Passive terahertz imaging with lumped element kinetic inductance detectors

Sam Rowe

A thesis submitted for the degree of
Doctor of Philosophy

Defended: 18 December 2015

School of Physics and Astronomy

Cardiff University

Summary of Thesis

Progress towards large format, background limited detector arrays in and around the terahertz or sub-millimetre region of the electromagnetic spectrum has – until recently – been hampered by the complexities in fabrication and cryogenic electronic readout associated with increasing pixel counts. Kinetic inductance detectors or KIDs are a superconducting pair breaking detector technology designed to overcome these complexities.

Traditionally, KID arrays have been developed for imaging in astronomy. However, the high sensitivities, broadband responses, fast time constants and high detector counts that are achievable – along with the simplicity of fabrication and readout compared to other contemporary technologies – make them suitable (and in fact desirable) for a variety of applications.

This thesis documents the development of a concept instrument to demonstrate KID technology for general purpose imaging applications. Specifically, I present the design, construction and performance of a near background limited, quasi-video rate, passive imaging system based on arrays of Aluminium lumped-element KIDs. The camera operates in two atmospheric windows at 150 GHz (2 mm) and 350 GHz (850 μm) with 60 and 152 pixels, respectively. Array fabrication was achieved with a single photolithographic cycle of thin film deposition, patterning and etching. Full array readout is with a single cryogenic amplifier and room temperature FPGA based frequency domain multiplexing electronics.

The camera is the first of its kind in applying KID arrays to imaging systems outside of pure astrophysics research and is the result of efforts from the staff and students of the Astronomy Instrumentation Group (AIG) in the School of Physics and Astronomy with support from QMC Instruments Ltd. The system exemplifies the AIG's world-leading expertise in the development of far-infrared/sub-mm instrumentation as well as QMCI's vision to provide the highest quality terahertz optical components and detector systems to the global marketplace.

Declaration of Authorship

- DECLARATION:

  This work has not previously been accepted in substance for any degree and is not concurrently submitted in candidature for any degree.

  Signed: .... (candidate) Date: ....

- STATEMENT 1:

  This thesis is being submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy (PhD).

  Signed: .... (candidate) Date: ....

- STATEMENT 2:

  This thesis is the result of my own independent work/investigation, except where otherwise stated. Other sources are acknowledged by explicit references. The views expressed are my own.

  Signed: .... (candidate) Date: ....

- STATEMENT 3:

  I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

  Signed: .... (candidate) Date: ....

Contents

List of Figures

List of Tables

List of acronyms

AC alternating current.

ADC analog to digital converter.

ADR adiabatic demagnetisation refrigerator.

AIG Astronomy Instrumentation Group.

API application programming interface.

BCS Bardeen, Cooper, Schrieffer theory.

BP band-pass filter.

CAD computer aided design.

CCD charge-coupled device.

CMB cosmic microwave background radiation.

CMM coordinate measuring machine.

CNC computer numerical control.

CORDIC coordinate rotation digital computer.

CPU central processing unit.

CPW coplanar waveguide.

CW continuous wave.

DAC digital to analog converter.

DAQ data acquisition system.

DC direct current.

DDC digital down converter.

DDS direct digital synthesis.

DFT discrete Fourier transform.

DMA direct memory access.

DSB double side-band.

DSP digital signal processor/processing.

ENOB effective number of bits.

FDM frequency division multiplexing.

FFT fast Fourier transform.

FIFO first in first out.

FIR far infrared.

FOV field of view.

FPA focal plane array.

FPGA field programmable gate array.

FPS frame(s) per second.

FTS Fourier transform spectrometer.

FWHM full width at half maximum.

GPGPU general purpose graphics processing unit.

GPIO general purpose input/output.

GR generation-recombination.

GSPS giga-samples per second.

GUI graphical user interface.

HBT heterojunction bipolar transistor.

HDPE high density polyethylene.

HEB hot-electron bolometer.

HEMT high electron mobility transistor.

HFSS High Frequency Structural Simulator.

HPF high-pass filter.

IDE integrated development environment.

IEEE Institute of Electrical and Electronics Engineers.

IF intermediate frequency.

IP internet protocol.

IQ in-phase and quadrature.

IRAM Institut de Radioastronomie Millimtrique.

JTAG Joint Test Action Group.

KID kinetic inductance detector.

KSPS kilo-samples per second.

LED light emitting diode.

LEKID lumped element KID.

LNA low noise amplifier.

LO local oscillator.

LPE low-pass edge.

LPF low-pass filter.

LSB lower side-band.

LUT lookup table.

MIR mid infrared.

MSPS mega-samples per second.

NEP noise equivalent power.

NET noise equivalent temperature.

NETD, NEDT noise equivalent temperature difference.

NIKA New IRAM KID Arrays.

NIKEL New IRAM KID Electronics.

NIR near infrared.

NSD noise spectral density.

PIL Python imaging library.

PLL phase locked loop.

PSD power spectral density.

PSF point spread function.

PSU power supply unit.

PTC pulse-tube cooler.

QCL quantum-cascade laser.

RAM random access memory.

RF radio frequency.

RMS root mean squared.

ROACH reconfigurable open architecture computing hardware.

SMA subminiature version A connector.

SNR signal to noise ratio.

SOI silicon on insulator.

SPS sample(s) per second.

SQUID superconducting quantum interference device.

SSB single side-band.

SSH secure shell.

TCP transmission control protocol.

TDM time division multiplexing.

TES transition edge sensor.

TLS two-level systems.

TTL transistor-transistor logic.

UDP user datagram protocol.

USB universal serial bus - upper side-band.

VGA video graphics array.

VHDL VHSIC hardware description language.

VHSIC very high speed integrated circuit.

VI virtual instrument.

VNA vector network analyser.

List of symbols

$A\Omega$ optical throughput or étendue, m^2 sr.

β ratio of fractional frequency response and dissipative response.

B_ν blackbody radiance, W sr^1 m^2 Hz^1.

C capacitance, F.

c speed of light in vacuum, 2.99792458×10^8 m s^{-1}.

D lens diameter, mm.

Δ superconducting energy gap, typically µeV.

δf KID frequency response, Hz.

$\left.\frac{dI}{df}\right|_{f_0}$ gradient of sweep in I measured at f_0.

$\left.\frac{dQ}{df}\right|_{f_0}$ gradient of sweep of Q measured at f_0.

δQ^{-1} KID dissipative response.

$x = \delta f / f_0$ KID fractional frequency response.

η_e quasiparticle production efficiency.

ϵ emissivity.

η_{det} detector absorption efficiency.

η_{opt} optical efficiency.

F focal length, m.

f frequency, Hz.

f_0 KID resonant frequency, Hz.

$f/$ f-number.

f_r frequency of readout tone, Hz.

f_s sample rate/frame rate, samples per second.

Δf_s post detection bandwidth, Hz.

γ photon.

G_{DAC} voltage gain on NIKEL DACs.

G_{Tone} voltage gain on NIKEL tones.

h Planck constant, $6.62607004 \times 10^{-34}$ m^2 kg s^{-1}.

IMA image plane in Zemax simulations.

(I_c, Q_c) co-ordinates of the centre of a resonator circle, (V, V).

$\{I(t), Q(t)\}$ in-phase and quadrature components of a readout tone, {V, V}.

$\{A(t), \phi(t)\}$ amplitude and phase components of a readout tone, {dB, rad}.

j imaginary unit.

K lens conic constant.

k_B Boltzmann constant, $1.38064852 \times 10^{-23}$ m^2 kg s^{-2} K^{-1}.

L inductance, H.

L1 the objective lens.

L2 the refocussing lens.

L3 the cold lens.

λ optical wavelength, m.

M the oscillating fold mirror.

N number of detectors.

N_0 normal state single spin density of states, μm^{-3} eV^{-1}.

N_{qp} number of quasiparticles.

n_{qp} quasiparticle density, m$^-3$.

ν optical frequency, Hz.

$\Delta\nu$ optical bandwidth, Hz.

$\mathcal{O}()$ the order of ().

OBJ object plane in Zemax simulations.

P optical power, W.

P_g electrical power from the signal generator, dBm.

ϕ_0 phase of readout tone on resonance, measured from IQ circle centre, rad.

P_FP optical power reaching the focal plane, W.

P_s optical power from the source, W.

Q_c coupling quality factor.

Q_i internal quality factor (unloaded Q).

Q_r resonator quality factor (loaded Q).

$\mathcal{R} = \frac{\nu}{\Delta \nu}$ spectral resolution.

$\mathtt{R_c}$ lens radius of curvature, mm.

R_RJ Rayleigh-Jeans responsivity ($\frac{dP}{dT}$), W K^{-1}.

R_x detector responsivity in x, x W^{-1}.

S_{11} complex input reflection coefficient scattering parameter.

S_{12} complex reverse voltage gain scattering parameter.

S_{21} complex forward voltage gain scattering parameter.

S_{22} complex output reflection coefficient scattering parameter.

σ_x^2 variance in x, units of x^2.

$S_\varphi(f)$ phase noise, dBc Hz^{-1}.

$\mathtt{STO}$ aperture stop plane in Zemax simulations.

$S_x(f)$ single sided power spectral density of x, x^2 Hz^{-1}.

$\mathtt{T}$ lens thickness, mm.

T temperature, K.

T_a amplifier noise temperature, K.

τ integration time, s.

τ_0 electron-phonon interaction time, ns.

T_c superconducting critical temperature, K.

Δt sample period, s.

T_eff effective temperature, K.

$\mathrm{TE_{m,n}}$ the (m, n)$^\mathrm{th}$ transverse electric mode.

$\mathrm{TM_{m,n}}$ the (m, n)$^\mathrm{th}$ transverse magnetic mode.

τ_qp quasiparticle lifetime, s.

V_L inductor volume, μm^3.

ω angular frequency, rad s^{-1}.

ω_r KID resonant frequency, rad s^{-1}.

Z impedance, Ω.

Z_0 characteristic impedance, usually $|Z_0| = 50\,\Omega$.

Chapter 1

Introduction and Background

1.1 Thesis outline

This thesis presents the design, implementation, characterisation and overall performance of the first passive terahertz video camera based on superconducting lumped element kinetic inductance detectors (LEKIDs). Superconducting detectors such as LEKIDs or TES bolometers offer vast improvements in sensitivity over common, commercially available, uncooled devices. LEKIDs have further advantages over similar cooled devices, including the relative ease of fabrication of large format arrays and the reduced cryogenic complexity of the readout electronics. The key result of this thesis has been to demonstrate that LEKID technology is now suitable for application in a range of fields where ultra-high sensitivity to terahertz radiation is required.

The following chapters present the various results of my research that have lead to the successful completion and integration of the core components of a LEKID based system into a fully operational camera. The thesis concludes with a comparison of the overall theoretical and measured system performances, and a comparison of the final system specifications with those of contemporary systems based on other detector technologies.

The chapter-by-chapter layout is as follows:

Chapter 1: Introduction and Background

    This summary and some additional background on terahertz imaging, kinetic inductance detectors and the KID camera. Concluding with a short list of some my original contributions to the fields of terahertz imaging and general KID research, as described in this thesis.

Chapter 2: Instrument Overview

    A complete overview of the camera, including the design and performance of the instrument in its final configuration at 350 GHz, as adapted from the instrument paper published in the peer reviewed journal Review of Scientific Instruments.

Chapter 3: The Optical System

    The design and performance of the optics system and field scanning mechanism,

including results from preliminary analyses through to the final characterisations of both the 150 GHz and 350 GHz channels. Fitting a low-cost, high magnification optical system into an existing cryostat was challenging and issues with LEKID susceptibility to stray light were identified. Results from a far field beam mapping system show improved optical coupling efficiency with the addition of a feed-horn array at the focal plane.

Chapter 4: The Detector System

The design and performance of the detector arrays, including both the 60 pixel array at 150 GHz and the 152 pixel array at 350 GHz. The noise, responsivity and microwave performance of the detectors were characterised in the group's low background test cryostat as a function of base temperature with early indications of photon noise limited performance. A study of the final detector performance was completed under high background in the camera cryostat as a function of the incident optical power from a variable temperature blackbody source at the focal plane. Results indicate that the LEKID device sensitivities are more than sufficient for background limited imaging.

Chapter 5: The Electronic Readout System

A brief study of KID readout techniques, including single pixel and multiplexing systems, as well as the design, implementation and characterisation of the two multiplexing systems utilised in this instrument. After successful imaging with the 60 pixel system, the initial architecture of the FPGA system was modified to include additional time division multiplexing capabilities for the characterisation of the 152 pixel array. A high-performance NIKEL system was then acquired, setup and configured to run at frame rates suitable for video rate imaging.

Chapter 6: The Software System and Image Generation

Overviews of the novel software systems for camera control, data acquisition, image generation and real-time video display, including calibrations, inputs, processes, outputs and data storage. Fully automated algorithms and routines are presented that operate in real-time for setting readout tones, sweeping out LEKIDs, tone power level optimisation, response calculations, flat-field calibrations, image generation and video display. A user friendly GUI system was developed to provide a front end for the operation of the camera while background processes collect and interpret the detector output signals, providing sufficient feedback into the input probe signals to optimise detector performance across the array. Single-shot images, time-lapses and recordings of live video streams are presented. These successfully demonstrate LEKID technology's potential for background limited passive imaging in the terahertz band.

Chapter 7: Camera Performance and Future Considerations

The performance of the camera in its final configuration with the 350 GHz array in

place, including a summary of system noise and sensitivity properties, a comparison with similar commercially available systems, and some considerations for future developments. Detailed models of the key system components – including an electro-thermal analysis of photon noise and stray light propagation through the camera, field-of-view scanning strategies, detector noise (G-R and TLS), and readout electronics/amplifier noise floors – are combined to indicate that the system operates within only a factor of two in NEP from the ideal background limited sensitivity.

1.2   Terahertz or sub-millimetre radiation

The terahertz region of the electromagnetic spectrum sits between the microwave and infrared regions. Definitions of the precise extents vary from a broad 0.1–10 THz range to a narrower 0.3–3 THz range. The latter band – with wavelengths between 1 and 0.1 mm – is known as the sub-millimetre (sub-mm) and extends from the millimetre (mm-wave) to the far-infrared (FIR). Terahertz photons have energies ranging from ∼0.12 to ∼12 meV and effective colour temperatures between 3 and 30 K.

The region is notable for the so-called terahertz gap, where technology for the generation, transmission, and detection of radiation has lagged behind that of the rest of the spectrum. Astronomers and atmospheric scientists have been keen to adopt the newest technological advances at these frequencies since the late 1950s and early 60s, and continue to be major drivers in the development of terahertz technology to the modern day [Siegel, 2002]. In recent decades, interest in spectroscopy and imaging in the terahertz has reached fields far and wide, from art history and the biomedical sciences, to the cosmetics industry, defence, and beyond.

1.2.1   Sources

Natural sources of terahertz radiation are generally limited to thermal blackbody emission from warm bodies – indeed, half of the total luminosity and 98% of photons emitted since the big-bang fall in the FIR. Giant molecular clouds, Bok globules, protostellar cores, massive star forming regions, planetary nebulae, supernova remnants and even the region around the supermassive black hole at the galactic centre, are all interspersed with 'cold' gas and dust where thermal emission peaks in the FIR. At these frequencies we are offered un-obscured views of star formation and stellar evolution across the Milky Way.

Furthermore, observations from the Herschel Space Observatory indicate that up to 90% of the galaxies in the local universe have spectral energy distributions with peak intensities in the FIR, and, as we look further out, these peaks are redshifted into the sub-mm and mm-wave regimes. The distributions and morphologies of these galaxies, combined with information from gravitational lenses, the cosmic microwave background and the Sunyaev-Zel'dovich effect – all available in the terahertz – allows us to probe galactic evolution, map out large-scale structure, and develop refined cosmological models of the progression of the universe throughout cosmic time.

Here on Earth, progress towards compact, coherent, high power terahertz sources is well under way. Some devices emit directly at terahertz frequencies such as gas lasers and quantum cascade lasers (QCLs), while lower frequencies – such as those emitted by Gunn diode or other parametric oscillators – can be converted up to the terahertz by fast diode or varactor based frequency multipliers. Continuous wave (CW) coherent sources such as these are typically very narrow band and only available at specific frequencies. Tunable frequency sources are becoming more accessible such as tunable QCLs, free-electron

Figure 1.1: *The Rosette molecular cloud, seen by the Herschel PACS and SPIRE instruments at 70, 160 and 250 microns is a beautiful example of astronomical imaging. Image source: ESA/PACS & SPIRE Consortium/HOBYS Key Programme Consortia [Motte et al., 2010].*

lasers, and other complex heterodyne based systems. Modern vector network analysers (VNAs) with off the shelf frequency multipliers can operate in CW or swept frequency mode at up to ~2 THz.

Broadband terahertz radiation is most commonly generated by illuminating certain semiconductor materials with ultra-short (femtosecond) optical laser pulses or by the optical rectification of femtosecond laser pulses. Terahertz pulsed imaging techniques such as that presented in Figure 1.2 enable the structural analysis of a variety of materials. Spectral analysis using pulsed sources is also common.

Very high powered (many megawatts) coherent terahertz radiation can now be generated with gyrotron devices. These high powered sources are being used to probe the plasmas in experimental nuclear fusion reactors and there are even experiments under way into using high powered pulsed THz radiation for launching payloads into orbit around the earth [Fukunari et al., 2014].

Low power (tens of milliwatts) sources are now being microfabricated and coupled via micromachined transmission lines to terahertz transistors, mixers and other components in single chip systems in a terahertz extension to monolithic microwave integrated circuits. Amongst other things, these compact systems are desirable for future portable medical diagnosis systems[1].

The number of applications involving the measurement of the transmission and reflection of locally generated terahertz waves is growing rapidly. Structural analysis of bulk materials and 3-D tomography, remote sensing, chemical and molecular analysis (organic and/or inorganic), and high speed telecommunications are but a few of the techniques enabled by the terahertz sources mentioned.

1.2.2 Detectors

In general, techniques for the detection of electromagnetic radiation can be classified according into two main schemes depending on which side of the terahertz gap one is observing. In the (high-energy) optical scheme, direct detectors absorb discrete packets of electromagnetic energy (photons) and convert them into measurable signals such as electrical currents. This is how a CCD sensor in a digital camera works. In the (low-energy) radio scheme, antennae are employed to coherently couple electromagnetic waves into electronic circuits where they can be operated upon by any number of solid state components such as transistors or diodes. This is how an FM radio works.

Classically, the challenge with the direct detection of terahertz photons has been that the photon energy, $h\nu$, is too low to promote electrons across the band gaps of common optical and infrared semiconductor photosensors. Bolometric sensors have become the direct detectors of choice as these rely on the measurement of thermal rather than electrical excitations. Kilopixel arrays of superconducting bolometers with transition edge sensors (TESs) and time domain or frequency domain SQUID readouts are the

---

[1] http://tricorder.xprize.org/about/overview
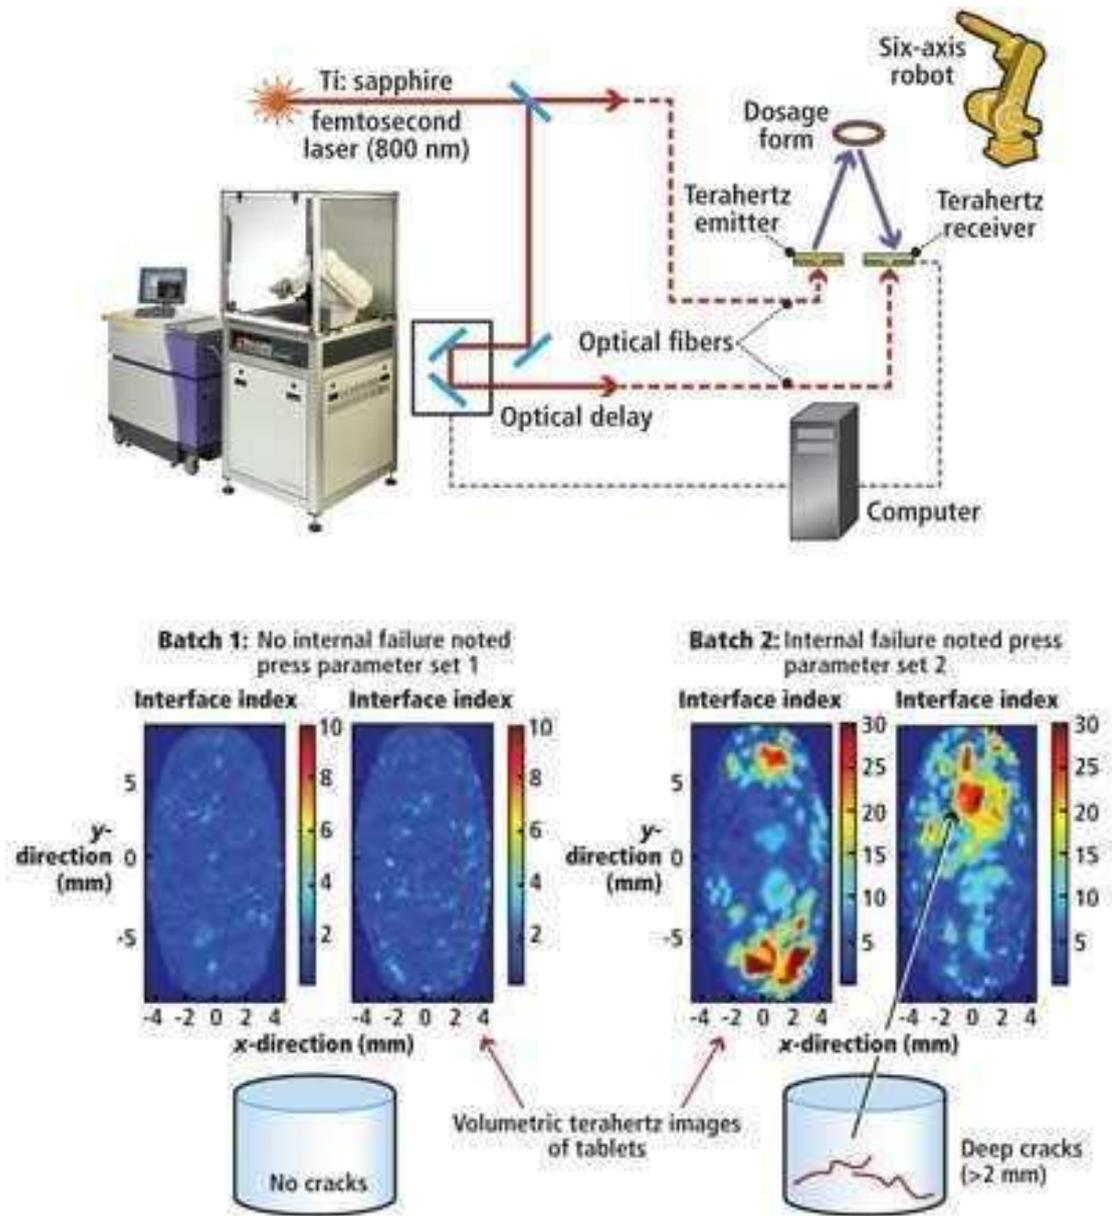
Figure 1.2: *Time domain pulsed THz imaging: defect detection in tablet coatings with a TeraView Imaga 2000 system operating at 2 THz. Image source: Vision-Systems article by A. Portieri, TeraView Ltd. [Portieri, 2012].*

present state of the art in terms of broadband sensitivity to terahertz radiation.

On the other hand, the challenge with high sensitivity coherent detection of terahertz waves has been that the oscillation frequencies of the electromagnetic fields are too fast for common solid state components to respond to. Heterodyne downconversion is the standard method for demodulating high frequency signals and bringing them into a usable intermediate frequency (IF) bandwidth, however, tremendously high speed local oscillators and mixers are required as well as low noise and high bandwidth IF amplifiers. State of the art systems utilise high power continuous wave sources such as QCLs (or lower frequency sources bumped up with frequency multipliers) as the local oscillators, ultra wide band low noise cryogenic HBT or HEMT based IF amplifiers, and critically, very low noise terahertz frequency mixers based on superconductor-insulator-superconductor (SIS) junctions or hot electron bolometer (HEB) devices.

