, Dec 2007. doi: 10.1051/0004-6361:20077888.

R. Morganti, T. Oosterloo, J. B. R. Oonk, W. Frieswijk, and C. Tadhunter. The fast molecular outflow in the Seyfert galaxy IC 5063 as seen by ALMA. A&A, 580:A1, Aug 2015. doi: 10.1051/0004-6361/201525860.

F. Müller-Sánchez, R. Nevin, J. Comerford, R. Davies, G. Privon, and E. Treister. Two separate outflows in the dual supermassive black hole system ngc 6240. Nature, 556, 04 2018. doi: 10.1038/s41586-018-0033-2.

H. Murakami, H. Baba, P. Barthel, D. L. Clements, M. Cohen, Y. Doi, K. Enya, E. Figueredo, N. Fujishiro, H. Fujiwara, M. Fujiwara, P. Garcia-Lario, T. Goto, S. Hasegawa, Y. Hibi, T. Hirao, N. Hiromoto, S. S. Hong, K. Imai, M. Ishigaki, M. Ishiguro, D. Ishihara, Y. Ita, W.-S. Jeong, K. S. Jeong, H. Kaneda, H. Kataza, M. Kawada, T. Kawai, A. Kawamura, M. F. Kessler, D. Kester, T. Kii, D. C. Kim, W. Kim, H. Kobayashi, B. C. Koo, S. M. Kwon, H. M. Lee, R. Lorente, S. Makiuti, H. Matsuhara, T. Matsumoto, H. Matsuo, S. Matsuura, T. G. Mller, N. Murakami, H. Nagata, T. Nakagawa, T. Naoi, M. Narita, M. Noda, S. H. Oh, A. Ohnishi, Y. Ohyama, Y. Okada, H. Okuda, S. Oliver, T. Onaka, T. Ootsubo, S. Oyabu, S. Pak, Y.-S. Park, C. P. Pearson, M. Rowan-Robinson, T. Saito, I. Sakon, A. Salama, S. Sato, R. S. Savage, S. Serjeant, H. Shibai, M. Shirahata, J. Sohn, T. Suzuki, T. Takagi, H. Takahashi, T. Tanab, T. T. Takeuchi, S. Takita, M. Thomson, K. Uemizu, M. Ueno, F. Usui, E. Verdugo, T. Wada, L. Wang, T. Watabe, H. Watarai, G. J. White, I. Yamamura, C. Yamauchi, and A. Yasuda. The Infrared Astronomical Mission AKARI*. Publications of the Astronomical Society of Japan, 59(sp2):S369–S376, 10 2007. ISSN 0004-6264. doi: 10.1093/pasj/59.sp2.S369. URL `https://doi.org/10.1093/pasj/59.sp2.S369`.

K. Nakanishi, S. K. Okumura, K. Kohno, R. Kawabe, and T. Nakagawa. Dense and Warm Molecular Gas between Double Nuclei of the Luminous Infrared Galaxy NGC 6240. PASJ, 57:575–586, Aug 2005. doi: 10.1093/pasj/57.4.575.

O. Noroozian, J. Gao, J. Zmuidzinas, H. G. LeDuc, and B. A. Mazin. Two-level system noise reduction for Microwave Kinetic Inductance Detectors. In B. Young, B. Cabrera, and A. Miller, editors, American Institute of Physics Conference Series, volume 1185 of American Institute of Physics Conference Series, pages 148–151, Dec 2009. doi: 10.1063/1.3292302.

Y. Ohyama, M. Yoshida, and T. Takata. Superwind-driven Intense H_2 Emission in NGC 6240. II. Detailed Comparison of Kinematic and Morphological Structures of the Warm and Cold Molecular Gas. AJ, 126(5):2291–2298, Nov 2003. doi: 10.1086/378709.

H.-A. Pan, L. Lin, B.-C. Hsieh, T. Xiao, Y. Gao, S. L. Ellison, J. M. Scudder, J. Barrera-Ballesteros, F. Yuan, A. Saintonge, C. D. Wilson, H. S. Hwang, I. D. Looze, Y. Gao, L. C. Ho, E. Brinks, A. Mok, T. Brown, T. A. Davis, T. G. Williams, A. Chung, H. Parsons, M. Bureau, M. T. Sargent, E. J. Chung, E. Kim, T. Liu, M. J. Michałowski, and T. Tosaki. The effect of galaxy interactions on molecular gas properties. The Astrophysical Journal, 868(2):132, dec 2018. doi: 10.3847/1538-4357/aaeb92. URL https://doi.org/10.3847%2F1538-4357%2Faaeb92.