Heterodyne systems are naturally better suited for narrow band detection applications such as high resolution spectroscopy. On the contrary, direct detectors are better for broad band detection applications such as continuum observations and low resolution spectroscopy. Direct detectors have the relative advantage of ease of fabrication into multi pixel arrays (e.g. the SCUBA-2 multi kilopixel arrays).

Heterodyne systems will always be limited in sensitivity by quantum noise in the mixers. Direct detectors do not suffer from this and commonly reach photon noise limited sensitivities. Cooling the detectors is key for achieving the highest sensitivities – not only does cooling reduce the thermal optical background, it opens the door for superconducting techniques for detection and readout.

1.2.3 Imaging

Imaging systems in the terahertz come in two flavours, passive and active. Passive systems capture the radiation that is naturally emitted by or scattered from a target. Active systems employ a terahertz source to illuminate a target and then capture either the transmitted or reflected rays. Figures 1.3 and 1.4 illustrate the differences in the particular case of concealed item detection in security applications. Figure 1.3 demonstrates two techniques for passive imaging: the first image (taken in an indoor setting) makes use of the thermal emission from the target gentleman to provide contrast against the opaque items concealed beneath his jumper; the second image (taken outdoors) derives its contrast from reflections of the cold, opaque atmosphere. The jumper and the newspaper clearly exhibit high transparency at terahertz frequencies. The active imaging system in Figure 1.4 (a body scanner now commonly found at airport security checkpoints) utilises high intensity terahertz sources and room temperature detectors to scan and measure the reflectance of the target.

Aside from the obvious privacy concerns with terahertz and mm-wave body scanners, active systems for the imaging of people (and living organisms in general) are seen as controversial as the health effects of exposure to high intensity terahertz/mm-wave

Figure 1.3:   *Passive imaging for security applications: The left image is from the passive mm-wave imaging system from Millitech in 1995 and took 30 minutes to acquire with a single pixel 94 GHz Schottky diode heterodyne receiver. The image on the right was taken with a similar passive mm-wave scanner from QinetiQ in the mid 2000s. Image sources: Airline Passenger Security Screening, National Academies Press [National Research Council, 1996], and [Appleby et al., 1994].*
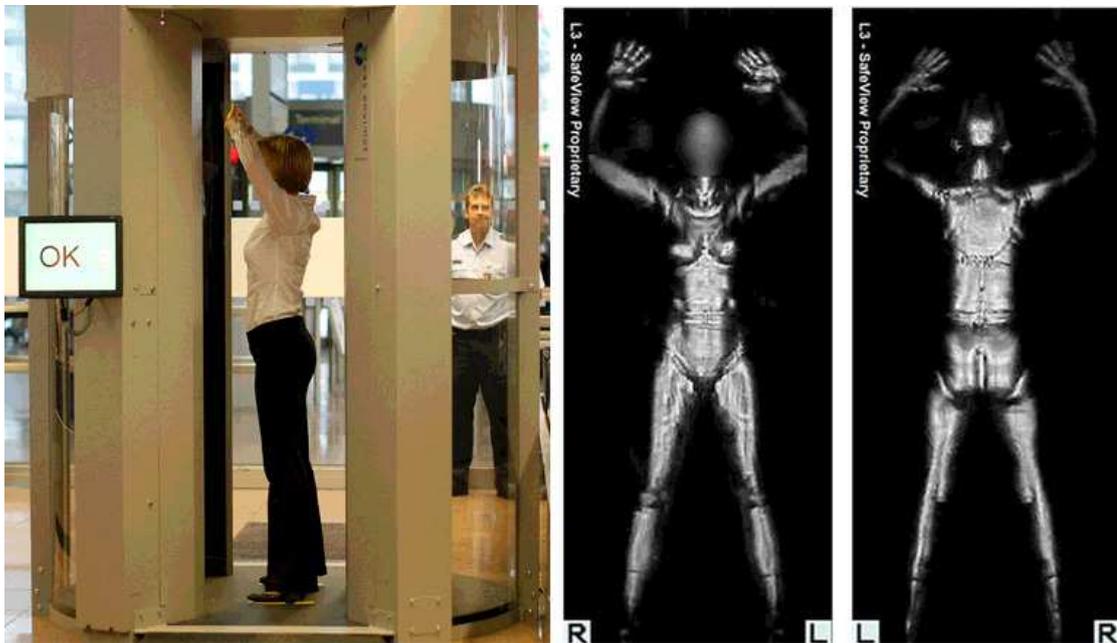
Figure 1.4:   *Active imaging for security applications: the Provision system, from L3 Communications, found in many airports and other security checkpoints scans a target in roughly one second by measuring the reflectance from a high intensity THz source. Concerns with active systems have been aired regarding the invasion of privacy and the safety of exposure to the high intensity mm-waves used. Image source: C. Charisius, Reuters; and, ITWorld article by S.J. Vaughan-Nichols.*

radiation have not been studied thoroughly. This aside, there are emerging markets for both active and passive systems. Applications include the analysis of sub-surface features in historical art works, imaging and spectroscopy in the biomedical sciences, structural and molecular analysis of cosmetics products and other pharmaceuticals (as in Figure 1.2), remote sensing in automatic navigation systems and in civil defence threat assessment; the list grows.

A number of time domain (pulsed imaging and spectroscopy) systems such as the TeraView Ltd. system shown in Figure 1.2 are commercially available, as are some active body scanners such as those from L3 Communications[2] (see Figure 1.4), Smiths Detection[3] and Rhode & Schwarz[4]. Additionally, some general-purpose uncooled passive systems are now available off the shelf, for example, NEC Corporation produces a hand held, passive, broadband (1-7 THz) terahertz camera[5] with a 320×240 element microbolometer focal plane array, a quoted sensitivity of 10^{-10} W (unspecified bandwidth) noise equivalent power (NEP) and a 30 FPS frame rate. Cooled passive imaging systems can provide many orders of magnitude greater sensitivity than this, and at the time of writing, passive systems based on cryogenic bolometer arrays with ~ 100 sensors have been developed although no commercially available products are available as of yet.

This thesis presents progress towards a full video rate passive imaging system based on arrays of direct detecting kinetic inductance detectors (KIDs) with sensitivities on the order of 10^{-16} W/$\sqrt{\text{Hz}}$. These detectors take advantage of superconducting microresonator structures to vastly simplify the fabrication and electronic readout of large format arrays. The incredibly reduced cryogenic complexity of KID systems makes them an attractive option for future scientific and commercial imaging systems that will require $\gg$ 1000s of detecting elements at terahertz frequencies.

1.3   Kinetic Inductance Detectors

Kinetic inductance detectors are thin film superconducting microwave microresonator structures that act as thermal (or square-law) detectors with optical responses that are proportional to the incident optical power. Absorbed photons break Cooper pairs in the superconductor causing variations in the complex surface impedance of the film. These variations affect the resonator quality factor Q_r and the resonant frequency f_0 which are both monitored by measuring the forward transmission S_{21} of a microwave readout signal that passes along a nearby feed line. A KID's response can only be measured when the readout tone frequency f_r is close to f_0.

Multiple KIDs, each with a different f_0, may be constructed along the same feed line, then they can all be monitored simultaneously with just one microwave signal that is the

---

[2]L£ Provision: `http://www.sds.l-3com.com/advancedimaging/provision-2.htm`

[3]Smiths Detection 'eqo' scanners: `http://www.smithsdetection.com/`

[4]R&S QPS100 system: `http://www.microwavejournal.com/ext/resources/pdf-downloads/EuMW2014DefenseForum/RandS-def-Forum.pdf`

[5]NEC passive video camera: `http://www.nec.com/en/global/prod/terahertz/`

sum of each of the individual resonator's required tones. In practice, thousands of KIDs can be read out along a single line before the up/downstream dynamic range or electronic bandwidth of the readout electronics system is exceeded – the very latest backend units have advanced DSP systems that can process over 2000 detector channels simultaneously. A-MKID [Janssen et al., 2014] is a 350 and 850 GHz dual colour instrument – currently in the commissioning phase at the the APEX telescope in the Atacama Desert – that has a total of ~24,000 KIDs that are read out with multiple feed lines and parallel sets of multiplexing backends.

Absorption of radiation can be either direct or via antenna. In the case of distributed KIDs, incoming radiation is captured with an antenna or some other absorbing element and directed into the resonant structure of the KID for detection. With lumped element KIDs (or LEKIDs), the absorbing element is itself part of the resonant structure and there is no need for additional coupling components. In either case, radiation can only be detected if the incident photon energy $h\nu$ is greater than twice the superconducting gap energy 2Δ, which for common aluminium devices sets a lower limit of $\nu \sim 90$ GHz (~ 3 mm). There is no upper limit to the energy that can be directly absorbed, so KID systems are practicable all the way up to x-ray frequencies and beyond. And, since the number of quasiparticles generated in a detection event is proportional to $h\nu$, the possibility opens up for energy resolving single photon counting imaging spectroscopy systems. The ARCHONS instrument [Mazin et al., 2013] operates on this principle at infrared to UV frequencies with a 1024 pixel aluminium optical lumped element KID (OLEKID) array.

The fabrication of entire KID arrays can be very straightforward. Once the pixel geometries and array layout have been optimised for the application in hand, the designs can be patterned into a thin metal film in a single lithographic cycle. The arrays tested in this thesis were etched from a single Al layer on a Si substrate with $> 3\,\mu$m feature sizes, and they operate with close to background limited sensitivities. More complex processes can be applied as with recent on-chip spectrometer devices that require many materials and layers. For example, SuperSpec [Shirokoff et al., 2014] is a prototype mm-wave superconducting spectrometer based on a broadband waveguide coupled horn antenna with a lithographed transition probe for coupling to a niobium microstrip feedline. 74 niobium 'U'-shaped spectral stripline filters, 74 corresponding titanium nitride LEKIDs, and a further CPW readout line enable the concurrent readout of each of the $\mathcal{R} = 700$ spectral channels centred at 250 GHz. The 20 nm thick TiN LEKIDs are partially covered with a protective SiO_2 layer and the optical and readout feedlines each have a Nb ground plane deposited over them with an internal Si_3N_X dielectric layer. [Barry, 2014]. The entire structure is temporarily mounted to a handle layer for the etching of the transition probe on the back side of the SOI substrate. All in all this is a relatively complicated design but it would be considerably more so were it not for the relative simplicity of the superconducting stripline filters, and the LEKIDs.

The fundamental ease of both fabrication and multiplexed readout makes KIDs the

ideal detectors for future instruments where large format arrays ($\sim 10^6$ detectors) will be desired. The complexity of array readout lies in the backend electronics and although the requirements increase proportionally with the pixel count, the cryogenic requirements will remain negligible and the capabilities of cutting edge DSP systems will continue progress in a Moore's Law fashion. The only real limit with fabrication and deployment of such large arrays is the physical size of the arrays given the finite dimensions of individual pixels.

Systems based on other comparably sensitive detectors, such as TES bolometers, will not be able to compete at such high pixel counts. They require highly intricate and elaborate fabrication processes as well as complex cryogenic SQUID based multiplexing systems. This results in very expensive systems when compared with KIDs. Schematics of the cross-sections of the LEKIDs tested in this thesis and the TESs utilised in the SCUBA-2 astronomical camera are presented in Figure 1.5 to highlight the comparative complexity of TES devices. For example, deep trenches and fragile SiN membranes are required to maintain strong thermal isolation between individual elements. Scaling this kind of technology – even to multi kilopixel arrays as with SCUBA-2 – is a highly challenging. Figure 1.6 presents schematics of the cryogenic readout electronics required for typical KID and TES systems. The amount of cryogenic wiring required for the TESs is vast, as is the number of Josephson junctions required for the SQUIDs. TES readouts based on SQUIDs coupled to kinetic inductance resonators are being developed to reduce the amount of wiring required between the cold stages however these systems will remain complex compared the pure KID based systems.

The historical trend of increasing pixel counts in direct detecting astronomical instruments operating across the sub-mm band is shown in Figure 1.7. The very first CMB probes of the late 80s to early 90s had a handful of individual bolometric detectors and the count per instrument has increased steadily ever since. KIDs are catching up very quickly with up and coming instruments – such as A-MKID, the NIKA-2 facility instrument at the IRAM 30 m telescope and the stratospheric polarimeter that is The Next Generation Blast Experiment – poised to become some of the most sensitive astronomical cameras to date in the far-infrared/sub-mm.

1.4 The KID camera

The Cardiff KID camera is a passive terrestrial terahertz imaging system based on lumped element KIDs. The mission was to demonstrate that KID technology is maturing beyond the niches of experimental astrophysics research to the point where it can be easily applied in a typical laboratory setting and that commercial exploitation in a variety of applications should be viable in the near future.

The camera has been a success thanks to the Astronomy Instrumentation Group's experience in various forms of terahertz technology, particularly in detector physics, terahertz optics, and instrument design. However, the camera would not be here were it

TES:

LEKID:
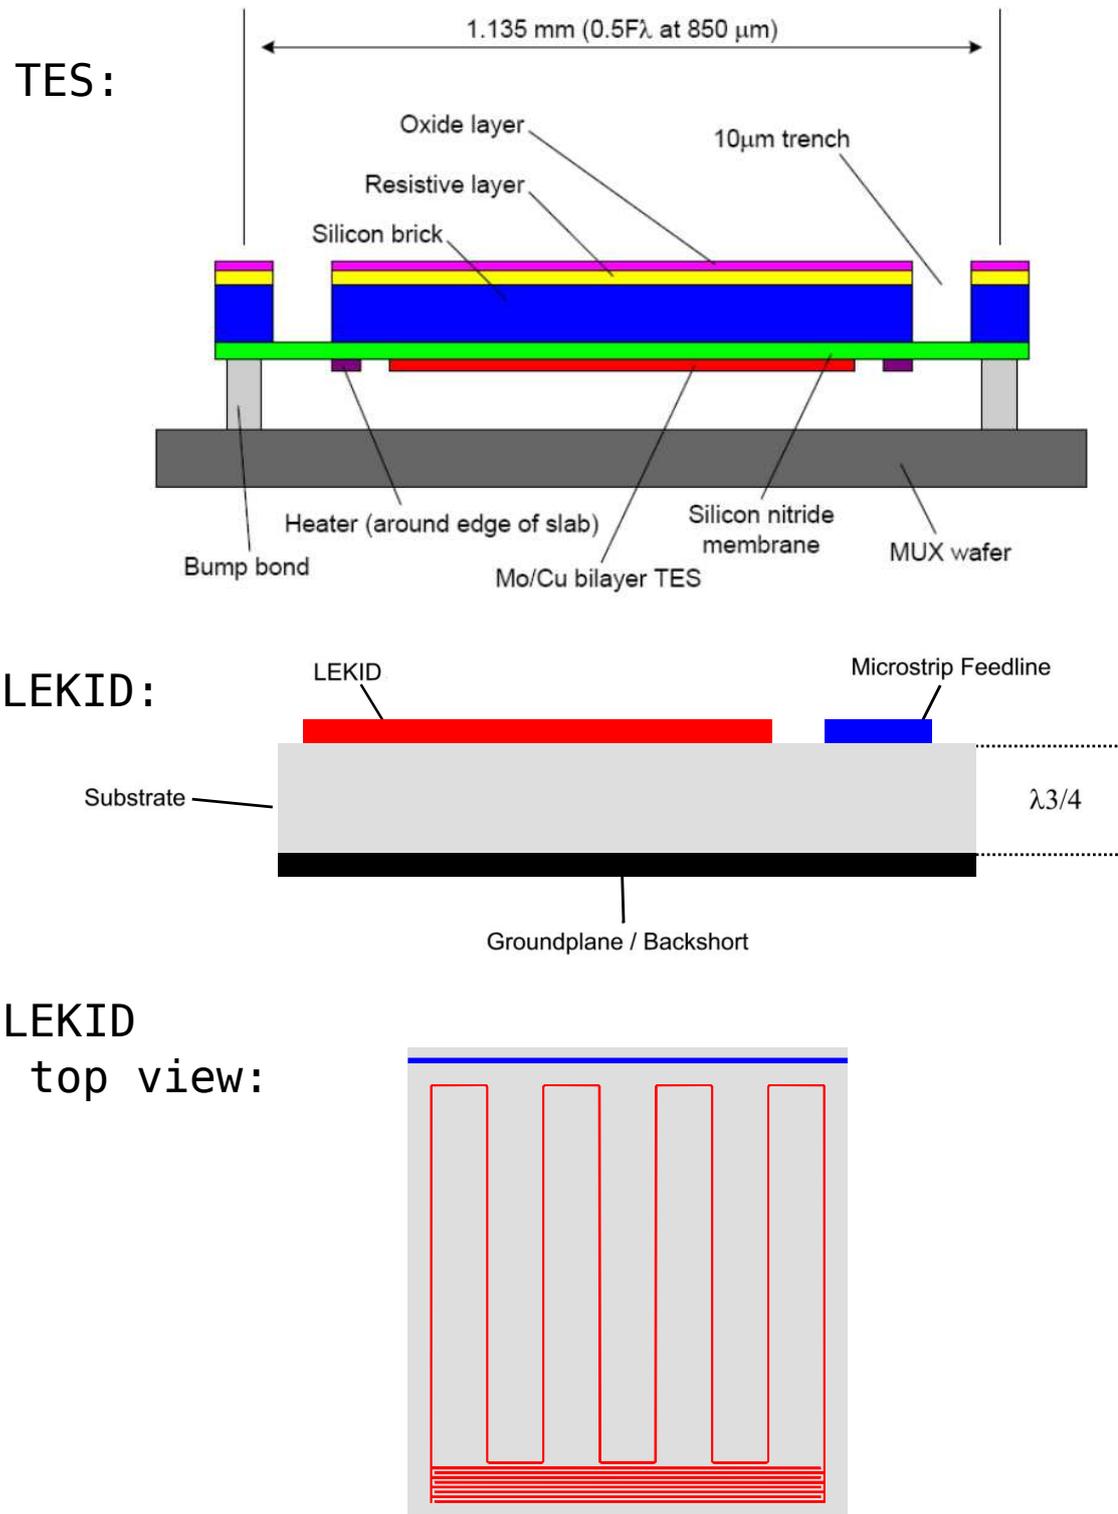
LEKID
 top view:

Figure 1.5: *A comparison of detector cross sections for a SCUBA-2 TES bolometer and a Cardiff LEKID. Fabrication of KID detectors can be remarkably simple compared to TES detectors. Image sources: [Hollister, 2009] and [Doyle, 2008].*
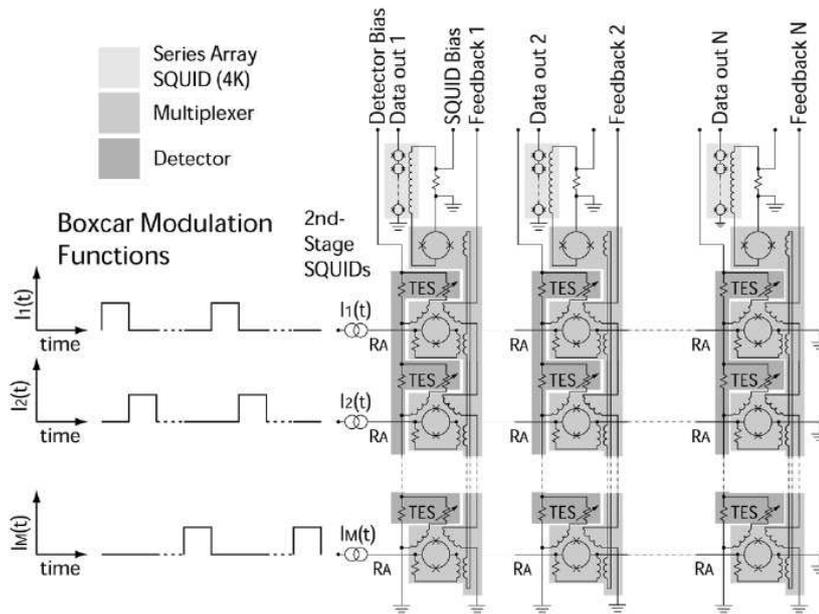
(a) *A schematic of a common cryogenic SQUID based time domain multiplexing readout system developed at NIST for arrays of TES bolometers. Four readout lines are needed per column and precisely synchronised timing signals are required to switch between the various rows. For a square array of n elements this method reduces the wiring requirements to $\mathcal{O}(\sqrt{n})$ wires rather than the $\mathcal{O}(n)$ that would be necessary without multiplexing. Image source: [Zmuidzinas and Richards, 2004]*
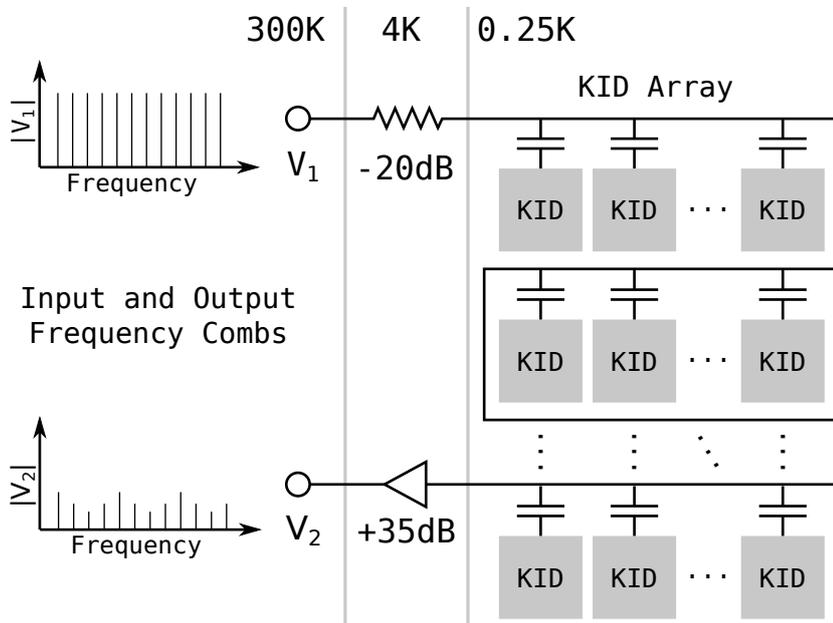
(b) *A schematic of the cryogenic frequency domain readout of typical KID arrays. Each detector is capacitively coupled to a transmission line and modulates a small portion of the input signal bandwidth. The only required cryogenic components are an attenuator and a low noise amplifier, no additional electrical signalling is required and only $\mathcal{O}(1)$ wires are needed into the cold stage.*

Figure 1.6: *A comparison of common multiplexed readout schemes for KIDs and TESs. The cryogenic requirements of the KID readout are absolutely minimal.*

Figure 1.7: *Instruments for far-infrared/sub-mm astronomy. A graph of detector count per instrument over time. Semiconductor and TES bolometer based instruments shown in red and gold respectively, while KID instruments are in green. The trends of exponential increase in pixel count are indicated by the coloured lines. The Cardiff passive imager is included for comparison although it is not an astronomical instrument.*

not for QMC Instruments Ltd. and their drive and dedication to provide their world wide client base with the latest advances in terahertz technology.

A brief overview of the development and progression from project conception to the fully operational camera is given here. First of all, a set of requirements were defined for the kind of demonstrator we wanted to make. Initial design reviews concluded with an approach based on a refractive optics system with a linear detector array and linear field scanning mechanism. Major modifications were made to an existing cryostat, including the installation of a very large window in the base, and the incorporation of cold optics and a band defining filter stack. Some preliminary detector arrays were constructed to operate in the 2 mm (150 GHz) atmospheric window but were limited to, at most, 60 pixels. This was well matched to an existing multiplexing readout system that I helped to developed as part of my Masters Degree, however, the DACs and ADCs were limited in capacity and a full video rate camera would require much faster components.

First light was achieved a year or so into the project but image quality was not . Issues with the optics system were identified and an extensive evaluation was performed. The results indicated inconsistencies between the design and manufactured lens parameters as well as the presence of significant stray light on the focal plane. A remodelling and consequent adjustment of the refractive system and the inclusion of optical baffles and a feedhorn coupling plate at the focal plane greatly improved the overall performance. The second round of imaging – both still and quasi-video rate – showed a marked improvement.

The detector array was upgraded to a horncoupled 152 pixel array operating in the 850 μm (350 GHz) window and an agreement was made with the NIKA team for access to one of their high performance 'NIKEL' multiplexing electronics systems. Prior to the arrival of the new system, extensive modifications were made to the original multiplexing system in order to characterise the new detector array – the effective bandwidth was increased at the expense of reducing the final frame rate. The new system was eventually incorporated into our own to great effect.

At this stage, the VeriCold pulse tube cooler (PTC) failed. An upgrade was ordered but incorporation of the new PTC required extensive modifications to the cryostat as well as a the application of novel mechanical and galvanic isolation system, upgraded thermometry, and re-optimised refrigeration cycles. After a final verification of the optics, electronics and detector systems, focus was shifted to improving the real-time video and imaging software. Some fine tuning of the calibration and image processing methods worked to clean up the images and videos. Beyond this point, no further improvements were possible without the application of major upgrades to the optics system, the scanning mechanism, the detector array and/or the readout electronics.

The camera serves well its purpose as a demonstrator of KID technology. It has attracted the attention of people and organisations around the world who are keen to apply high sensitivity passive imaging systems to their particular fields of interest. The camera also serves as useful testbed for KID systems in general as it is very easy to

swap out the focal plane arrays. A number of investigations into device physics have been performed with our system. Finally, the system is an excellent working example of how technology transfer can occur from purely academic research in astronomy and experimental astrophysics into terrestrial applications in the scientific, industrial, and commercial sectors.

1.5 Original contributions to the field

An adapted version of the published instrument paper is given in the following chapter. It describes the camera, its performance and its significance in the context of the field of terahertz imaging. A list of my primary original contributions to the field is given here:

- Demonstration of an ultra-high sensitivity camera for background limited, video rate, passive terahertz imaging.

- Development of a cryogen-free cooling platform suitable for ultra-high sensitivity terahertz imaging.

- Development and characterisation of arrays of lumped element kinetic inductance detectors suited for terahertz imaging in room temperature environments.

- Demonstration of significant improvements in optical coupling to LEKID arrays and stray light reduction in high magnification, terahertz frequency optical systems through the inclusion of feed-horn arrays and optical baffles.

- Development of a multiplexing electronic readout system with a novel time-domain switching capability that extends the effective IF bandwidth of the system and enables the simultaneous readout of very large KID arrays without the need for very high bandwidth components.

- Development of real-time control and acquisition systems for automatic detector readout, including IQ calibrations, optical response calculations and data storage/display.

- Demonstration of significant improvements in KID array yield and overall system sensitivity with the optimised fine-tuning of individual detector probe signal levels. Including an innovative automated process to run the optimisation during system operation.

- Development of efficient real-time image generation/processing software and novel flat-field calibration schemes for real-time terahertz imaging with LEKIDs.

- Demonstration of real-time stand-off scanning and concealed object identification using LEKID technology.

Chapter 2

Instrument Overview

This chapter presents a complete overview of the design and performance of the terahertz camera in its final configuration, as adapted from the instrument paper published in the AIP journal, Review of Scientific Instruments: http://dx.doi.org/10.1063/1.4941661.

2.1 Introduction

Modern astronomy requires state-of-the-art technology for the efficient detection of the faintest light from the farthest reaches of the universe. It is not uncommon for the technologies developed by astronomers to find uses in everyday life. Rosenberg et al [Rosenberg et al., 2014] have compiled numerous examples of such technology transfer including, but not limited to: CCDs – popularized by the *Hubble Space Telescope* and now used in practically every digital camera; wireless local area networking – utilizing algorithms from image processing in radio astronomy; computerized tomography in modern medical scanners – based on aperture synthesis techniques from radio interferometry; and gamma ray spectrometers for lunar/planetary surface composition analysis – now used to probe historical buildings and artefacts.

Ongoing successes in sub-millimeter astronomy (e.g. the *Herschel* [Pilbratt et al., 2010, Griffin et al., 2010, Poglitsch et al., 2010, de Graauw et al., 2010] and *Planck* [Planck Collaboration et al., 2011, Lamarre et al., 2010, Bersanelli et al., 2010] space telescopes) and the ever present demand for instruments with improved sensitivities and mapping speeds at terahertz (THz) frequencies have spurred the development of highly sensitive detectors, sophisticated optical components, cutting edge electronics, and advanced data processing techniques.

Kinetic Inductance Detectors (KIDs) are contemporary superconducting pair-breaking detectors that operate across the spectrum from x-ray to sub-THz frequencies [Day et al., 2003, Mazin et al., 2006, Mazin et al., 2013, Monfardini et al., 2010]. Compared to alternative THz technologies such as semiconductor or Transition Edge Sensor (TES) bolometers, KIDs are relatively simple to fabricate and read out. As such, they provide a practical and cost-effective solution to the manufacture and operation of the large

format arrays required for advances in many fields of THz astronomy. A variant known as the Lumped Element KID, or LEKID [Doyle et al., 2008], has been demonstrated to provide state-of-the-art performance at millimeter-wavelengths [Mauskopf et al., 2014] and has seen first light as part of the NIKA [Monfardini et al., 2014] instrument at the IRAM 30-m telescope. Projects such as The Next Generation Blast Experiment [Galitzki et al., 2014], NIKA-2 [Monfardini et al., 2014] and A-MKID [1] are currently under way to incorporate multi-kilopixel KID arrays into astronomical cameras with the potential for THz megapixel imaging within the next decade.

Beyond astronomy, the THz region of the electromagnetic spectrum (0.1-10 THz) has applications in a range of fields – academic and industrial [Siegel, 2002]. In addition to the presence of a multitude of interesting spectral features, many typically opaque materials become transparent when viewed in this frequency range. Various disciplines – including biomedical sensing, non-destructive testing, and security screening – now have the opportunity to benefit from the highly sensitive and highly multiplexable detector technology being developed by astronomers.

For example, THz radiation is being used to study protein dynamics [Knab et al., 2007], to investigate interactions between THz waves and human DNA [Titova et al., 2013], and as a potential imaging modality for the improved identification of skin cancers [Woodward et al., 2003]. However, there are currently no off-the-shelf THz imaging spectrometers or cameras available to help proceed more rapidly with these investigations.

The analysis and restoration of cultural artefacts benefits from the unique differential penetration of THz radiation, making it ideal for the non-destructive investigation of the internal paint layers in pieces of art [Seco-Martorell et al., 2013]. Time domain techniques have been used to show that unique information can be gleaned at THz frequencies to verify the age, chemical composition and structure of works of art.

Far larger potential demand is associated with the detection of hidden objects (such as land mines [Osiander et al., 2003]), process control in manufacturing [Fitzgerald et al., 2005], and security screening [Grossman et al., 2010]. Active mm-wave scanners are now widely deployed in airports across the globe and large format KID arrays could be used to produce systems with improved sensitivity at a comparable cost. The capacity for truly passive imaging, and the fast time-response of LEKID detectors (typically $< 10^{-4}$ sec) enables, for the first time, the possibility of capturing images at video rate for so called "walk through" systems. This is regarded as desirable by ECAC (the European Civil Aviation Conference) and opens up the possibility of use in situations where requiring people to stand one-by-one in a booth is not practical [European Commission, 2010]. Furthermore, multi-spectral observations would improve image contrast and reduce the number of false positives – a common occurrence with current active systems. However, the efficacy of passive terahertz imaging for security applications has yet to be fully demonstrated under and a range of environmental conditions and a systematic

---

[1] http://www3.mpifr-bonn.mpg.de/div/submmtech/bolometer/A-MKID/a-mkidmain.html

study into the direct and indirect effects of temperature, humidity and precipitation on the appearance of the subjects undergoing screening would be a worthy endeavour. Unfortunately, this remains outside of the scope of this article.

To demonstrate the capabilities described above, we have built and characterized a simple field scanning camera based on a 152 element linear array of LEKIDs operating at 350 GHz. The detectors have been optimized to perform under the optical loads present at ambient room temperatures ($\sim$300 K), which are substantially higher than the backgrounds present during astronomical observations. The instrument, in its present configuration, is comparable in performance to other recent passive imaging systems including those based on room temperature microbolometer focal plane arrays (FPAs) [Shumaker et al., 2013, Oda et al., 2014], cooled bolometer arrays [Grossman et al., 2010], and superconducting TES arrays [Becker et al., 2010, Heinz et al., 2012].

This chapter describes the camera and its achieved performance as a quasi-video-rate system. It concludes with discussions of the improvements which will be implemented for the next generation camera in order to achieve a full video-rate, photon-noise limited imaging system.

2.2 Requirements

Our goal was to demonstrate a video rate scanner capable of imaging variations in the thermal THz radiation received from a moving target (a person) with sufficient sensitivity to detect and identify concealed objects - akin to airport-style security scanners or other stand-off scanning instruments. The basic requirements were for a simple-to-use system with the necessary spatial resolution, scanning speed and sensitivity to identify objects of a few cm in size concealed behind clothing or inside bags or other luggage.

The camera is designed to provide a 1×2 m useful field of view (typical of body scanners) with operation at a distance of 3-5 m from the target and a linear resolution of roughly 1 cm. The camera observes in the 350 GHz transparent atmospheric window with a $\leq 10\%$ wide band to minimize the loading and thermal fluctuations present at less transparent frequencies. As a demonstration system, quasi-video rate imaging with frames updating at least every second was deemed acceptable although the goal would be to reach full 25 Hz video rate. A 0.1 K noise equivalent temperature difference, NEΔT, in each frame is required as this enables high fidelity imaging and will ease the identification of the shape of concealed objects. Deitlein et al [Dietlein et al., 2006] neatly present the detrimental effects of increased image noise on the ease of object identification. Finally, the superconducting detectors need to be operated at around 0.25 K, so the camera requires a completely dry cryogenic system which, unlike wet systems, can be easily deployed in the field.

Ideally, the noise performance of the system would be limited by variance in the arrival of photons from the source rather than from any components of the camera itself. The noise equivalent power due to photons measured at the detector focal plane in a

diffraction limited optics system is given by [Lamarre, 1986]

$$NEP_{\text{photon}} = \sqrt{2Ph\nu + \frac{2P^2}{m\Delta\nu}}\,, \tag{2.1}$$

where h is Planck's constant, ν is the frequency, $\Delta\nu$ is the optical bandwidth, and $m = 2$ for a detector absorbing light in both polarizations. P is the power in a band of width $\Delta\nu$:

$$P = \eta_{\text{opt}}\eta_{\text{det}}A\Omega\epsilon(\nu)B_\nu(T)\Delta\nu\,, \tag{2.2}$$

where $A\Omega$ is the camera throughput or étendue, $\epsilon(\nu)$ is the emissivity and $B_\nu(T)$ is the blackbody radiance of a source at temperature T, while η_{opt} and η_{det} are respectively the optics and detector efficiencies. It is important to note that the two terms in Equation 2.1 account for both the familiar photon shot noise (signal variance caused by the random distribution of independent light particles) and the additional photon wave noise (a signal variance introduced by the bunching of photons in highly coherent quantum states).

2.3 Camera Design

The camera employs a linear array of detectors housed in a research cryostat retrofitted with a large (250 mm diameter) window in the base. Incoming radiation is coupled to the detectors through a refractive optics system and a flat beam-folding mirror. A thin horizontal section of the object plane is observed in any one instant and this section is scanned continuously in the vertical direction by oscillation of the beam-folding mirror. This is illustrated in Figure 2.1. The orientation of the mirror is recorded with an absolute encoder and images are reconstructed in real time by the acquisition electronics in a scheme no dissimilar to that of a common office desktop scanner.

2.3.1 Optics

A fast ($f/0.9$) triplet of high-density polyethylene (HDPE) lenses was designed to keep the optics simple and compact given limitations on where the focal plane array could be situated within the cryostat. To achieve the desired resolution at these frequencies, a large diameter (440 mm) primary lens, L1, was chosen. The focal distance of the camera is designed to be adjustable between 3 and 5 meters depending on the position of the secondary lens, L2. At a distance of 3.5 m, the scanned field of view is 0.8×1.6 m and the working depth of field is approximately ±150 mm inside and outside of the focus. The third lens, L3, visible in the CAD model in Figure 2.2, is housed within the cryostat behind the HDPE window and a number of thermal blocking filters [Tucker and Ade, 2006] (not shown in the figure).

The lens and window absorptivities were measured in band and are non-negligible, with combined losses of up to 45% expected through the optics chain. Furthermore, the HDPE components are not anti-reflection coated and are uncooled (except L3).
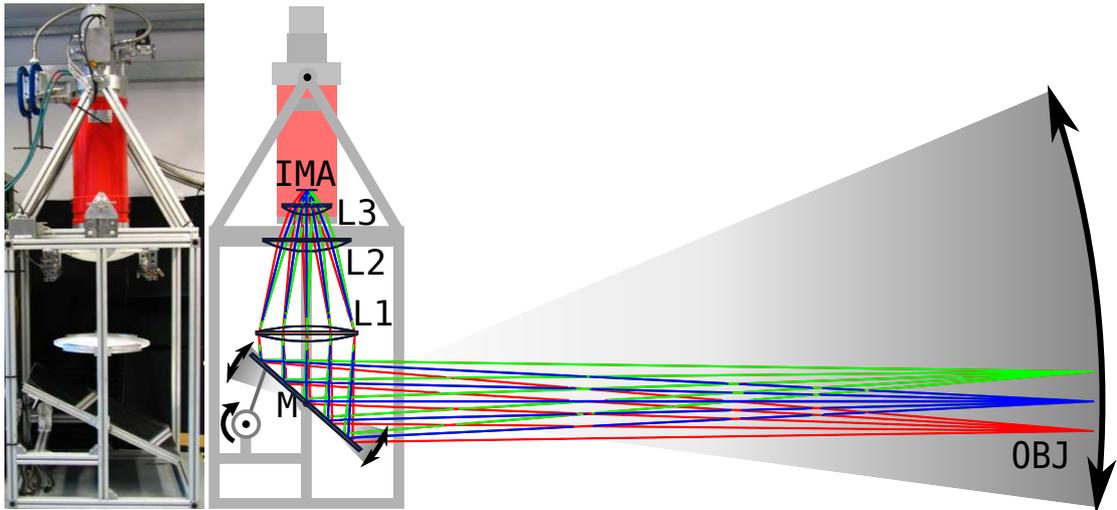
Figure 2.1: *A photo of the current system (left) and a schematic (right) of the system overlaid with the Zemax model and ray traces. The camera sees a small horizontal strip of the object plane in one instance with the oscillating fold mirror providing full sampling in the vertical direction.*

Consequently, stray light from these sources contributes significantly to detector loading.

The oscillating beam-folding mirror is constructed from a thin, polished sheet of aluminium (800 mm long by 550 mm wide) braced with strut profile and mounted to the camera's main frame via a set of bearings on the central horizontal axis. The oscillation is brought about by a crank wheel driven by a servo motor located behind the mirror. A small steel rod with bearings at each end connects the mirror to the wheel. The oscillation rate is controlled by a motor driver that is configured via USB from the control station. This mechanism can modulate the field of view at a maximum frequency of 2-3 frames per second, this ultimately limits the video rate output.

A series of quasi-optical metal-mesh filters [Ade et al., 2006] define the optical bandwidth of the system. Currently, three low-pass edges with cut-offs at 630 GHz, 540 GHz, and 450 GHz and two 10% wide band-pass filters define a combined 6% wide band centred at 347 GHz. The additional bandpass filter was added as a precaution against detector saturation with the effect of reducing the overall bandwidth and the camera optical efficiency. The filter profiles were measured by a Fourier Transform Spectrometer (FTS) from 200 GHz to 1 THz with 1 GHz resolution and are displayed in Figure 2.3. Inset to the figure is a plot of the total transmittance of the filter stack. The peak in-band transmission is 45% and the out-of-band rejection at high frequency is better than 100 dB.

The large cryostat window and the requirement for fast optics make the focal plane susceptible to off axis radiation. To lessen any stray light effects, SiC blackened metal baffles are arranged at the entrances to the three radiation shields and a feedhorn plate is mounted to the detector array at the focal plane, see Figure 2.2. The back-to-back copper feed-horns are each approximately $f/1.3$ and whilst this helps prevent stray light
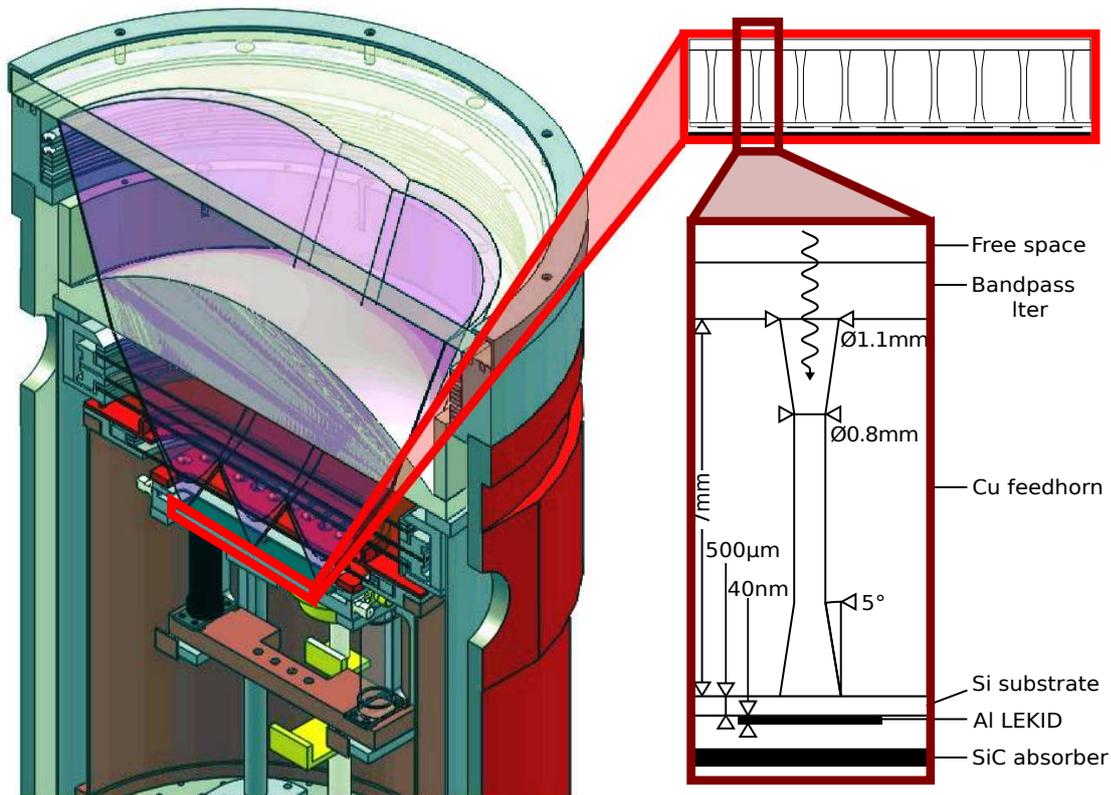
Figure 2.2: *Cross sections of the cryostat and focal plane assembly. The model on the left shows the focal plane array (teal) mounted via various thermally isolating support structures to the 4 K base plate and connected to the He-10 fridge (yellow). Optical paths from the focal plane, through the quasi-optical filters, baffles (not shown), cold lens and the cryostat window are indicated the transparent cones for pixels in the centre and edges of the array. The schematics on the right shows sections of the horn coupling system with a zoom in of a single array element. The aluminium LEKID device is back-illuminated through the silicon substrate via a back-to-back conical feedhorn (for off-axis stray-light reduction) and the final band-defining filter. Any radiation that transmits though the detector is caught by a layer of silicon carbide infused epoxy at the back of the array packaging.*
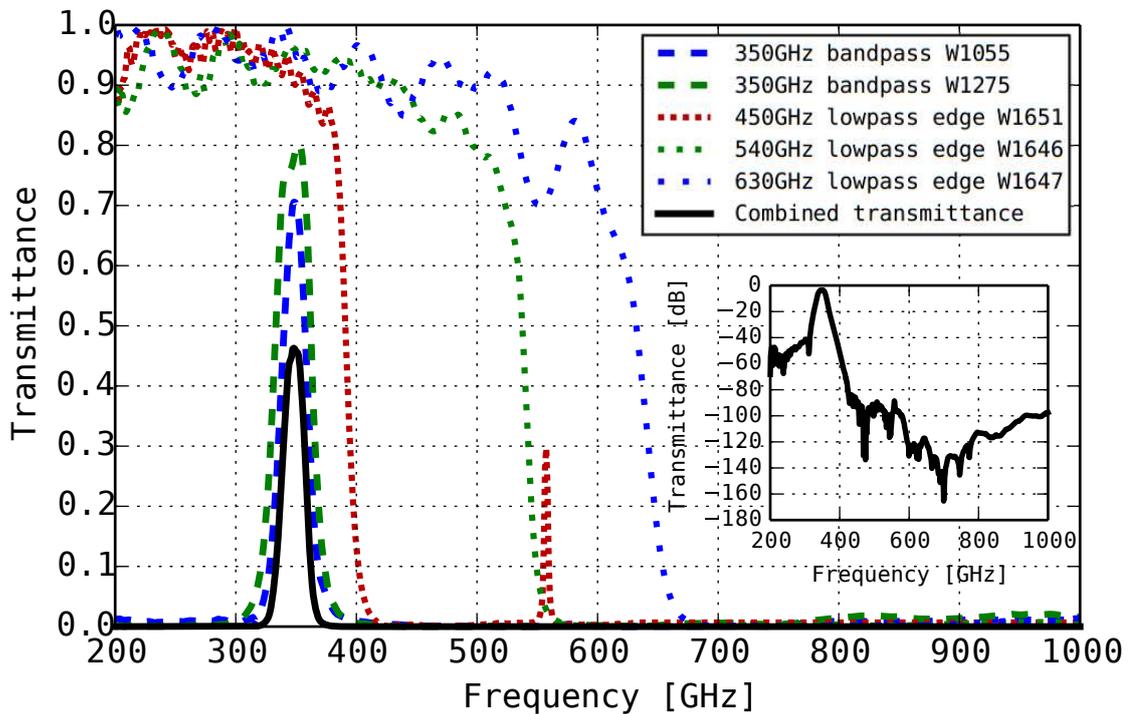
Figure 2.3:   *The spectral transmittance of the band-defining quasi-optical filters as measured by a Fourier transform spectrometer. Thermal filters are excluded for clarity as these have ~100% transmission in band. (Inset) The product of these measured profiles is computed numerically and presented on logarithmic scale to indicate the expected high out of band rejection of source power.*

reaching the detectors, there is a slight mismatch with the $f/0.9$ refracting optics. The cylindrical waveguides connecting the back-to-back horns admit at most two transverse electromagnetic modes, the TE_{11} and TM_{01} modes.

2.3.2 Array definition

The number of detectors in the array needed to achieve the required performance is estimated. Each lens in the system is characterized by its emissivity and transmission properties. L1 and L2 and the window operate at $300\,K$, while the cold lens, L3, operates close to $150\,K$. Having measured the HDPE transmission (the absorption coefficient is $0.127\,m^{-1}$ at $350\,GHz$), we estimate the overall lens transmission using Zemax[2]. The overall instrument efficiency, including the filters, is 23%. Upon consideration of thermal emission from the each of the optical components, this corresponds to an expected load of $131\,pW$ per detector at the focal plane assembly. The photon noise (including both the shot noise and wave noise components) at the focal plane is then calculated from Equation 2.1 to be $3.0\,mK\sqrt{s}$. This allows us to estimate that in order to achieve an image sensitivity of $\sim 0.1\,K$ per frame at a $25\,Hz$ frame rate and a $1\,cm$ resolution, 150 detectors are sufficient, which is compatible with the available space in the cryostat focal plane and electronic readout limitations. Note that to first order this estimate is independent from the detector and optical efficiencies, as the photon wave noise dominates the noise budget with $\frac{NEP_{shot}}{NEP_{wave}} = 0.23$.