A. Paskin and A. D. Singh. Boundary Conditions and Quantum Effects in Thin Superconducting Films. Physical Review, 140:1965–1967, Dec. 1965. doi: 10.1103/PhysRev.140.A1965.

Pilbratt, G. L., Riedinger, J. R., Passvogel, T., Crone, G., Doyle, D., Gageur, U., Heras, A. M., Jewell, C., Metcalfe, L., Ott, S., and Schmidt, M. Herschel space observatory* - an esa facility for far-infrared and submillimetre astronomy. A&A, 518:L1, 2010. doi: 10.1051/0004-6361/201014759. URL https://doi.org/10.1051/0004-6361/201014759.

L. K. Pollack, C. E. Max, and G. Schneider. Circumnuclear Star Clusters in the Galaxy Merger NGC 6240, Observed with Keck Adaptive Optics and the Hubble Space Telescope. ApJ, 660 (1):288–300, May 2007. doi: 10.1086/512777.

J.-L. Puget, A. Abergel, J.-P. Bernard, F. Boulanger, W. B. Burton, F.-X. Desert, and D. Hartmann. Tentative detection of a cosmic far-infrared background with COBE. A&A, 308:L5, Apr. 1996.

G. Rodighiero, E. Daddi, I. Baronchelli, A. Cimatti, A. Renzini, H. Aussel, P. Popesso, D. Lutz, P. Andreani, S. Berta, A. Cava, D. Elbaz, A. Feltre, A. Fontana, N. M. F. Schreiber, A. Franceschini, R. Genzel, A. Grazian, C. Gruppioni, O. Ilbert, E. L. Floch, G. Magdis, M. Magliocchetti, B. Magnelli, R. Maiolino, H. McCracken, R. Nordon, A. Poglitsch, P. Santini, F. Pozzi, L. Riguccini, L. J. Tacconi, S. Wuyts, and G. Zamorani. THE LESSER ROLE OF STARBURSTS IN STAR FORMATION ATz= 2. The Astrophysical Journal, 739(2):L40, sep 2011. doi: 10.1088/2041-8205/739/2/l40. URL https://doi.org/10.1088%2F2041-8205%2F739%2F2%2Fl40.

W. Rujopakarn, G. H. Rieke, D. J. Eisenstein, and S. Juneau. MORPHOLOGY AND SIZE DIFFERENCES BETWEEN LOCAL AND HIGH-REDSHIFT LUMINOUS INFRARED GALAXIES. The Astrophysical Journal, 726(2):93, dec 2010. doi: 10.1088/0004-637x/726/2/93. URL https://doi.org/10.1088%2F0004-637x%2F726%2F2%2F93.

H. R. Russell, B. R. McNamara, A. C. Fabian, P. E. J. Nulsen, F. Combes, A. C. Edge, M. T. Hogan, M. McDonald, P. Salomé, G. Tremblay, and A. N. Vantyghem. Close entrainment of massive molecular gas flows by radio bubbles in the central galaxy of Abell 1795. MNRAS, 472:4024–4037, Dec. 2017. doi: 10.1093/mnras/stx2255.

D. B. Sanders and I. F. Mirabel. Luminous Infrared Galaxies. ARA&A, 34:749, 1996. doi: 10.1146/annurev.astro.34.1.749.

D. B. Sanders, B. T. Soifer, J. H. Elias, B. F. Madore, K. Matthews, G. Neugebauer, and N. Z. Scoville. Ultraluminous infrared galaxies and the origin of quasars. ApJ, 325:74–91, Feb. 1988. doi: 10.1086/165983.

D. Schiminovich, J. H. van Gorkom, J. M. van der Hulst, and S. Kasow. Discovery of Neutral Hydrogen Associated with the Diffuse Shells of NGC 5128 (Centaurus A). ApJ, 423:L101, Mar. 1994. doi: 10.1086/187246.