The detector array in use for this demonstration system is composed of 152 LEKIDs arranged in 8 rows of 19 columns. The columns are skewed such that the instantaneous field of view is Nyquist sampled in the horizontal direction (see Figure 2.4).

2.3.3 Detector system

In general, a kinetic inductance detector (KID) is fabricated by patterning a thin film of superconducting material in such a way as to create an LC resonant circuit with frequency $f_0 = 1/(2\pi\sqrt{LC})$. The inductance of the superconductor, L, has two key components, $L = L_{geometric} + L_{kinetic}$. These depend, respectively, on the shape of the patterned detector and the density of Cooper pairs in the film. Photons that couple into the resonator with sufficient energy to overcome the superconducting gap will break Cooper pairs into unbound pairs of quasiparticle excitations, leading to a decrease in f_0. Then, any variations in incident optical power are monitored by measuring the variations in f_0. This is achieved by monitoring the complex transmission of a probe signal that is fed through a microwave transmission line adjacent to the resonator. Multiple resonators, each with a different f_0, may be coupled to the same transmission line and read out simultaneously with a superposition of probe signals. This inherent multiplexing capability considerably reduces the requirement for complex cryogenic circuitry.

---

[2]http://www.zemax.com
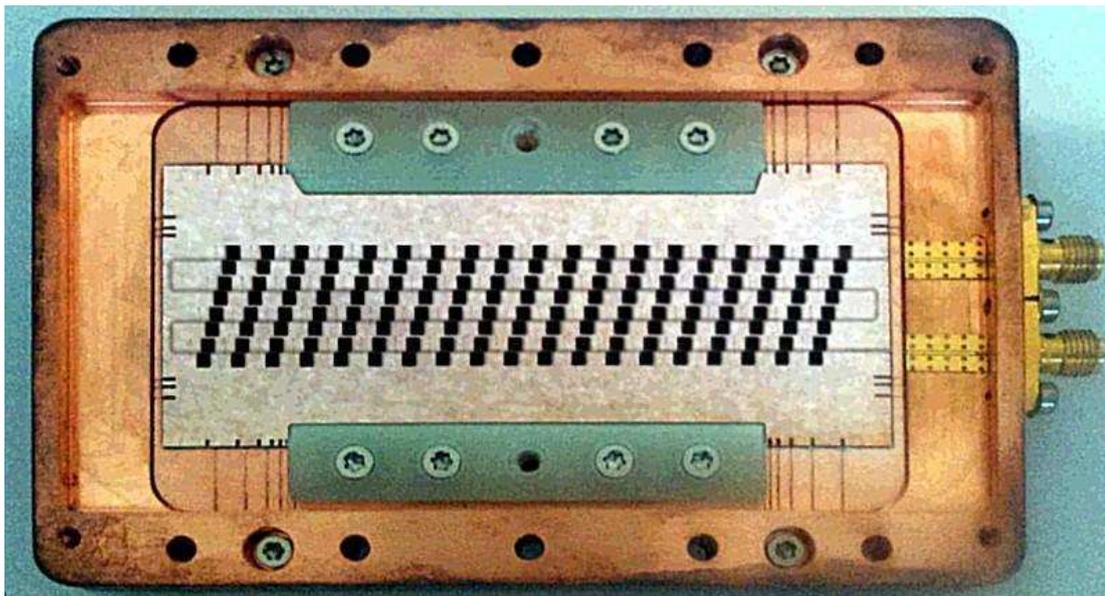
Figure 2.4: *The array and packaging. The detectors are arranged to Nyquist sample a horizontal section of the object plane. For readout, each detector modulates a small range of the total bandwidth of the probe signal that propagates along the feedline. The feedline can be seen winding between the rows of detectors and is terminated to SMA type connectors at each end.*

Lumped-element KIDs – as opposed to distributed KIDs – are designed such that the absorbing element of the detector is part of the resonator structure itself. In this configuration it is possible to achieve very high filling factors in focal plane arrays without the need for additional coupling optics such as microlens or feedhorn arrays. Note that the feed-horns used in this system are for stray light reduction only and would not be necessary in a fully baffled optical system.

Each lumped resonator in the current focal plane array has three sections: an inductor, an interdigital capacitor and a coupling capacitor. These are highlighted by the different coloured sections in the design and the equivalent circuit in Figure 2.5a. The inductor section is a 4th order Hilbert curve which efficiently couples to both orthogonal polarizations of incoming radiation [Roesch et al., 2012]. Variations in the length of the interdigital capacitor sections have been designed to set a range of resonant frequencies centred at $1.5\,\mathrm{GHz}$ and each separated by $3\,\mathrm{MHz}$. The detectors are capacitively coupled to a coplanar waveguide (CPW) feedline, with the length of the coupling capacitor section and its distance from the feedline limiting the Q-factor of the resonators to be of the order of 10,000.

The array is fabricated from a $40\,\mathrm{nm}$ aluminium film deposited by thermal evaporation onto a $500\,\mathrm{\mu m}$ high resistivity float-zone silicon wafer. The array design was patterned into the aluminium in a single photolithographic step with a wet etch of orthophosphoric acid, nitric acid, and water in a 25:2:6 ratio. The CPW line is cross-bonded with wire bridges at regular intervals to ensure a constant potential across the ground plane, thus
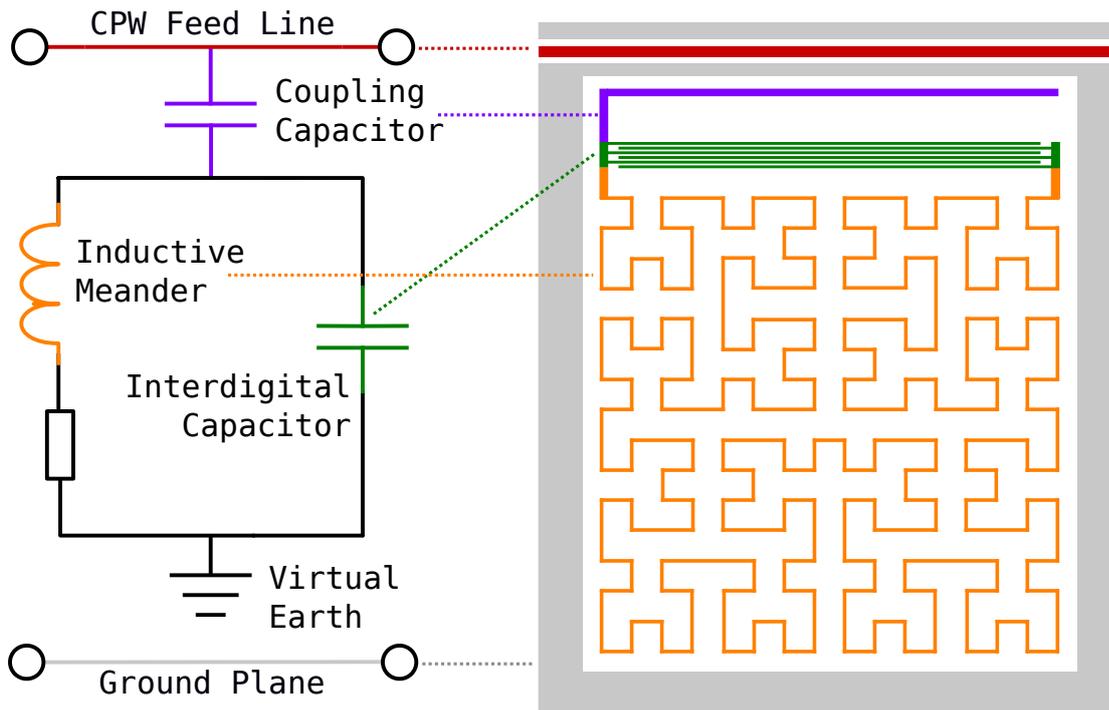
Figure 2.5: *The LEKID design used in this camera. The absorbing inductive meander (orange), the interdigital capacitor (green), the coupling capacitor (violet), the CPW feedline (red) and the ground plane (grey) are etched from a 40 nm Al film on a Si dielectric substrate (white).*

inhibiting problematic slotline modes in the CPW line.

Figure 2.2 includes a cross section of a single detecting element in the focal plane assembly. Optical coupling is optimised by back-illumination of the detectors through the silicon substrate.

2.3.4 Cryogenics

Thin film aluminium has a superconducting transition temperature of $T_c \sim 1.5\,\mathrm{K}$ and KID arrays require cooling to at least $T_c/6$ in order to sufficiently reduce the density of quasiparticles in the superconducting film. The current system utilizes a Cryomech[3] PT400 series pulse-tube-cooler (PTC) and air-cooled compressor unit that operate off mains electricity only so that no liquid cryogens are required. A closed-cycle He-10 adsorption fridge from Chase Cryogenics [Chase Research Cryogenics Ltd.,] cools the focal plane assembly to the required sub-Kelvin temperatures. Thermometry and fridge-cycling are fully automated and may be monitored/controlled remotely.

Cool-down of this thermally unoptimised demonstration system from room temperature takes around 36 hours with the PTC cold head settling at 3.2 K. The optical baffles on the radiation shields settle at 4.2 K and 60 K respectively, and the cold lens settles with a radial temperature gradient ranging between 100-150 K. In the present

---

[3]http://www.cryomech.com
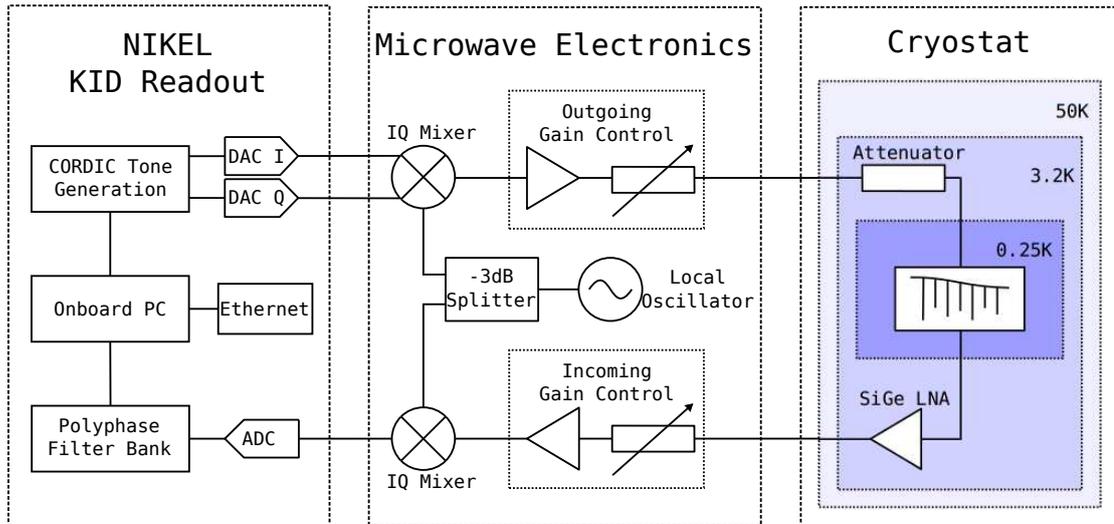
Figure 2.6: *An overview of the readout system. Tones are output by the NIKEL digital electronics system, mixed up to the required KID resonant frequency range, and passed into the cryostat through a single coaxial cable. The probe signal is attenuated before reaching the detectors and then boosted by a Si-Ge cryogenic low noise amplifier at the 4 Kelvin stage. Outside of the cryostat the signal is mixed back down to the DAQ band and read back in to the NIKEL electronics for spectral decomposition. The I and Q components of the 152 readout signals are sent, via Ethernet, to the control computer for processing and image generation.*

(un-optimised) cryogenic configuration, the fridge runs for approximately 16-18 hours at a time at 250 mK and require 3 to 4 hours for recycling. Although not incorporated in this demonstrator, continuous cooling can be readily achieved at these temperatures with so-called 'push-me-pull-me' or tandem refrigerators. Klemencic [Klemencic et al., 2015] presents a novel and economical closed cycle design that can provide steady cooling power over time scales of several months (at least) of operation.

2.3.5 Electronic Readout

The electronic readout system consists of cryogenic, warm and digital components (see Figure 2.6), as well as a suite of software to control the camera components, to monitor the housekeeping system and to generate and display images in real-time. The nature of multi-channel KID read out is such that the complexity of the cryogenic electronics is reduced to an absolute minimum. Aside from the detector array itself, a single attenuator, a single low noise amplifier, and a single pair of coaxial cables are the only components required within the cold stages.

The ends of the aluminium CPW transmission line on the array wafer are wire-bonded to SMA connectors mounted to the copper array packaging. A pair of semi-rigid coaxial cables then feed out to the 4 K stage where a cold RF attenuator on the input channel reduces the power (and the thermal noise) in the multiplexed probe signal prior to the

detector array, and A Caltech CITLF4 SiGe low noise amplifier (LNA) [4] – with <7 K noise temperature – sits on the output channel boosts the probe signal prior to readout. Further semi-rigid coaxial cables then feed out to hermetic SMA connectors on the cryostat exterior. Stainless steel coaxial cables are used between the 250 mK and 4 K stages to minimise the thermal load on the fridge head – the additional cable attenuation introduced prior to the LNA does not significantly affect the signal quality. Copper coax cables are used through the rest of the system where thermal loading requirements are more relaxed. A schematic of the cryogenic readout system is presented in Figure 2.5(b).

A room temperature analog mixing circuit converts the probe signal to and from the 1.25-1.75 GHz detector readout band and the 0-500 MHz digital electronics band. An R&S SMF100A signal generator is used as the LO input for a pair of Marki IQ mixers and a combination of amplifiers and variable attenuators are in place to balance the incoming and outgoing power levels.

The digital system is a NIKEL [Bourrion et al., 2012] (New IRAM KID Electronics) frequency domain multiplexing system developed for the NIKA astronomical camera. It has the ability to output the in-phase and quadrature (I and Q) components of the superposition of up to 400 CORDIC-generated tones across 500 MHz of DAC bandwidth. A single ADC feeds into a polyphase filter bank and the resultant 400 independent decomposed I and Q time-streams (as well as the mirror encoder values and other housekeeping data) are decimated and sent via the on-board computer over Ethernet to the control station. The sample rate is limited to 477 Hz which provides a data rate of 24 Mbps. The control station is a desktop computer equipped with a custom software suite for control of the readout electronics, data acquisition, image generation and graphical display. The readout electronics system is initialized with commands sent over UDP to the NIKEL on-board computer.

The detector responses (variations in f_0, otherwise referred to as δf or df) are computed from linear transformations of the raw I and Q time-streams using coefficients from frequency sweep data taken across the resonators during initialization. A flat field calibration is performed at the start of each run where the detector responses are measured between a 30°C glow bar and a room temperature section of the field of view. Low frequency gain variations between detectors will eventually have a detrimental impact on the final video frame quality, so the flat field coefficients can be remeasured on demand. This generally required after roughly 30 minutes of continuous operation.

The raw I and Q, the transformed amplitude, phase, and df, and the calibrated response time-streams can be accessed and displayed alongside their power spectral densities using the real-time plotting software KST[5]. Otherwise, image generation is performed on a scan-by-scan basis by reading the latest data, applying the transformations and calibrations, binning these products into a map, and updating the graphical interface with a new frame. Broken or poorly performing detectors can cause blank or noisy

---

[4]http://radiometer.caltech.edu/datasheets/amplifiers/CITLF4.pdf
[5]https://kst-plot.kde.org/
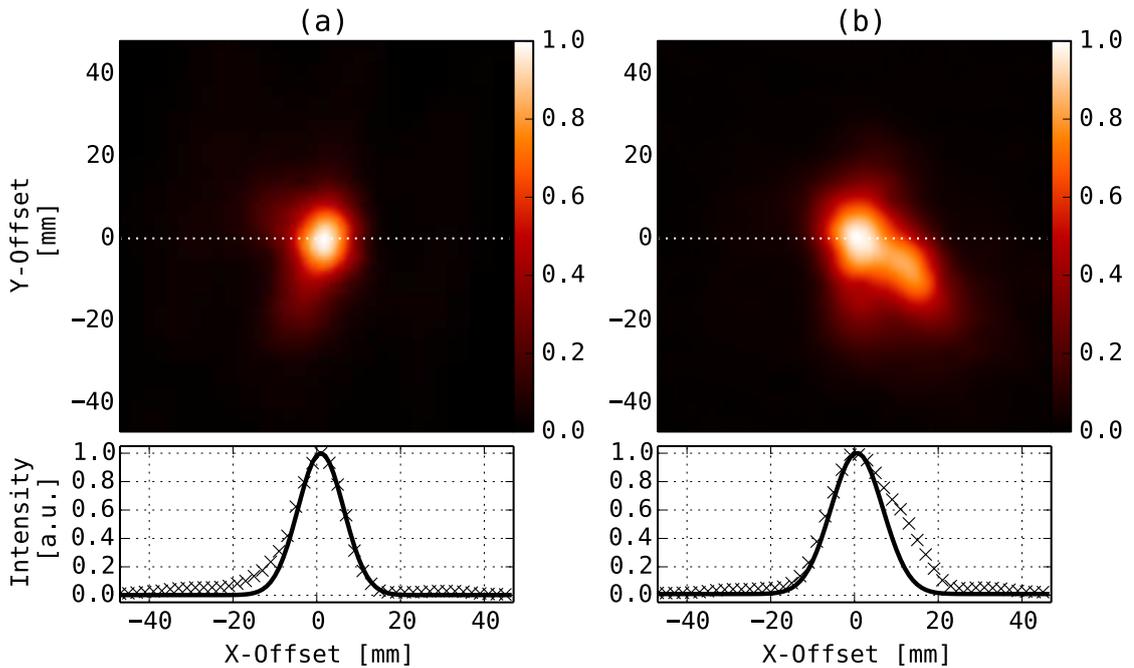
Figure 2.7: *Beam profile maps and central slices from two channels measured with a 50°C chopped blackbody source of 10 mm circular aperture. (a) is from a typical detector with a measured FWHM of 13 mm at a distance of 3.5 m. (b) is from a detector on the right edge of the array with a broader 15 mm FWHM beam and some strong aberration. Both beams show some off-centre low-level response attributable to light leaks from neighbouring feed-horns.*

columns in the image frames however these can be digitally filtered or interpolated over in real-time to improve the overall image quality.

2.4 Performance

The optics system was tested with a measurement scheme based on raster scans of a chopped 50°C blackbody across the object plane. Maps of the beam profiles for each working detector were made down to a 25 dB signal to noise level. A typical beam (Figure 2.7a) is approximately Gaussian and the full width at half maximum (FWHM) is 11 mm at 3.5 m after deconvolution of the 10 mm diameter source aperture. This provides a resolution close to that expected for a diffraction limited system in this configuration, although, some channels show mild broadening and aberrations (Figure 2.7b), particularly at one edge of the focal plane. There is also some indication of localized leakage from adjacent feed-horns at a level typically less than 5-10% of the main beam level.

The operational yield of the current detector array is 85% with the majority of unusable pixels suffering from resonator overlap due to non-uniformities of the thickness of the aluminium film. Aside from this resonator clash there is no indication of any other electromagnetic cross coupling between resonators down to the measured 25 dB level.

A noise power spectrum for a typical detector channel sampled at the maximum
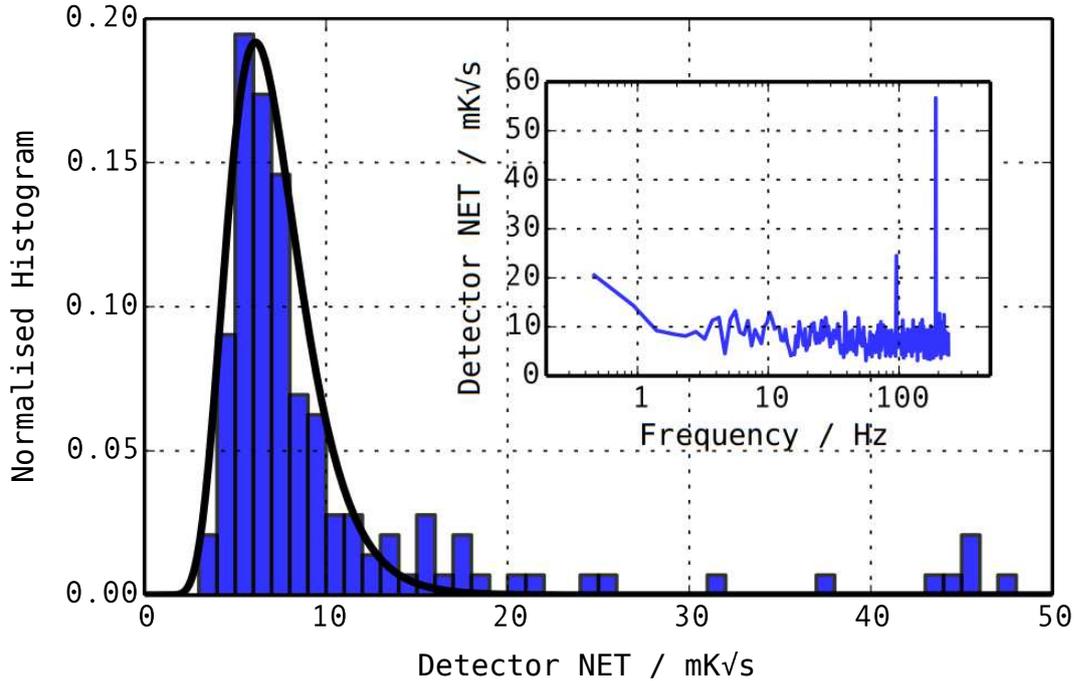
Figure 2.8: *A normalized histogram of noise equivalent temperature over each of the N = 152 detectors, sampled at 100 Hz where white noise is predominant. The distribution is well approximated by a log-normal function (black curve), the modal value of which is 6.1 mK√s. (Inset) An example noise spectrum as measured from a typical detector timeline across the fully sampled post-detection bandwidth.*

rate of 477 Hz is presented in the inset to Figure 2.8. The spectrum shows white noise down to ∼1 Hz which is typical across all of the detectors. The excess below this knee frequency is attributed to the warm electronics system, as are the spurious components at 95.5 Hz and 191 Hz. These unwanted narrowband features are digitally filtered from the detector timelines prior to image generation. The filters are implemented as fifth-order, Butterworth bandstop filters that operate on the timelines in the time domain on a frame-by-frame basis.

The distribution of NETs sampled at the white noise frequencies (sampled at around 100 Hz) is indicated in the histogram in Figure 2.8. The distribution is approximately log-normal with a peak NET value of 6.1 mK√s, a factor of 2 higher than the expected limit from photon noise in this system. The excess is thought to be due to stray infrared radiation leaking from the 4 K stage.

The constraint set by the scanning mechanism and the higher than expected noise currently limit the update rate to 2 frames per second for an NEΔT of 0.1 K per frame with the camera in its present configuration. Figure 2.9 shows a single frame taken from a combined "three-colour" video. The sensitivity is clearly sufficient to identify objects that are invisible to thermal NIR cameras and standard digital video cameras.
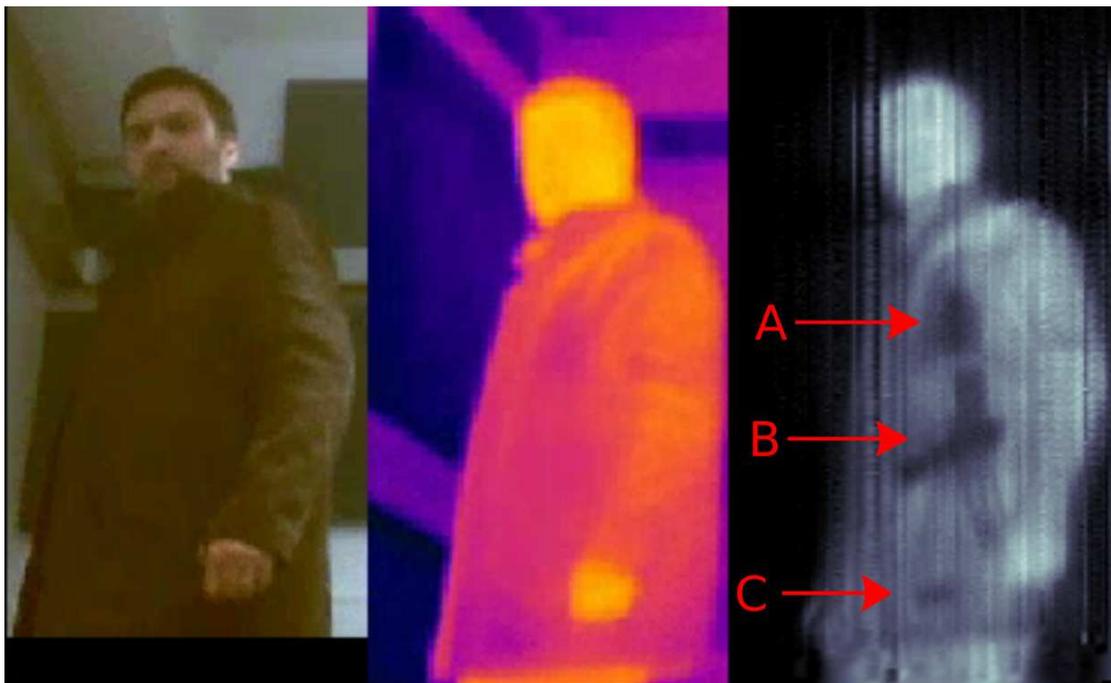
Figure 2.9: *A snap shot from a 2 FPS video in which the 350 GHz frames (right) were displayed simultaneously with frames from a standard webcam (left) and a thermal NIR camera (centre). Objects such as (A) a wallet, (B) an air pistol, and (C) some loose change, are hidden by the high opacity of the coat at higher frequencies but become apparent at 350 GHz.*

2.5 Discussions / Future Development

In most respects, the camera presented here has achieved the required specifications. The presence of parasitic optical loading on the detector array limits the noise performance so that full video rate could not be achieved – even if a faster field modulation system was employed. However, a second generation system could overcome this in a number of ways. For example, by utilizing a reflective optics approach, especially one with a cold Lyot-stop, such as that used in the BLAST-TNG telescope [Galitzki et al., 2014]. This would help to inhibit stray light loads on the detectors and also, in this case, eliminate the requirement for the feedhorn coupling.

Additionally, the field scanning mechanism of the present system is purely linear and thus does not employ any cross linking between detector channels. As such, the video frames suffer from vertical striping due to broken/noisy detectors and low frequency gain fluctuations between individual detectors. Transitioning to a dual-axis circular or Lissajous style scanning strategy would remedy this and is an advisable approach for any future system.

A general purpose instrument similar to that presented here would benefit from a modular (rather than fixed) optics system. Providing an additional image plane located externally to the cryostat would enable fast turnaround between a variety of application specific imaging formats without the need for any modification to the cryogenic platform.

2.6 Conclusions

Kinetic inductance detectors originally developed for far-infrared astronomy are now suitable for use in a range of applications requiring high sensitivity and/or fast mapping of objects at terahertz frequencies.

The instrument presented here mimics stand-off imaging systems for the detection of concealed items but could easily be transformed for other applications by modification of the optics platform. This LEKID based system operates close to the ideal photon noise limited sensitivity estimated with Equation 2.1 and is comparable in performance to the latest passive THz imaging systems.

The development of larger KID arrays is ongoing and next generation instruments will benefit from order of magnitude increases in detecting elements with no considerable penalty in array fabrication or readout complexity.

Chapter 3

The Optical System

In any camera or imaging device, the optics system is responsible for directing the light emitted or scattered by some object (within a well defined field of view, FOV) onto some focal plane so that it may be viewed and/or measured. This chapter presents the design, implementation, characterisation, and overall performance of the optics system employed in our kinetic inductance camera.

It is worth noting here that, after the requirements for the system had been decided (see chapter 2), two important optics design issues arose. Firstly, given the modest detector count (60 pixels at 150 GHz and 152 pixels at 350 GHz), what kind of scanning strategy should be employed to fully sample the FOV. Secondly, given limited freedom in the positioning of the detector focal plane array within the cold stage of a pre-existing cryostat, where on the cryostat should the window be installed and thus, what imaging optics could be employed to couple light though to the focal plane array.

As with many of the design choices in this project, the solutions were chosen to minimise the cost and complexity of the instrument – it is, after all, a demonstration system. Consequently, these solutions – although they were effective in the end – fundamentally limited the overall performance of the system.

The author acknowledges supporting contributions from a number of people including Dr P. Hargrave for optics design and modelling; Prof P.A.R. Ade and Dr R. Sudiwala for support with preliminary characterisation measurements; Dr W. Granger for designing the scanning mirror system; and, Prof C. Tucker and Prof. P.A.R. Ade for the band-defining and thermal filters.

3.1 Design

The focal plane array consists of a linear array of detectors that see out of the base of cryostat through a large HDPE window, via an oscillating fold mirror located beneath that scans the FOV. The detector arrays were designed to Nyquist sample the FOV in the (horizontal) x-direction while the mirror oscillates at a sufficient rate to Nyquist sample the the FOV in the (vertical) y-direction. Images are reconstructed after sampling over
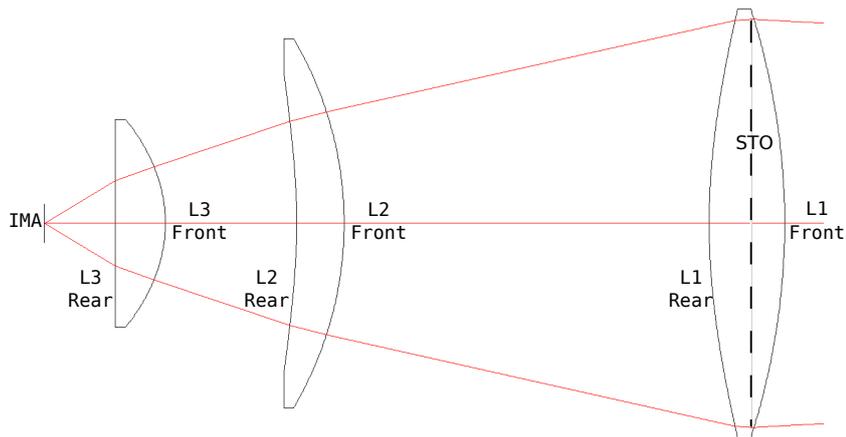
(a) *Lens design and configuration.*
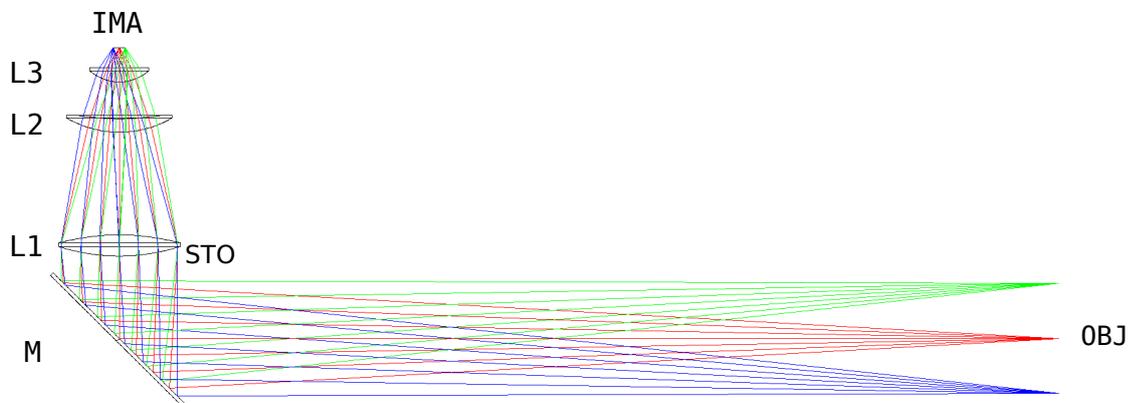
(b) *Ray traces from the object plane to the image plane.*

Figure 3.1: *Views of the nominally prescribed optics model. See table 3.1 on page 38 for the parameters defining the precise shapes, sizes and positions.*

one or more half-periods of revolution of the mirror. See figure 2.1 on page 23 for the schematic of the scanning system.

The mirror, M, is constructed from a thin, flat sheet of polished aluminium supported by aluminium strut profile and connected by a steel arm at one end to a servomotor driven circular cam. An absolute rotary encoder is mounted on the central axle of the mirror to report the orientation to the data acquisition system.

A half-meter primary aperture is required to achieve the desired diffraction-limited spatial resolution of 1 cm at a distance of 4-5 m ($0.13°$ angular resolution) when operating in the 350 GHz (850 µm) band. The resolution increases linearly with wavelength so in the 150 GHz (2 mm) band a 2.5 cm resolution is expected at the same distance, or $0.3°$. A biconvex HDPE lens, L1, was chosen but manufacturing processes limited its diameter to 0.44 m.

An additional feature included in the design was the ability to adjust the focal length of the system in order to re-focus on objects that were slightly closer or further away than the standard 4-5m distance (in the z-direction). To that end, a secondary HDPE lens, L2, with adjustable height was included behind the primary lens.

A third HDPE lens, L3, is included inside the cryostat in order to keep the optics chain compact. The lens mounts onto the end of the 50 K radiation shield and required an extension to the vacuum can in order to fit it in.

This triplet of lenses was optimised using the Zemax modelling software and is presented in figure 3.1 on the preceding page. Table 3.1 on the following page is taken from the Zemax file representing the nominal system. It contains all the parameters required to describe each surface in the optics system starting from the object, OBJ, and ending at the image, IMA, including the primary aperture stop, STO.

The low effective image-space f-number, $f/0.9$, and the high primary aperture diameter make for quite a fast convergence of rays from L3 to the image at the focal plane. This, in combination with the large window size and a lack of any pupil or cold stop, makes the system especially vulnerable to stray light. Detector response and optical resolving power are shown to be improved by the inclusion of optical baffles within the cold section of the optics chain and by horn-coupling to the detector array. More detail on this is included later in this chapter.

3.2   Lens manufacture

The first measurements with the camera were taken shortly after the initial assembly of the optics system. The 150 GHz detector array was in position within the cold cryostat, the multiplexing readout system was operational, and the data acquisition software was presenting detector response timelines live on the computer display. Characterisation of the the system began with the controlled scanning of a bright source across the field of view. This would help to confirm such parameters as the focal length and the spatial resolution. This was, however, when we realised that something was not quite right with

Table 3.1: *The nominal optics design as taken from the Zemax model.* R_c *is the radius of curvature,* T *is the thickness;* Glass *is the refractive index,* D *is the diameter,* K *is the conic constant, and* $X - $Tilt *is the axial tilt.*

Surf.	Comment	R_c/mm	T/mm	Glass	D/mm	K	X-Tilt/°
OBJ		∞	3700.0		800	0.000	0
1	Coord. Break		0.0				+45
2	Fold Mirror	∞	0.0	MIRROR	338.3	0.000	0
3	Coord. Break		-300.0				+45
4	L1 Centre	∞	-33.6				0
STO	Aperture Stop	∞	33.6		420	0.000	0
6	L1 Front	704.2	77.0	1.52	440	-1.962	0
7	L1 Rear	-690.4	295.5		440	-10.072	0
8	L2 Front	384.8	47.0	1.52	380	-0.113	0
9	L2 Rear	636.8	168.8		380	-46.996	0
10	L3 Front	147.7	51.2	1.52	214	-0.597	0
11	L3 Rear	∞	72.9		214	0.000	0
IMA		∞	–		77	0.000	0

the optics. See figure 3.2 for the measured beam profiles for a selection of detectors.

The scans were taken at the position of greatest response for the central pixel. This was found to be 2.9 m from L1, considerably closer than the expected focal length of 3.7 m. Furthermore, the angular widths of the beams were typically 0.5°, almost a factor of 2 broader than the 0.3° expected for the 150 GHz band. And finally, the beams showed strong off centre anomalies or 'ghosts'.

To troubleshoot these problems, the manufactured lens dimensions were checked against the design specification. The lenses and the cryostat window had been cut from single blocks of HDPE by CNC lathe at TK Instruments Ltd. and were checked by CMM
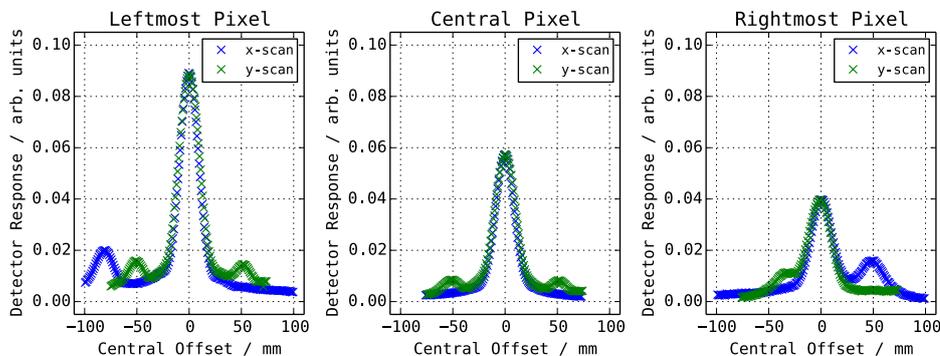
Figure 3.2: *Initial results from scanning a bright source across detector beams in the horizontal (x) and vertical (y) directions, for three KIDs spread across the 150 GHz array*

Table 3.2: *Nominal versus manufactured lens parameters.* R_c *is the radius of curvature,* T *the thickness, and* K *the conic constant.*

Surface	R_c/mm		T/mm		K	
	Nominal	Manufact.	Nominal	Manufact.	Nominal	Manufact.
L1 Front	704.2	698.55	77.0	77.12	-1.96	-10.243
L1 Rear	-690.4	-708.30	295.5		-10.07	-0.590
L2 Front	384.8	373.22	47.0	48.36	0.11	-0.038
L2 Rear	636.8	621.66	168.8		-46.99	-45.221
L3 Front	147.7	148.21	51.2	50.97	-0.60	-0.609
L3 Rear	∞		72.9		0.00	

Table 3.3: *Re-optimised lens positions to compensate for errors in manufacture.*

Distance/mm	Nominal	Compensated
L1 Rear to L2 Front	295.487	384.650
L2 Rear to L3 Front	168.763	122.471
L3 Rear to Image	72.857	72.857

(coordinate measuring machine) in the department. Subtle but significant differences were found, as indicated in table 3.2.

Fortunately, a suitable re-optimisation of the optics model based on the newly measured lens parameters was found to compensate for this error. The simulated performance of the re-optimised system in terms of focal length and beam size was sufficiently close to the nominal performance that the project was able to move on without requiring any modifications to the lenses. See table 3.3 for the re-optimised configuration.

Before continuing back with the camera, a preliminary analysis of the optics system was performed externally from the cryogenic system. Measurements would be made with a single, fully optically characterised bolometer in place of the KID array. This would confirm the optics performance against the newly optimised specification, and hopefully, shed some light on the cause of the ghosts by eliminating any degeneracies that might have been introduced by the cryogenic and/or detector configurations.

3.3   Preliminary evaluation

3.3.1   Measurement setup

A measurement system was set up to test the optics under various conditions. The optical components were arranged outside of the cryostat by laser alignment as per the re-optimised Zemax model and a bright source was scanned across the object plane with

a receiver at the image plane to measure the optical response. Figure 3.3a presents a basic schematic of the setup.
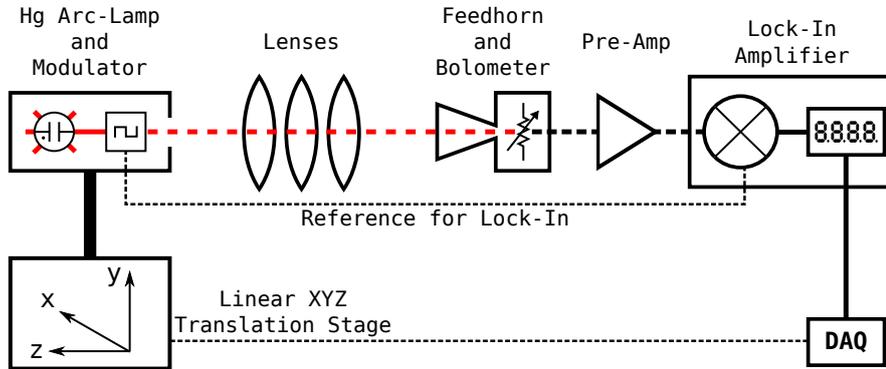
(a) *A schematic of the experimental setup.*
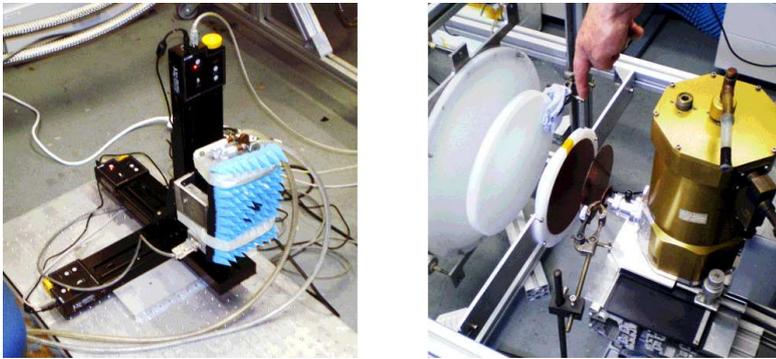
(b) *Photographs of the experimental setup. The source (left) and the receiver (right).*

Figure 3.3: *The measurement setup for the preliminary characterisation of the re-optimised optics system.*

The receiver, a 4 K InSb hot electron bolometer, was housed in a portable liquid He^4 cryostat with a coupling feedhorn of 1.6 mm aperture and a 350 GHz optical band-defining filter. The cryostat was placed on a manual translation stage for measurements as a function of detector position across the image plane. See the photo on the right of figure 3.3b. The bolometer readout was performed with a standard, low noise readout circuit and preamplifier.

The source was a Hg vapour arc-lamp with a blackbody temperature of $\sim$1500 K and a 10 mm aperture chopped at 28 Hz by a vibrating vane modulator. Its position was variable in three dimensions with a triple stage 'XYZ' scanner (3 $\times$ single Thorlabs linear scanners). The remotely controlled scanners were stepper motor driven with a maximum throw of 300 mm and an accuracy no better than 5 µm in each dimension. The exterior of the arc-lamp casing was covered in high-absorption foam to minimise reflections. See the photo on the left part of figure 3.3b.

The bolometer output and the chopper reference signal were fed into a Stanford Research Systems SR810 DSP Lock-In Amplifier and the demodulated response was digitised using a National Instruments ADC. A Labview Virtual Instrument recorded

the XYZ positions and bolometer responses into ASCII text files for analysis.

Individual measurements typically required an initial alignment of the source to the centre of the receiver's beam. This was performed by manually adjusting the XYZ scanner until the detector response (displayed digitally on the SR810) reached a maximum. The object plane would then be scanned in two perpendicular sections (x and y) orthogonal to the optical axis (z), with each section crossing the centre of the beam.

Where necessary, the heights, widths and centres of the beam profiles were estimated by fitting to a Gaussian function of the form

$$g(x) = A_0 + A\, \exp\left(-\frac{1}{2}\left(\frac{x - x_0}{\sigma_x}\right)^2\right),$$ (3.1)

where x is the distance along the direction of the scan, x_0 is the central offset, A_0 is the baseline level, A is the height or amplitude, and σ_x is the dispersion. The term 'width' shall be used refer to the full width at half maximum, which is calculated from the dispersion by

$$\mathrm{FWHM} = 2\sqrt{2 \ln 2}\ \sigma_x.$$ (3.2)

A suite of tests, each involving multiple repetitions of this measurement scheme, was carried out. The results are presented below.

3.3.2 Diffraction limited resolution and depth of field

The aim of the first set of tests was to characterise the beam shape and size with the receiver centred in the image plane at $(x', y') = (0, 0) = (x_0', y_0')$. The source was positioned, initially, at the location of maximum response in the object plane, defined here as the origin, $(x, y) = (0, 0) = (x_0, y_0)$. The focal length, z_0, was found to be $3.35\,\mathrm{m}$ normal to the front surface of L1.

Scans were completed across $x = x_0 \pm 50\,\mathrm{mm}$ and $y = y_0 \pm 50\,\mathrm{mm}$ with $1\,\mathrm{mm}$ steps and repeated at $z = z_0 \pm 150\,\mathrm{mm}$ to estimate the depth of field of the system. All the beam profiles were normalised against the peak response of the $z = z_0$ case. The signal to noise (peak response over standard deviation at the baseline level) in the $z = z_0$ case was measured as approximately 2500 or, equivalently, $34\,\mathrm{dB}$.

Gaussians were successfully fitted to the beam profiles and the measured profiles and fits are displayed in figure 3.4 along with residuals as percentages of the measured values. The normalised amplitude and beam FWHM of each scan is presented in table 3.4.

The results were promising, with $\sim 13\,\mathrm{mm}$ wide in-focus beams corresponding to an angular resolution of $0.22°$. A better estimate of the minimum resolvable feature size (i.e. the width of the point spread function) was found by the deconvolution of the beam profiles with the circular aperture of the source ($10\,\mathrm{mm}$ diameter). This resulted in $11\,\mathrm{mm}$ at $3.35\,\mathrm{m}$ or $0.19°$, in line with but not exactly matching the expected limit from diffraction of $8\,\mathrm{mm}$ at $3.35\,\mathrm{m}$ or $0.14°$.

Depth of field is loosely defined as the distance between the nearest and furthest
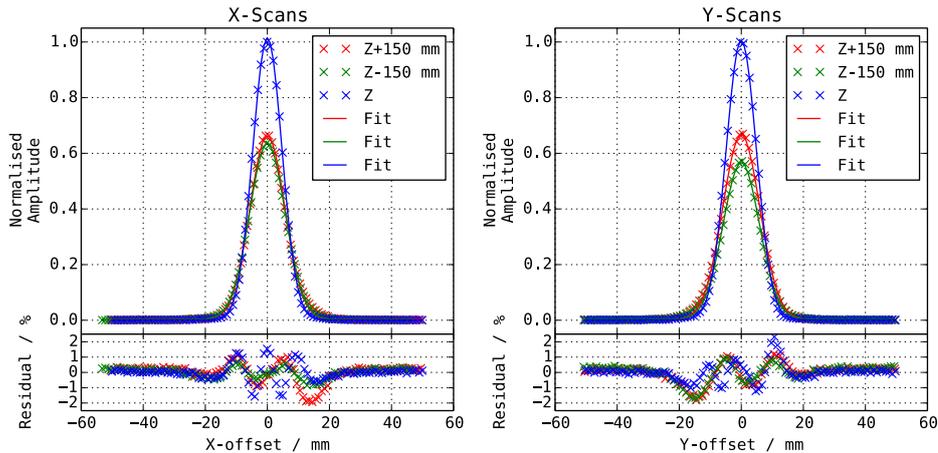
Figure 3.4: *Normalised detector response and Gaussian fits as source is scanned through image plane in x-direction (left) and y-direction (right). Blue is with the source in focus at $z = 3350\,mm$. Red and green respectively are with the source $\pm 150\,mm$ out of focus.*

Table 3.4: *Preliminary beam profile parameters measured with the source $3350\,mm$ from the front surface of L1.*

Z-offset/mm	Gaussian fit to X-scans		Gaussian fit to Y-scans	
	Amplitude	FWHM/mm	Amplitude	FWHM/mm
−150	0.66	13.2	0.66	14.1
0	1.01	11.1	1.01	11.5
+150	0.63	13.5	0.56	13.8

objects that remain in sharp focus in an image. At $z = z_0 \pm 150\,mm$, the beam widths increase by 20%. The camera requirements only specify a sufficient depth of field to image a human target and, without a well defined limit on an acceptable circle of confusion, the defocus across this $300\,mm$ range is considered, rather subjectively, to be acceptable for this system.