N. Scoville, K. Sheth, F. Walter, S. Manohar, L. Zschaechner, M. Yun, J. Koda, D. Sanders, L. Murchikova, T. Thompson, B. Robertson, R. Genzel, L. Hernquist, L. Tacconi, R. Brown, D. Narayanan, C. C. Hayward, J. Barnes, J. Kartaltepe, R. Davies, P. van der Werf, and E. Fomalont. ALMA Imaging of HCN, CS, and Dust in Arp 220 and NGC 6240. ApJ, 800:70, Feb. 2015. doi: 10.1088/0004-637X/800/1/70.

E. A. Shapoval. Critical Temperature of Small Superconductors. ZhETF Pisma Redaktsiiu, 5:57, Jan. 1967.

M. D. Shaw, J. Bueno, P. Day, C. M. Bradford, and P. M. Echternach. Quantum capacitance detector: A pair-breaking radiation detector based on the single Cooper-pair box. Phys. Rev. B, 79(14):144511, Apr 2009. doi: 10.1103/PhysRevB.79.144511.

Y. L. Shirley. The critical density and the effective excitation density of commonly observed molecular dense gas tracers. Publications of the Astronomical Society of the Pacific, 127(949): 299–310, mar 2015. doi: 10.1086/680342. URL `https://doi.org/10.1086%2F680342`.

B. T. Soifer, G. Helou, and M. Werner. The spitzer view of the extragalactic universe. Annual Review of Astronomy and Astrophysics, 46(1):201–240, 2008. doi: 10.1146/annurev.astro.46.060407.145144. URL `https://doi.org/10.1146/annurev.astro.46.060407.145144`.

P. M. Solomon and P. A. Vanden Bout. Molecular Gas at High Redshift. ARA&A, 43:677–725, Sept. 2005. doi: 10.1146/annurev.astro.43.051804.102221.

J. Souchay, S. Mathis, and T. Tokieda. Tides in Astronomy and Astrophysics, volume 861. 2013. doi: 10.1007/978-3-642-32961-6.

J. P. Stott, D. Sobral, I. Smail, R. Bower, P. N. Best, and J. E. Geach. The merger rates and sizes of galaxies across the peak epoch of star formation from the HiZELS survey. Monthly Notices of the Royal Astronomical Society, 430(2):1158–1170, 01 2013. ISSN 0035-8711. doi: 10.1093/mnras/sts684. URL `https://doi.org/10.1093/mnras/sts684`.

T. Suzuki, P. Khosropanah, M. L. Ridder, R. A. Hijmering, J. R. Gao, H. Akamatsu, L. Gottardi, J. van der Kuur, and B. D. Jackson. Development of ultra-low-noise tes bolometer arrays. Journal of Low Temperature Physics, 184(1):52–59, Jul 2016. ISSN 1573-7357. doi: 10.1007/s10909-015-1401-z. URL `https://doi.org/10.1007/s10909-015-1401-z`.

L. J. Swenson, P. K. Day, B. H. Eom, H. G. Leduc, N. Llombart, C. M. McKenney, O. Noroozian, and J. Zmuidzinas. Operation of a titanium nitride superconducting microresonator detector in the nonlinear regime. Journal of Applied Physics, 113(10):104501-104501-9, Mar 2013. doi: 10.1063/1.4794808.

L. J. Tacconi, R. Genzel, M. Tecza, J. F. Gallimore, D. Downes, and N. Z. Scoville. Gasdynamics in the luminous merger NGC 6240. The Astrophysical Journal, 524(2):732–745, oct 1999. doi: 10.1086/307839. URL https://doi.org/10.1086%2F307839.

L. J. Tacconi, R. Genzel, I. Smail, R. Neri, S. C. Chapman, R. J. Ivison, A. Blain, P. Cox, A. Omont, F. Bertoldi, T. Greve, N. M. Förster Schreiber, S. Genel, D. Lutz, A. M. Swinbank, A. E. Shapley, D. K. Erb, A. Cimatti, E. Daddi, and A. J. Baker. Submillimeter Galaxies at z ˜ 2: Evidence for Major Mergers and Constraints on Lifetimes, IMF, and CO-H_2 Conversion Factor. ApJ, 680: 246–262, June 2008. doi: 10.1086/587168.