3.3.3   Field of view and plate scale

The aim of the second set of tests was to characterise the instantaneous field of view of the optics. The responses that would be measured by the linear array of detectors are simulated by measuring beam profiles as a function of receiver position in the x direction of the image plane. The x and y scans are performed with the receiver between $x' = -50\,mm$ and $x' = +40\,mm$ with $10\,mm$ steps and $y' = 0\,mm$. For reference, the maximum distance between the outermost detectors in the focal plane arrays is at most $80\,mm$.

Summaries of the results are included in figure 3.5 where the beam widths, amplitudes and central x-offsets are plotted against receiver offset.
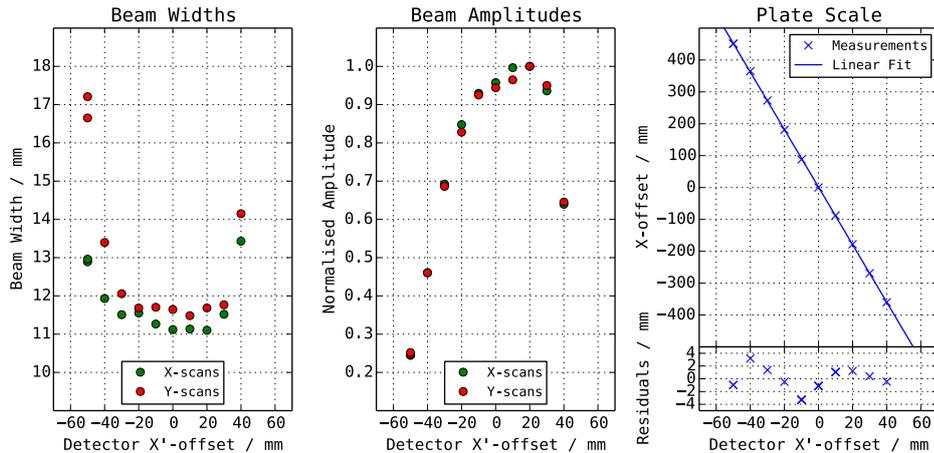
Figure 3.5: *Beam parameters measured with the receiver swept horizontally across the image plane to mimic the parameters expected across the KID focal plane array.*

The response of the beam widths to the receiver offset was mostly flat out to $x' = \pm 30$ mm however, the widths increase very rapidly beyond ± 40 mm. The response of the beam amplitudes was much less flat and not symmetric about the optical axis. Beam amplitudes greater than 90% of the peak value were observed in the range $(-10 \le x' \le +30)$, falling steeply outside of this range.

The calculated plate scale was 9.02, meaning that a 1 mm offset in the image plane corresponds to a 9.02 mm offset in the object plane. Plate scale is a measure of comparison between the angular size of an object viewed through a telescope and the linear size of the image presented at the focal plane. The angular object size is converted to a linear object size and the scale is calculated as the gradient of a linear fit of the detector x'-offsets versus the x-offsets of the centres of the measured beams. Given that the detector arrays have up to an 80 mm extent in the image plane, the horizontal extent in the object plane would be $80 \times 9.02 = 722$ mm. This is somewhat smaller than the 1 m specified in the requirements, but is still sufficient to image a person standing in the field of view.

3.3.4   Refocusing optics

The added feature of an adjustable secondary lens enables the refocussing of the camera onto objects that are offset from the nominal object plane. The nominal distance from the rear surface of L2 to the front of L3, taken from table 3.3 is 122 mm. Beam profiles were measured with L2 at this position and with ± 5 mm and ± 10 mm offsets. Data from the +5 mm test is unavailable.

In figure 3.6, variations in the in-focus position of the object are displayed on the left and variations in beam widths are on the right. Beam profiles were measured in-focus at $z = z_0$, as well as at $z = z_0 \pm 150$ mm out of focus. No negative effect on the shape or size of the beams was seen.
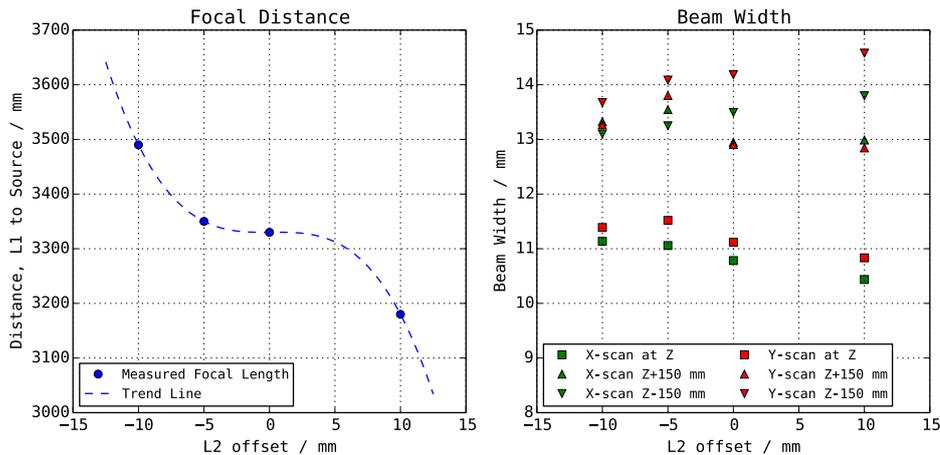
Figure 3.6: *Refocusing optics. Left: Variations of focal distance (left) and beam widths (right) upon shift of refocusing lens, L2.*

Moving L2 through a full 20 mm caused the focal length to vary by only 300 mm, so, further refocusing could be applied before any significant performance degradation would occur, should this ever be required.

3.3.5 Additional optical components

A final set of tests was carried out to check the performance of the optics with the addition of various components that would be present in the full system. The series of cumulative additions of components went as follows:

1. The basic optics system, as characterised in this section.
2. The cryostat window, requiring a +7 mm shift in L2 to compensate for defocus.
3. The 21 cm^{-1} low pass edge filter.
4. A copper annulus to mimic the detector array holder, placed 4 mm in front of the receiver.
5. The 15 cm^{-1} low pass edge filter, placed (a) 50 mm from L1, towards the receiver, (b) 15 mm from the receiver and (c) 1 mm from the receiver.
6. The 350 GHz (850 μm) band pass filter. Five scans were averaged to improve the signal to noise.

The amplitudes and widths of the x and y beam profiles are presented in table 3.5. Aside from the expected dimming introduced by some of the components, and some very minor beam broadening, no further issues became evident upon the completion of this test.

3.3.6 Summary

In summary, manufacturing errors in key components of the optical system meant that the performance in the nominally prescribed configuration did not meet the required

Table 3.5: *List of tests with additional optical components*

Test	Comment	Normalised Amplitude		FWHM / mm	
		X-Scans	Y-Scans	X-Scans	Y-Scans
1	Base system	1.00	1.00	11.26	11.76
2	+ cryostat window	0.85	0.86	11.36	11.79
3	+ 21 icm LPE	0.53	0.53	11.71	11.89
4	+ annulus	0.52	0.52	11.70	11.89
5a	+ 15 icm LPE (50mm)	0.38	0.38	11.96	12.19
5b	+ 15 icm LPE (15mm)	0.38	0.38	11.96	12.20
5c	+ 15 icm LPE (1mm)	0.39	0.39	12.11	12.30
6	+ 850 micron BP	0.04	0.04	11.89	11.65

specification for the camera. After a suitable re-optimisation was found, the key components were arranged externally to the cryogenic system, and an analysis was undertaken using a single pixel bolometric detector and a scanning optical blackbody source.

Aperture deconvolved beam widths, measured in the object plane at the new nominal focal length of 3.35 m, were found to be around 11 mm on the central axis. The beam widths expected from diffraction alone ought to be 8.3 mm. Beam profiles measured in the vertical direction were consistently but not significantly broader than those in the horizontal. The results were repeatable with the receiver positioned ±30 mm about the central axis in the image plane, however, the performance dropped off sharply outside of this range. The peak amplitudes of the beams dropped off quickly when the receiver was positioned beyond ±20 mm centred at +10 mm from the central axis in the image plane. The spatial distribution of KIDs in the focal plane arrays ranges between ±40 mm about the central axis so some optical attenuation and distortions are expected at the extremes of the array.

The beams were broadened by only 20% when measured ±150 mm out of focus, this indicates a sufficient depth of field as to allow for the imaging of a human target without significant blurring.

Refocusing of nearby or faraway targets was confirmed in a 300 mm range centred about the nominal focus without any measurable performance degradation. It is expected that this range could be extended, in practice, before any negative effects emerge.

The presence of additional optical components such as the band defining filters, the cryostat window and a copper ring to mimic the detector array housing, caused roughly a 10% broadening of the beams as well the expected dimming from optical absorption.

In all cases, beam profiles were successfully fitted by Gaussian functions with residuals at levels typically lower than 2%. In no cases were any distortions in the beam shapes noted nor was there any sign of ghosting as seen in the very first run of the fully integrated, non-optimised system.

3.4 Full system evaluation

The following sections present the evaluation of the camera optics in the fully integrated configuration based on the preliminarily confirmed, re-optimised model. Full details of the measurement setup and the methodologies for data processing and analysis are given first. Then the results and conclusions from the tests in the 150 GHz and 350 GHz optical bands will be presented.

3.4.1 Measurement setup
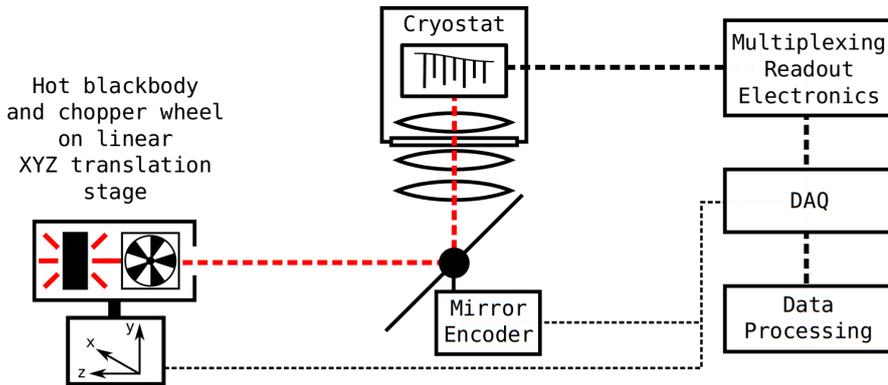
(a) *A schematic of the experimental setup.*
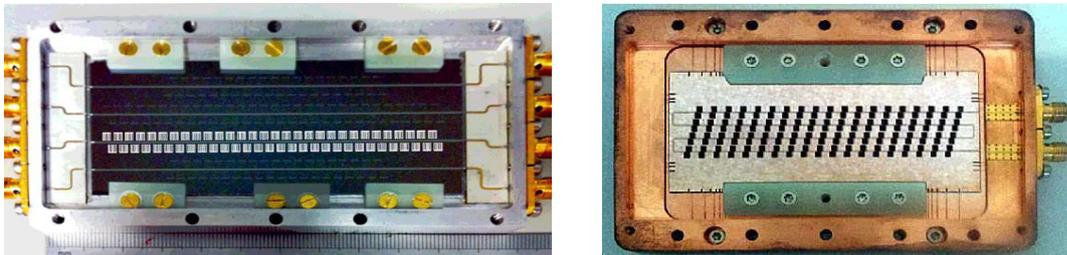
(b) *Photographs of the detector arrays for the 150 GHz band (left) and 350 GHz band (right). The detectors are highlighted to indicate the spatial arrangement.*

Figure 3.7: *Measurement setup for the final characterisation of the fully incorporated, re-optimised optics system.*

A similar strategy to that of the preliminary evaluation was employed, the key difference being the inclusion of actual KID arrays for measurement. The simultaneous multi-pixel readout enabled an efficient characterisation of the system, however, without any knowledge of the detector responsivities or absorption efficiencies, uncertainties would remain in the absolute optical power reaching the focal plane.

The original Hg arc-lamp that was used as the blackbody source for the preliminary measurements was unavailable during this characterisation, so a replacement system was put together. This consisted of a small metal plate painted black with SiC infused epoxy resin and heated by a series of resistors. The plate was warmed up to 125°C, as measured by FLIR thermal infrared camera, when supplied with 14 W of electrical power.

A 10 mm diameter aperture stop was placed in front of the hot plate along with a room temperature chopper wheel to modulate the source. This ΔT of 100 K was an order of magnitude lower than the 1200 K provided by the arc lamp so the optical power received would be, accordingly, also an order of magnitude lower.

A choice was made not to incorporate a reference signal from the new chopper wheel into the housekeeping systems of the multiplexing electronics. The resources required to build such an interface were seen as unnecessary given the convenience of a software based demodulation.

This new source was mounted to the triple stage linear scanner and the acquisition system for the XYZ position information was modified to transmit the values over UDP to the multiplexing electronics housekeeping systems. This resulted in a ~15 ms latency between detector data acquisition and XYZ acquisition - a small but significant factor that required careful consideration when analysing data from the scans.

Initial x and y scans indicated that the camera's focal length and beam size were considerably improved by applying the re-optimised configuration. However, the ghosts were back and they were present in all of the measured beams, in some cases at up to 50% of the main beam level. To gain further insight into the cause of these ghosts, a new scanning strategy was devised. One in which the source was raster scanned across the x, y plane to map out the entire instantaneous field of view of the camera. The XYZ scanner has a maximum range of 300×300 mm, so three separate side by side scans were required to be taken and stitched together in order to create a full map for every detector in each of the focal plane arrays.

A schematic of the measurement setup is given at the top of figure 3.7. Also included in the figure are photos of the detector arrays with the individual pixels highlighted to indicate their spatial distributions. Details of the map making process and the digital demodulation, are described below.

3.4.2   Data processing

Multiple processing stages were required in order to generate and analyse the 2-D maps of the detector beam profiles of the fully integrated camera. Firstly, the detector I and Q data had to be converted into detector response timelines. These timelines then had to undergo filtering to lock-in to the optical signal – as modulated by the chopper wheel. Then, using the XYZ scanner position data, these demodulated optical responses had to be binned into appropriately sized 2-D arrays. Finally, any gaps in these undersampled arrays were interpolated between. The beams were then available for inspection and, where appropriate, the beam parameters could be estimated by fitting Gaussian functions.

Timelines

For the purposes of this chapter, details of the readout electronics systems and 'raw' detector timelines will be skipped over. It suffices to consider that for an array of N

detectors, N files can be generated, each one containing regularly sampled values of one detector's response (measured in Hertz) over time. The responses should be assumed to vary linearly with respect to the optical power incident upon each detector.

An example of such a timeline is displayed in blue in the upper part of figure 3.8a. The modulated source can be seen passing through the beam of the detector over the course of a number of raster scan sub-scans. A close up of a single pass through a roughly Gaussian looking beam is inset to the figure. A further close up of the timeline reveals the square wave modulation imparted by the optical chopper wheel.

The power spectral density (PSD) corresponding to the timeline is given in green in the lower part of the figure. The PSD shows a strong component at the frequency of the optical chopper (28 Hz), as well as harmonics from the squareness of the chop, and other noise.

The XYZ positions are acquired from the linear scanning system at the same rate as the detector responses.

Lock-in

Demodulation of the chopped signal is presented in the figures 3.8b and 3.8c. The process involves an initial baseline removal to offset the effects of any detector drifts that might occur over the duration of the raster scans.

The baselined timelines are then passed though a zero-phase Butterworth bandpass filter[1], designed to let the 28 Hz chopper signal pass but to stop everything else. The rejection can be seen down to -200 dB in the implementation shown by green line in the PSD in figure 3.8b. The zoom in of the timeline in blue shows a relatively pure sinusoidal modulation compared to the square wave present prior to bandpass filtering.

Finally, the absolute value of the bandpassed timelines are taken and lowpass filtered, again with zero-phase Butterworth IIR filters. The lowpass cut-off frequency is designed such that the 28 Hz modulation is smoothed out without smoothing over the larger structure of the beams.

The end results of this process are timelines of detector response to the 28 Hz chopped optical signal which are ready to be binned into maps.

Map making

The whole beam mapping process was time consuming. The XYZ scanner could perform a full 300 mm sub-scan in, at fastest, 30 seconds, effectively providing one data sample per 20 μm, at a 500 Hz sampling rate, in the x-direction. Depending on the spatial extent and resolution desired in the y-direction, anywhere from 50 to 100 sub-scans

---

[1]Butterworth filters were chosen to keep the passband response as flat as possible. Since the minimal phase and hence the minimal group delay of a Butterworth filter is not naturally zero, each stage of filtering was completed in two passes, forward-and-backward, to cancel out the group delays introduced by each pass. Without a zero group delay, the detector responses would be shifted relative to the XYZ positions, meaning further processing would be required before the binning of any maps.
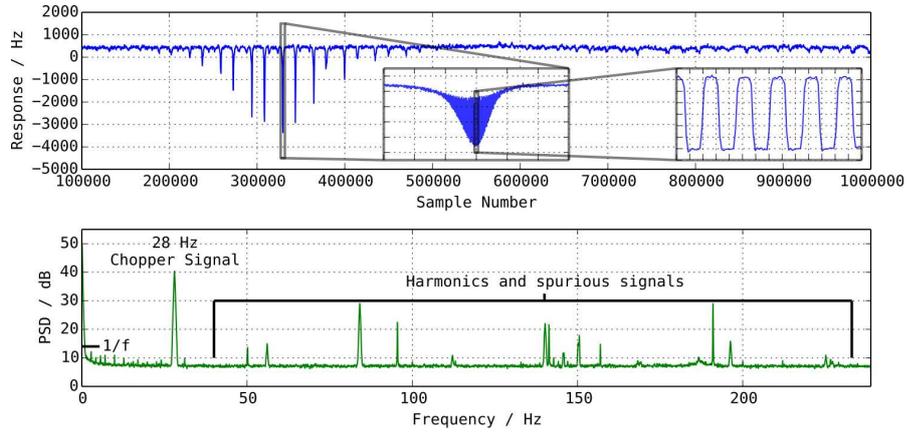
(a) *An unprocessed timeline of detector response over the course of the beam mapping process (upper), and the corresponding power spectrum (lower).*
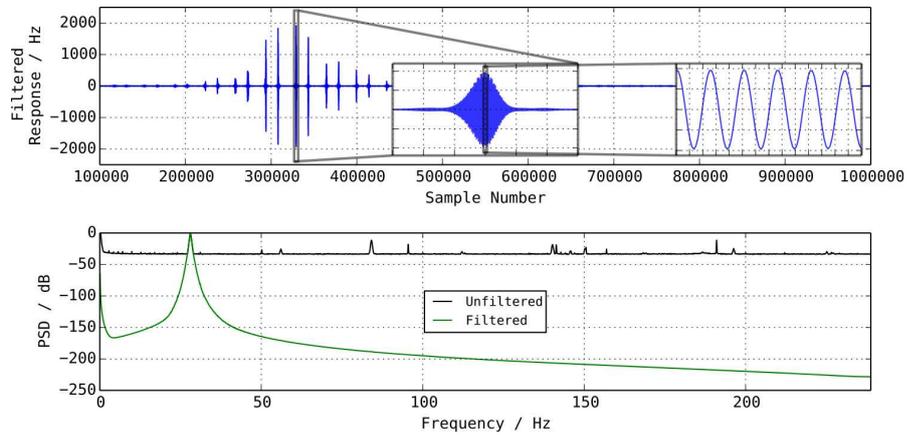
(b) *Locking-in to the modulated signal: the band-pass filtered response and spectrum.*
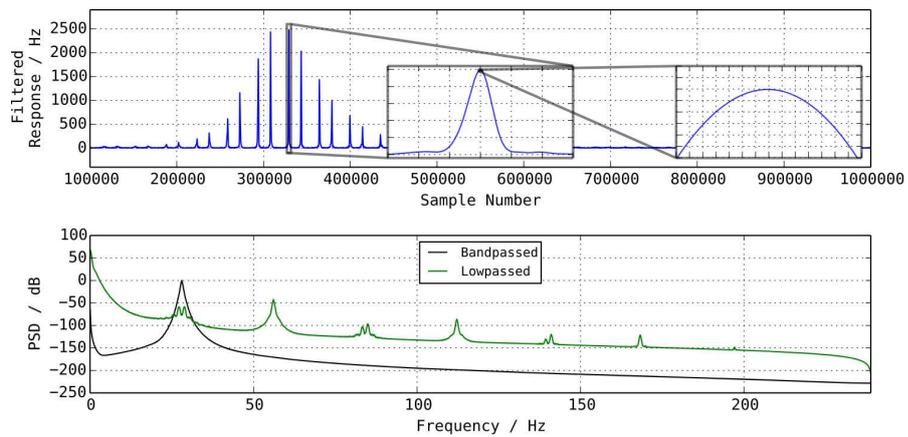
(c) *Finally, a lowpass filter on the absolute value of the bandpassed timelines. The timeline is ready for map binning.*

Figure 3.8: *Map making, part 1: Timeline filtering for the digital lock-in.*
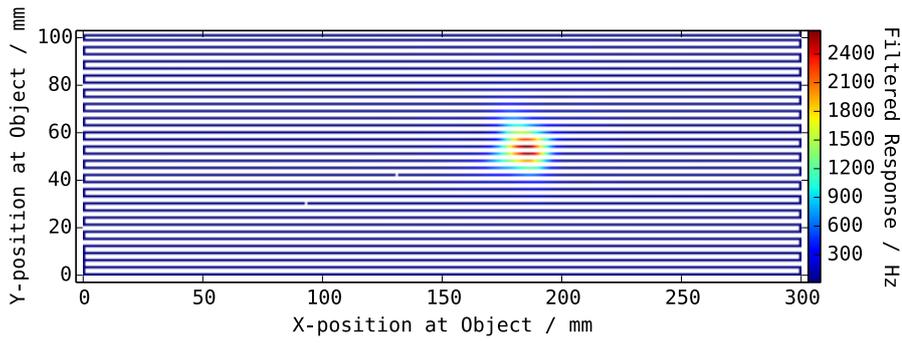
(a) *A section of a raw map. Note the raster scan path.*
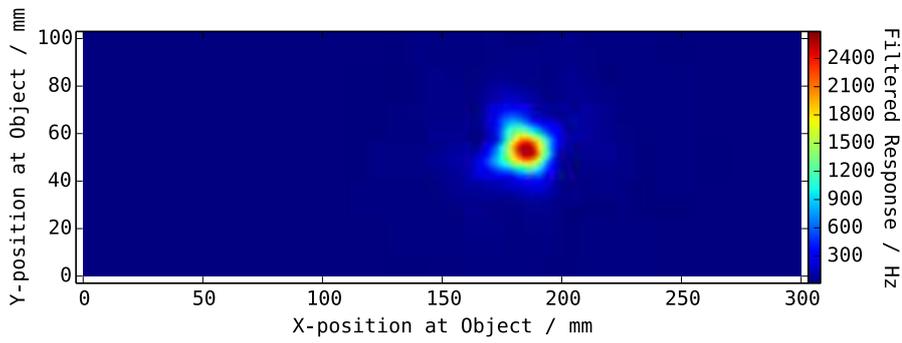
(b) *A section of an interpolated map with all gaps filled.*
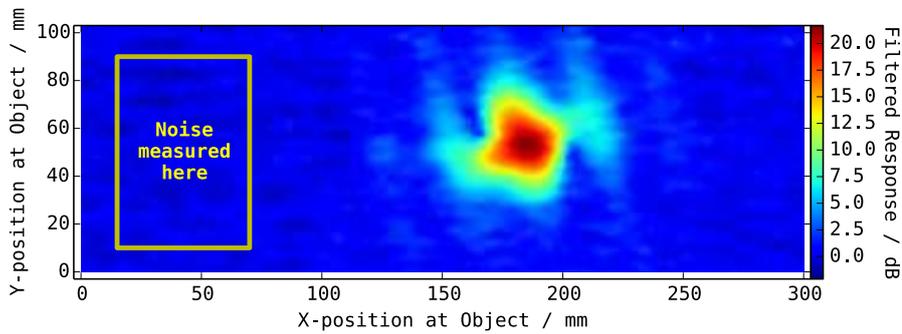
(c) *A logarithmic colour scale shows low level structure. The background noise in the map is measured in the given region.*
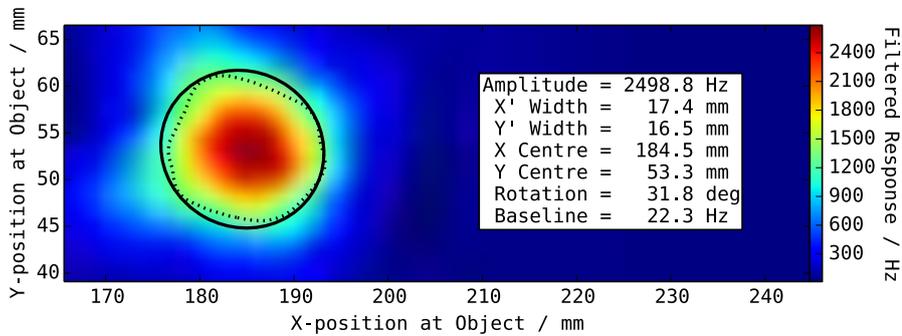
(d) *Beam parameter estimation with half power contours of the measured beam (dotted line) and the Gaussian fit (solid line).*

Figure 3.9: *Map making, part 2: Map binning, interpolation, and analysis.*

would be required to sample the beams well. Furthermore, mapping out each and every detector beam in the far field required three full raster scans due to the limited throw of the linear scanners. That makes nearly three hours to fully map out the field of view, including the time taken to manually offset the scanning system. The stitching of the three individual sets of data involved a simple concatenation of the filtered timelines. However, unevenness of the floor in the lab where the measurements were carried out, and small errors in the repositioning between raster scans led to inconsistencies between the reported XYZ positions and the actual spatial position of the source in the field of view. These were accounted for in the map making process with offsets and rotations applied to the XYZ positions.