M. Tecza, R. Genzel, L. J. Tacconi, S. Anders, L. E. Tacconi-Garman, and N. Thatte. Stellar dynamics and the implications on the merger evolution in NGC 6240. The Astrophysical Journal, 537(1):178–190, jul 2000. doi: 10.1086/309021. URL https://doi.org/10.1086%2F309021.

J. Thronson, Harley A., S. Majewski, L. Descartes, and M. Hereld. A Subdued Interpretation of the Visual and Infrared Emission from Merging Galaxies: Application to NGC 6240. ApJ, 364: 456, Dec 1990. doi: 10.1086/169430.

A. Toomre. On the gravitational stability of a disk of stars. ApJ, 139:1217–1238, May 1964. doi: 10.1086/147861.

P. Vignati, S. Molendi, G. Matt, M. Guainazzi, L. A. Antonelli, L. Bassani, W. N. Brandt, A. C. Fabian, K. Iwasawa, R. Maiolino, G. Malaguti, A. Marconi, and G. C. Perola. BeppoSAX unveils the nuclear component in NGC 6240. A&A, 349:L57–L60, Sep 1999.

J. Wang, E. Nardini, G. Fabbiano, M. Karovska, M. Elvis, S. Pellegrini, C. Max, G. Risaliti, V. U, and A. Zezas. Fast and Furious: Shock Heated Gas as the Origin of Spatially Resolved Hard X-Ray Emission in the Central 5 kpc of the Galaxy Merger NGC 6240. ApJ, 781:55, Jan. 2014. doi: 10.1088/0004-637X/781/1/55.

D. Watson, L. Christensen, K. K. Knudsen, J. Richard, A. Gallazzi, and M. J. Michałowski. A dusty, normal galaxy in the epoch of reionization. Nature, 519(7543):327–330, Mar 2015. doi: 10.1038/nature14164.

J. Wheeler, N. Rangwalla, A. Fyhrie, and J. Glenn. New Observational Insights on the Structure and Kinematics of the Nuclear Molecular Disks and Surrounding Gas. ApJ, 2019 submitted.

R. J. Williams, R. F. Quadri, and M. Franx. THE DIMINISHING IMPORTANCE OF MAJOR GALAXY MERGERS AT HIGHER REDSHIFTS. The Astrophysical Journal, 738(2):L25, aug 2011. doi: 10.1088/2041-8205/738/2/l25. URL https://doi.org/10.1088%2F2041-8205%2F738%2F2%2Fl25.

G. S. Wright, R. D. Joseph, and W. P. S. Meikle. The ultraluminous interacting galaxy NGC6240. Nature, 309(5967):430–431, May 1984. doi: 10.1038/309430a0.

M. Yoshida, M. Yagi, Y. Ohyama, Y. Komiyama, N. Kashikawa, H. Tanaka, and S. Okamura. Giant Hα Nebula Surrounding the Starburst Merger NGC 6240. ApJ, 820:48, Mar. 2016. doi: 10.3847/0004-637X/820/1/48.

T. T. Yuan, L. J. Kewley, and D. B. Sanders. The Role of Starburst-Active Galactic Nucleus Composites in Luminous Infrared Galaxy Mergers: Insights from the New Optical Classification Scheme. ApJ, 709(2):884–911, Feb 2010. doi: 10.1088/0004-637X/709/2/884.

M. S. Yun and J. E. Hibbard. Molecular Gas in Optically Selected Mergers. ApJ, 550(1):104–121, Mar 2001. doi: 10.1086/319697.

A. V. Zasov, A. S. Saburova, O. V. Egorov, and V. L. Afanasiev. Tidal bridge and tidal dwarf candidates in the interacting system Arp194. MNRAS, 462:3419–3429, Nov. 2016. doi: 10.1093/mnras/stw1905.

J. Zmuidzinas. Superconducting microresonators: Physics and applications. Annual Review of Condensed Matter Physics, 3(1):169–214, 2012. doi: 10.1146/annurev-conmatphys-020911-125022. URL `https://doi.org/10.1146/annurev-conmatphys-020911-125022`.