The spatial extent in the object plane of the instantaneous FOV of the camera was roughly 700×50 mm in the 150 GHz band and 700×150 mm in the 350 GHz band. Scans were continued beyond this range to typically 900×250 mm, in attempt to catch the ghosts. Empty arrays were allocated with typically 900×250 elements, providing a 1×1 mm spatial bin size. One array was allocated for each detector channel, as well as an array for the number of samples collected in each bin.

For each detector channel, the map binning process involved a single loop over each timeline. The precise XYZ positions were converted into the indexes of the empty arrays, and a running sum of the detector response was calculated for each bin. Each array of the total summed response was, finally, divided by the array of the hit counts per bin to generate the response maps.

Given the number of sub-scans relative to the spatial extent and bin size, as well as the occasional missing data packet, the response maps were not necessarily fully filled. A qualitative analysis then becomes challenging as seen with the raw map in figure 3.9a, so an interpolation is performed over the remaining empty elements. Figure 3.9b presents an interpolated map.

Analysis

A lot can be inferred by simply looking at the maps, however as with the preliminary analysis, more quantitative analysis techniques are available. For example, by choosing a region in the map where there is a distinct lack of signal, the background noise level can be measured, and thus, a signal to noise ratio can be calculated for the map. See figure 3.9c.

Also as before, the widths and other parameters of the beams can be estimated by fitting Gaussian functions. A more advanced analysis can be performed on 2-D maps rather than on 1-D sections by fitting 2-D Gaussian functions to the data. These functions look like ellipses when taken in slices of constant amplitude and have a rotational degree of freedom. The form of the seven parameter 2-D Gaussian functions typically used in

this analysis is

$$g(x, y) = A_0 + A \exp\left(-\frac{1}{2}\left(\left(\frac{x'}{\sigma_{x'}}\right)^2 + \left(\frac{y'}{\sigma_{y'}}\right)^2\right)\right), \qquad (3.3)$$

where A_0 is a baseline level, A is the amplitude, and $\sigma_{x'}$ and $\sigma_{y'}$ are the dispersions in the major and minor axes, respectively. The positions in the rotated coordinate system, x' and y', are given by

$$x' = (x - x_0)\cos\phi - (y - y_0)\sin\phi \qquad (3.4a)$$

$$y' = (x - x_0)\sin\phi + (y - y_0)\cos\phi, \qquad (3.4b)$$

where ϕ is the rotation angle and (x_0, y_0) is the location of the centre of the function in the unrotated map coordinates, x and y.

An example fit is presented in figure 3.9d with half power contours of the data (dotted line) and the fit (solid line) overlaid onto the map. The fit is performed as a non-linear least squares minimisation based on the Levenberg-Marquardt algorithm. In some cases the 2-D Gaussian model may not be optimal for the results in hand, however, it can and will provide a reasonable approximation for at least some of the parameters.

In this analysis, the FWHM may be presented either as a pair of values (major and minor axis widths) each calculated from the dispersions as in equation 3.2, or as a single value that is the mean of these.

Finally, where necessary, any glitches that have propagated through the data processing pipeline can be masked out and ignored.

3.4.3   Performance at 150 GHz

The first run with the system in its full configuration was with the 60 pixel 150 GHz array. Initial scans had indicated strong ghosting in the beams and the preliminary optics evaluation ruled out the lenses as the cause of this. 2-D beam maps over the full field of view confirmed the presence of the ghosts and hinted that they might be due to reflections from the edges of the array packaging. Stacking the beam maps by computing the average over every map clearly confirmed this (see figure 3.10a) by showing that the ghosts were each part of a larger more more general structure of the same shape as the entrance to the array holder.

The system was then modified to reduce these stray reflections by arranging blackened baffles at the entrances of the various radiation shields, and by installing a horn plate in front of the detector array. A second run of the beam mapping process showed considerably improved performance in terms of reduced contributions from stray reflections as well as improved beam widths and signal to noise ratios however some cross leakage from the other horns was evident. A third run was undertaken with horns of wider entrance apertures to further improve the optical coupling. Although the overall signal

to noise ratios were improved, the horn leakage was not effected significantly, stray light reflections from the optical baffles were visible, and the beams appeared distorted in some cases.

Example beams are presented in figure 3.10. The three parts of the figure correspond to the three cases mentioned: direct illumination; horn coupling; and horn coupling with widened apertures. In each part, maps of the leftmost, central and rightmost pixels are presented as well as the stacked average over all the maps. The maps on the left half of the figure are of signal to noise ratio normalised to the baseline level of each pixel, while the maps on the right half are of signal to noise ratio normalised to the peak response of each pixel and displayed with a logarithmic colour scale to highlight low level details.

The distributions of signal to noise ratios and beam widths are presented in the histograms in figure 3.11. The same results are presented in figure 3.12 but as a function of the detector position in the focal plane array. A similar trend to that found in the preliminary evaluation is present in that the beams are generally narrowest and strongest in the center of the focal plane with a particularly weak and broad region on one edge. The first horn coupling test seems anomalous in the this respect as the signal to noise ratios are somewhat lower in the centre of focal plane. However it should be noted that the signal to noise ratios here are dependant on the individual detector performances and are not direct indicators of the optical power present. Poor noise performance or reduced responsivity can skew the results considerably.

3.4.4   Performance at 350 GHz

Tests of the 350 GHz array were performed in the horn coupling mode only. Stray light reflections from the baffles and array packaging were notably absent, however, the beams showed a number of interesting distortions. While they were well described by Gaussians at their peaks, extensive non-Gaussian tails were present. Additionally, new ghosts were present in most of the maps at positions slightly offset in the horizontal from the main beam.

These new ghosts were found to be due a misalignment of the horn plate with respect to the detector array. A second run with the correct alignment eliminated the ghosts and improved the singal to noise level, although the low-level non-Gaussianity remained. These contributions (clearly visible in the logarithmic maps in figure 3.13b) are most likely the result of further light leakage from the neighbouring feed horns.

Finally, a coma-like aberration is present a number of channels localised to one edge of the array. This is a manifestation of the edge distortions noted in the 150 GHz and preliminary evaluations. The aberrations affect the outputs of the Gaussian fitting function and lead to the reporting of beams with high ellipticities. As a measure of the resolving power of the system, the histograms in the lower right part of figure 3.14 are then somewhat biased as indicated by excesses in the reported major axis widths.

The spatial distribution of the aberrations can be seen in figure 3.13c where some
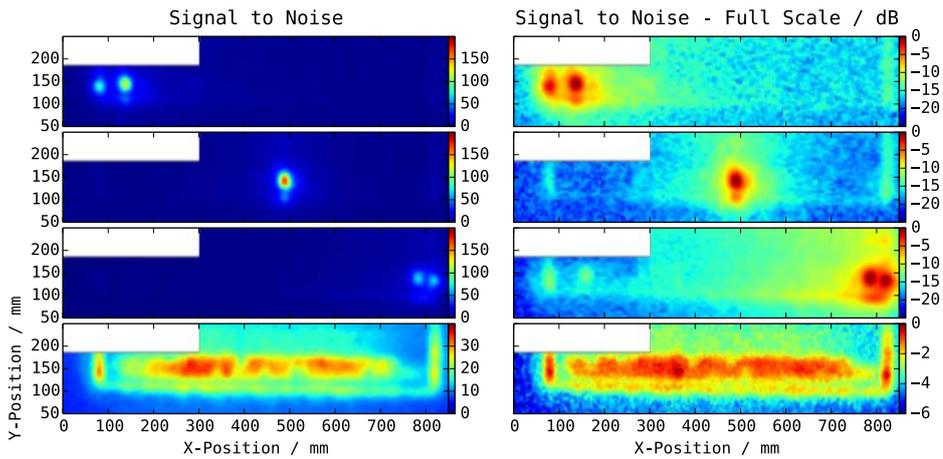
(a) *Direct illumination: Broader than expected and low SNR beams with considerable ghosting and off axis response.*
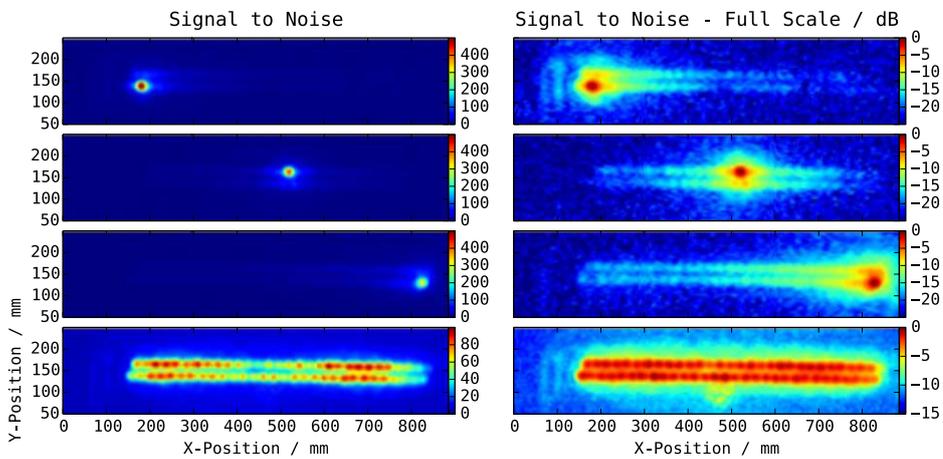
(b) *Horn coupling: reduced stray light, improved SNR and beam widths. Light leaks from neighbouring horns visible.*
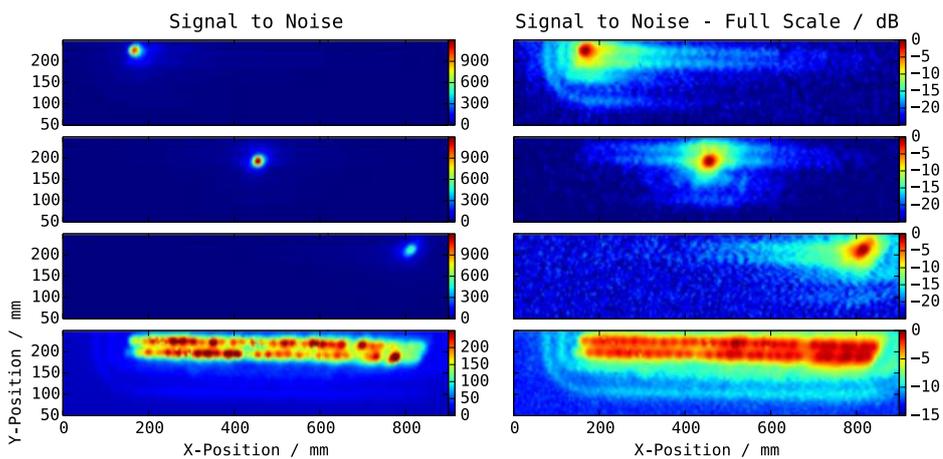
(c) *Horn coupling with widened horns: very good SNRs, some distortions in beam shapes and reduced horn leakage.*

Figure 3.10: *150 GHz maps: A selection of beam maps over three runs of the 60 pixel array. Horn coupling the detectors to the refractive optics system vastly improves the detector response and reduces stray light loading.*
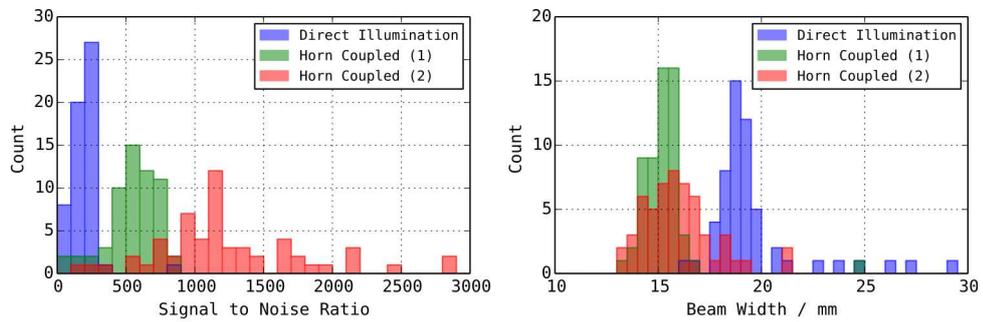
Figure 3.11: *Histograms of signal to noise ratio (left) and beam width (right) for the N = 58 working channels in the 150 GHz array.*
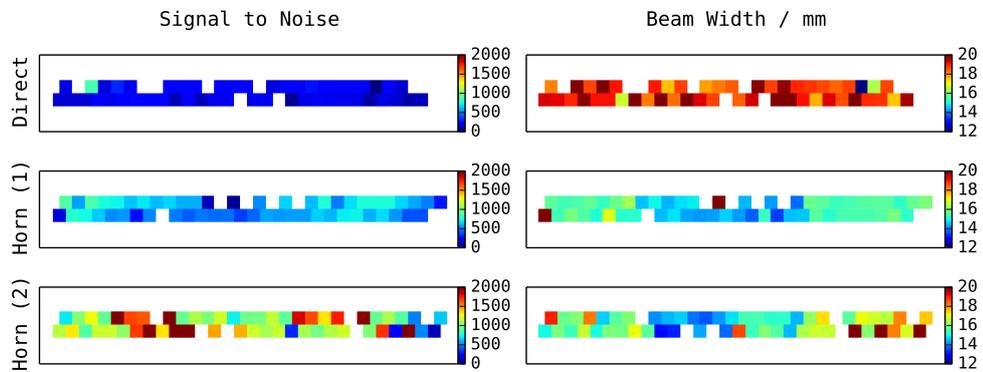
Figure 3.12: *150 GHz optics performance as a function of position across the focal plane.*
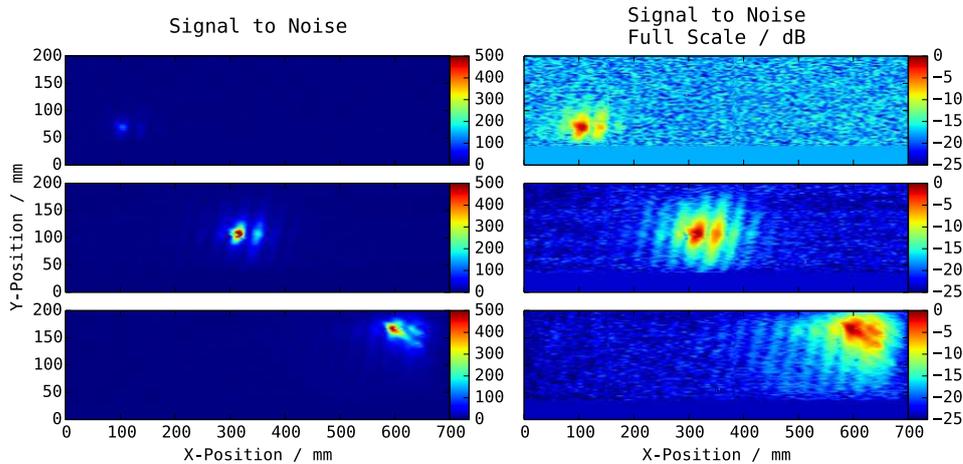
(a) *Misaligned horn plate: the main beams are accompanied by numerous minor beams originating from nearby feed horns.*
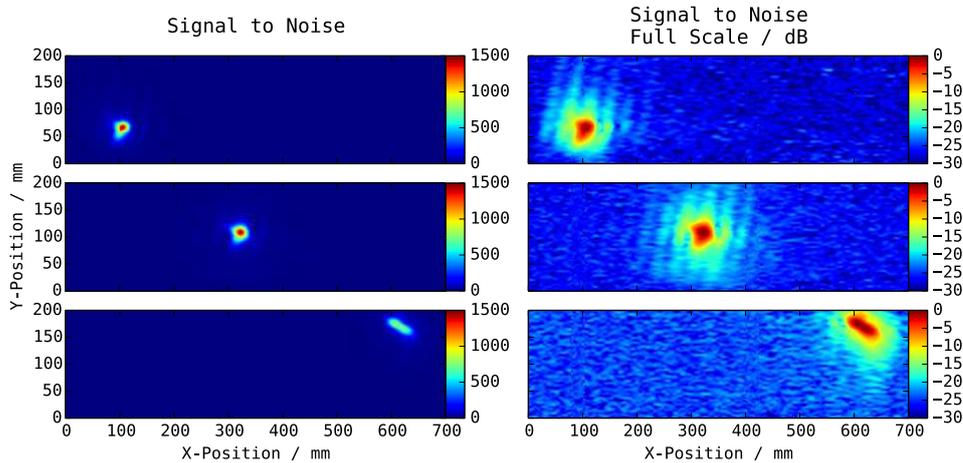
(b) *Properly aligned horn plate: the strongest minor side beams are gone but horn leakage is still present. The rightmost beam shows some considerable aberration.*
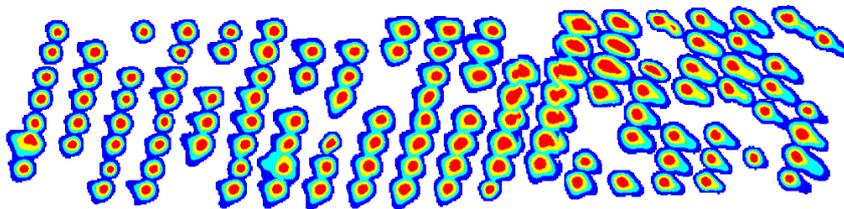
(c) *Properly aligned horn plate: the distribution of the optical aberrations is shown by overlaying the peaks of each beam.*

Figure 3.13: *350 GHz maps: A selection of beam maps over two runs of the 152 pixel array. Correct alignment if the horn plate to the detector array vastly improves the detector response and reduces contributions from horn leakage. Strong optical distortions are visible on one edge of the array.*
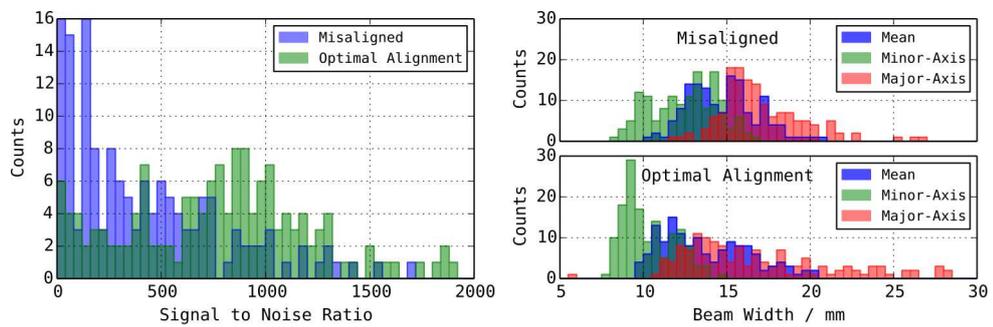
Figure 3.14:   *Histograms of signal to noise (left) and beam width (right) for the $N = 152$ detectors in the 350 GHz array. There are some ellipticities to the beams so the major, minor and mean beam widths are differentiated.*
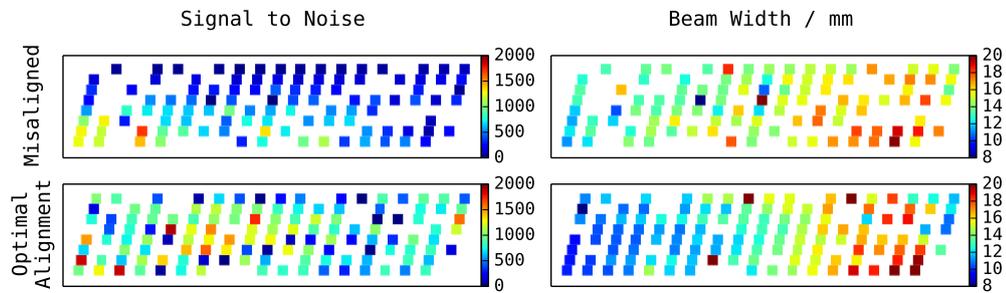
Figure 3.15:   *350 GHz optics performance as a function of position across the focal plane.*

preliminary results from a beam-mapping run are overlaid. The distorted beams do not exactly match the locations of the broadest width beams indicated in figure 3.15 because they were taken in a different run prior to minor optical adjustments, however the trend is the same: broader and more distorted beams on the right hand edge of the focal plane.

It should again be noted that the signal to noise ratios presented here are a product of the optics response of the system as well as the detectors' noise performance and responsivity. Further analysis is required to decouple the individual detector responses although it is believed that only in a few cases do the detector noise performances have a noticeable effect. One example of this is that there appear to be some regularly spaced pairs of detectors with reduced signal to noise in regions were the average signal to noise is high. This seems to be correlated with the locations of detectors with similar resonant frequencies.

3.5   Scanning mechanism

A graphical overview of the field scanning mechanism was given in figure 2.1 on page 23. The system involves rotating the main beam-folding mirror in an oscillatory manner to repeatedly sample the field of view. This is achieved by linking the back of the mirror via a short arm to a small circular cam that is driven by a feedback controlled servomotor.

The angular orientation of the mirror is reported back to the readout electronics system with an absolute rotary encoder and the timelines are used to generate the real time images. Figure 3.16 shows some example timelines for various frame rates up to 2 FPS. A close look at the high rate timelines reveals small bumps as the angle approaches the lower extreme of each cycle. This is caused by weakness in the bearings that connect the mirror to the arm and the arm to the cam. This small bump can provide quite a kick to the system and repeated kicks can cause the camera frame to oscillate and eventually lead to the cold stage of the cryostat warming up. As such, the frame rate is kept below 2 FPS.

The circular cam defines a sinusoidal scanning pattern and is not ideal, however no suitable alternative – such as a linear scan based on a heart-shaped cam – was implemented. In the sinusoidal case, less time is spent imaging the centre of the field of view and this is indicated in comparison to the ideal linear scan by the histograms in figure 3.17. At faster rates, unsampled spatial bins are interpolated over in the image generation software.

3.6   Optical efficiency

Given the parameters provided by the Zemax model of the optics system, along with other measured factors, it is possible to estimate the optical power present at the focal plane and to infer an optical efficiency. The diameter of the primary aperture is 0.44 m, the optimal focal length is 3.35 m, and a typical beam FWHM at focus is 11 mm. Zemax
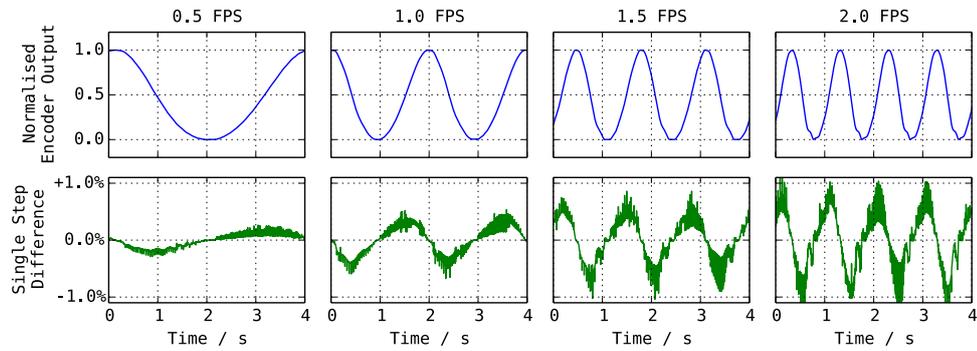
Figure 3.16: *Timelines from mirror encoder (blue) normalised between zero and one for increasing frame rates – note the bumping at higher frequencies that leads to shaking of the camera frame and warming of the cold stage. The single step difference (green) is a measure of the mirror velocity – note the clear departure from a sinusoidal form with increasing frame rate*
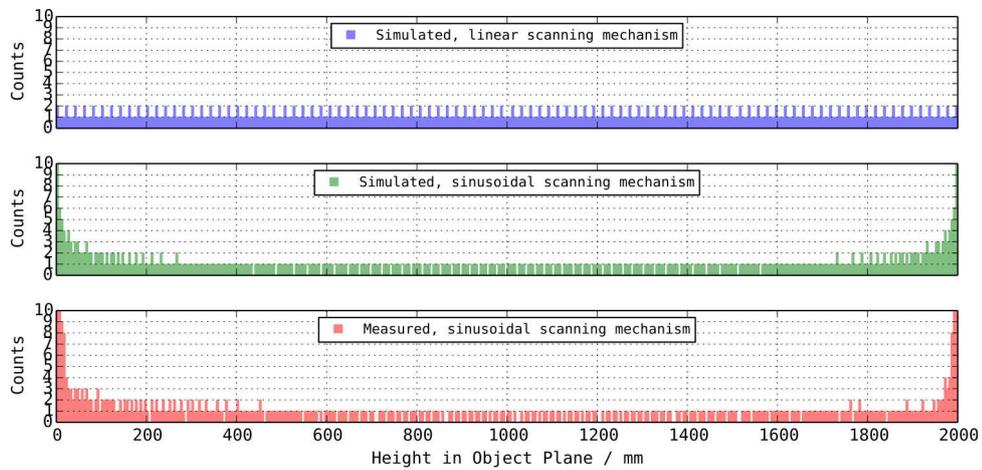
Figure 3.17: *Distribution of spatial field sampling for a simulated ideal linear scan rate (top), a simulated sinusoidal scan rate (middle), and the measured sinusoidal scan rate (bottom). In each case the frame rate is set to 1 FPS, the data rate is 500 Hz, the time duration is 1 second, and the spatial bin width is 5mm.*

Table 3.6: . *The measured transmission and emission properties of the optics system. For each of the key components, estimates are made of the power radiated by the component P_s and the power presented to a detector at the focal plane P_{FP}. The total power is used to calculate the photon noise limited sensitivity given in Table 7.1.*

Component	Thick. [mm]	Temp. [K]	Peak Trans. [%]	Peak Emiss. [%]	P_S [pW]	P_{FP} [pW]
Source	-	308	-	0.99	312	68
Mirror	-	293	-	0.01	0.9	0.2
L1	100	293	0.77	0.27	80	23
L2	65	293	0.84	0.17	50	16
Window	15	293	0.96	0.04	11	4.8
L3	70	150	0.83	0.19	28	12
630 GHz LPE	-	150	0.93	-	-	-
540 GHz LPE	-	150	0.94	-	-	-
450 GHz LPE	-	10	0.88	-	-	-
350 GHz BP	-	10	0.80	-	-	-
350 GHz BP	-	0.25	0.69	-	-	-
				Total signal power at horn entrance:		68
				Total stray light power at horn entrance:		56
				Total power at horn entrance:		124
Detector Efficiency	25%			*Total detected signal power:*		17
				Total detected stray light power:		14
				Total detected power:		**31**

informs us that the camera throughput is $A\Omega = 1.33 \times 10^{-6}\,\text{m}^2\,\text{sr}$, the measured filter profiles (see figure 2.3) indicate a $\Delta\nu = 6\%$ wide band centred at $\nu = 348\,\text{GHz}$ or $\lambda = 862\,\mu\text{m}$, so the number of spatial modes in the beam is 1.84, which is consistent with the two electromagnetic modes that are permissible through the feed-horn and waveguide to the detector (see figure 2.2. The transmission and emission properties of the optical components were measured and are summarised in Table 3.6 along with estimates of the in-band power from a human target present through the optics system onto each horn in the focal plane. The total incident power per horn is 124 pW, of which 68 pW is from the 308 K source and 56 pW is stray light from the optics. Given that the power emitted from the 308 K source into the main beam is 312 pW (estimated from Equation 2.2), the overall optical efficiency of the camera is $\eta_{opt} = 22\%$.

3.7 Summary

The design and performance of the optics system have been presented, including not only the final working performance but also some of the issues that were faced in getting

the system up and running. The optics are by no means perfect and the detectors are not uniformly illuminated across the field of view, however, the cost and complexity of the system are minimal and the performance is sufficient for a working camera.

The large window and the fast, refractive optics contribute to a considerable excess in optical loading on the focal plane array. Further discussions on this and other stray light effects are summarised in the noise budget analysis in Chapter 7.

Chapter 4

The Detector System

The ultimate goal for a general purpose imaging system operating in the terahertz would be to achieve high resolution, video rate imaging of large fields of view with background limited sensitivities. Kinetic inductance detectors have recently been demonstrated with background limited sensitivities under very low optical loading conditions [Yates et al., 2011, Mauskopf et al., 2014], they have fast response times, and relative to other comparable technologies – they are easy to manufacture into large arrays which are also relatively easy to read out. This makes them an ideal choice for terahertz imaging applications, and the instrument we have developed has proven that the technology is can be a viable and cost-effective option.

Initial runs with the camera were with a 60 pixel LEKID array operating in a narrow optical band at 150 GHz or 2 mm. The camera was then upgraded to a 150 pixel LEKID array at 350 GHz or 850 μm in order to improve the image resolution as set by the diffraction limited optics. The two arrays have worked well enough for successful demonstrations of the camera and at no point have they been a limiting factor in the performance of the system. As such, not a great deal of detector development has been involved in this project to date. Most of the development time has been spent on optical, electronic, cryogenic, and software engineering.

In this short chapter, I present the design and characterisation of the two LEKID arrays. Dr Doyle is acknowledged for the design and modelling. Dr Dunscombe is acknowledged for the fabrication. A side-by-side comparison of the differences between the arrays in the two bands is given in Table 4.1 on page 73, and some estimations of detector noise and responsivity are given towards the end.

4.1   150 GHz design

The 150 GHz pixels are based on the classical LEKID design [Doyle et al., 2008]:

- The aluminium is 40 nm thick and sits on a 525 μm high resistivity silicon substrate.

- Absorption is direct with the aluminium facing out of the focal plane. See Figure 4.1.
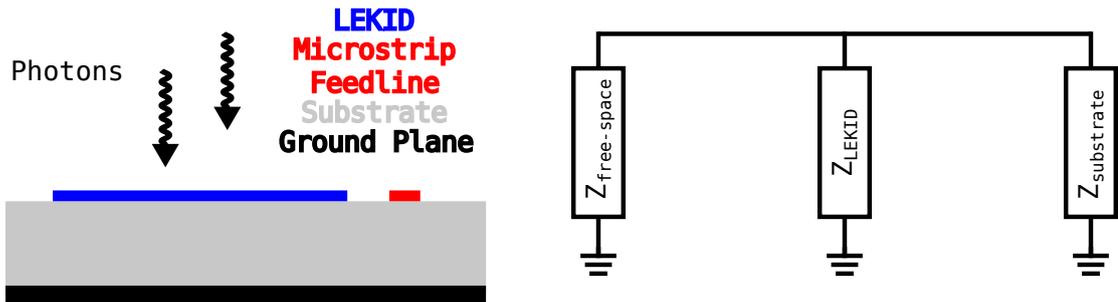
Figure 4.1: *The 150 GHz LEKID optical coupling design: (left) the aluminium LEKIDs sit on the top of the wafer and are directly illuminated by incoming photons, and (right) the equivalent transmission line model of electromagnetic wave propagation. Figure adapted from [Doyle, 2008] with permission.*
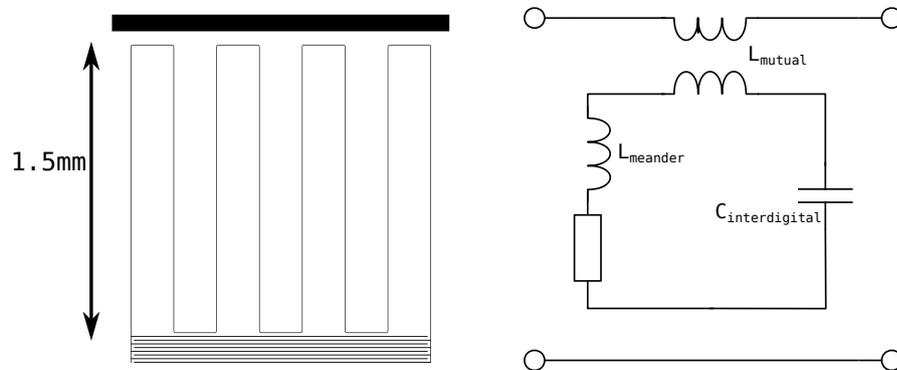
Figure 4.2: *The 150 GHz LEKID devices are based on the classic inductively coupled LEKID design: (left) the photomask, and (right) the equivalent microwave readout circuit.*
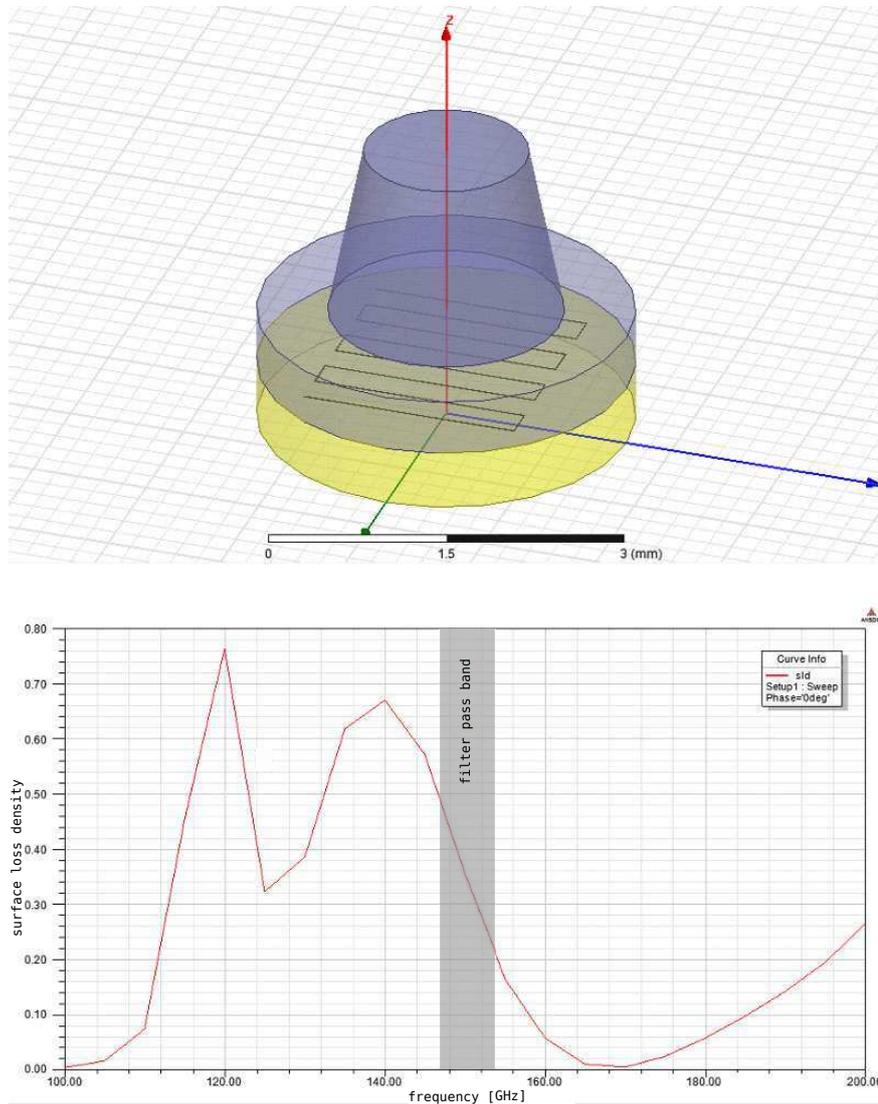
Figure 4.3:   *Optical coupling efficiency of the 150 GHz LEKID meander including a section of the feedhorn is modelled in HFSS: (top) the model, and (bottom) the results: ~30% efficiency at 150 GHz.*
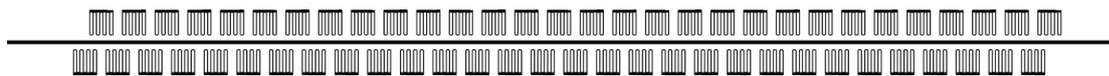
Figure 4.4:   *The 150 GHz array layout: Two rows of 30 LEKIDs are spaced evenly along the microstrip transmission line. The rows are offset by half of the single-row pixel separation. This separation is sufficient to Nyquist sample the horizontal axes of the image plane given the diffraction limited spot-size expected from the optics system at 150 GHz.*

- The 1875×1875 µm absorber element is a classical, straight, inductive meander that is sensitive to a single polarisation. The meander has 8 turns with a line width of 4 µm and a separation of 275 µm. See Figure 4.2.

- The HFSS model indicates around 30% coupling efficiency in a 10% wide band centred at 150 GHz. See Figure 4.3.

- The interdigital capacitor section has 7 fingers, each 6 µm wide and separated by 18 µm.

- The linear array is composed of 60 LEKIDs separated by $F\lambda/2$ in the horizontal axis so that the focal plane is Nyquist sampled. See Figure 4.4.

- The resonators are inductively coupled to a microstrip transmission line of width 240 µm at a distance of 500 µm. The stainless steel array packaging acts as the ground plane.

- The variable length of the the outer capacitor finger sets the resonant frequencies to be around 1.5 GHz with 2 MHz spacing between resonances, providing an overall readout bandwidth of 120 MHz.

- The total inductor volume is 2700 µm^3 and the spatial filling factor is 1.9%.

4.2   150 GHz characterisation

The array was installed into the camera and the microwave transmission was measured with a vector network analyser. The results are displayed in Figure 4.5. A total of 58 resonances were identified from the amplitude of S_{21} sweep. The data set was reduced to keep only the points close to each resonance and the individual resonant frequencies and quality factors were measured. The resonant frequencies range from 1.425 GHz to 1.545 GHz giving a total bandwidth of 120 MHz, as expected from the design. The separation between resonances varies considerably from as much as 8 MHz down to a few hundred kHz in some cases. If any resonances are too close together, the response of one can affect the response of the other and neither can be used for imaging. In total, four pairs of closely coupled resonators had to be ignored due this effect, bringing the total array yield to 50 out of 60, or 83.3%. The resonator quality factors are limited by coupling losses to around 20000. The quality factor distribution is uniform across the array.

A single pixel was selected for a more detailed analysis of optical response and noise. LEKID_03 ($f_{res} = 1.4313$ GHz) was chosen due to its central spatial position along the array, and also because it showed microwave properties typical of the majority of the detectors in the array. A variable temperature blackbody source was set up in the field of view of the camera at the point of focus of LEKID_03. Frequency sweeps were taken with a standard homodyne readout system as the blackbody was varied in temperature
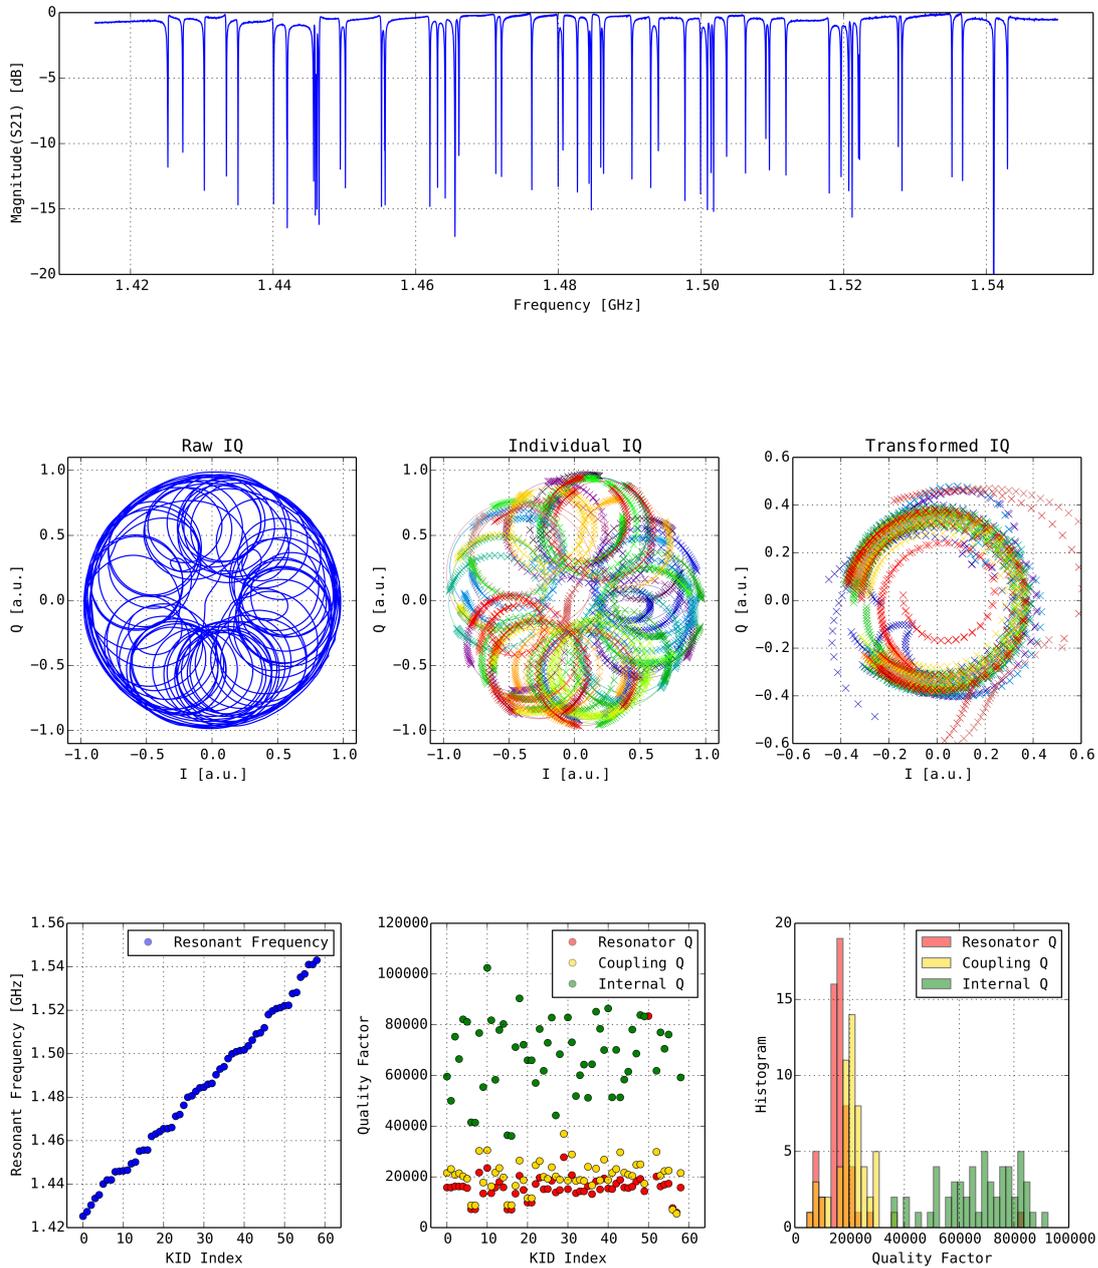
Figure 4.5: *Microwave characterisation of the 150 GHz LEKID array with a VNA sweep driven at −65 dBm: (top) the magnitude of $S_{21}(f)$, (middle) the IQ plane representation, and (bottom) extracted resonant frequencies and quality factors.*
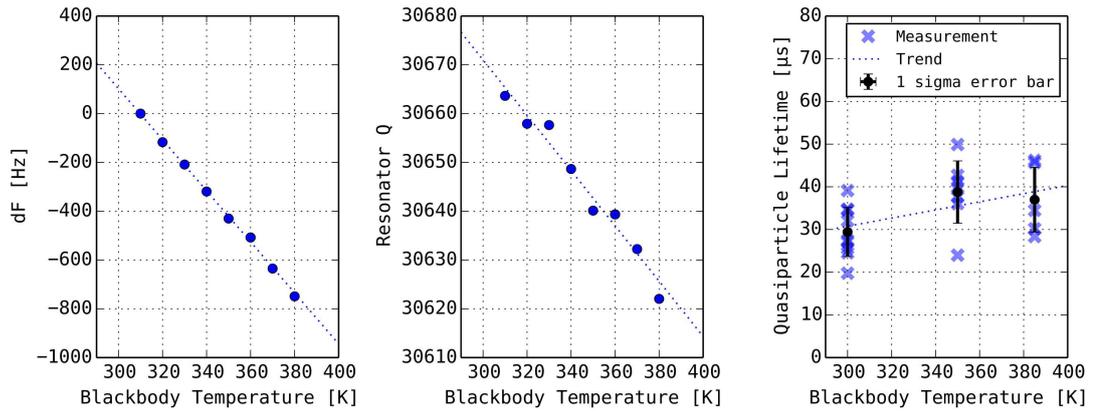
Figure 4.6: *The response to variations in optical power for LEKID_03 at 1.4313 GHz. (left) the frequency response, (centre) the dissipative response, and (right) the variation in quasiparticle lifetime.*
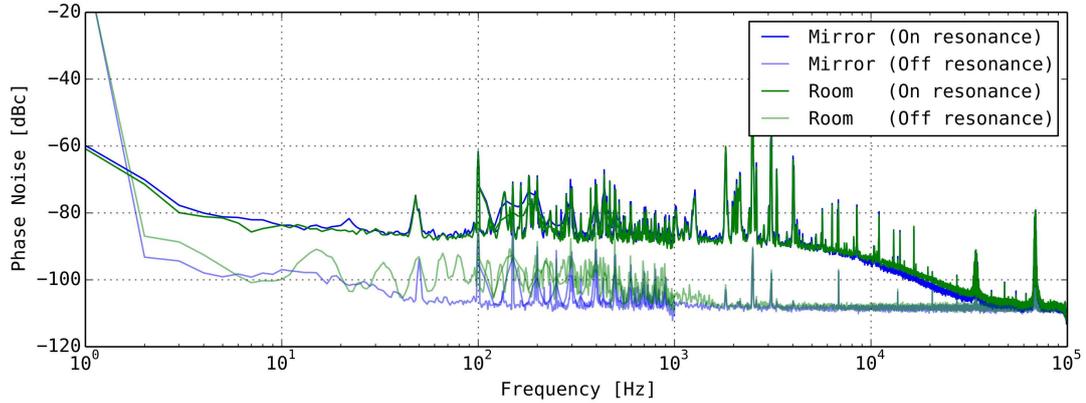
Figure 4.7: *Noise power spectra for LEKID_03 measured with a fixed tone on resonance and a fixed tone off resonance: (green lines) the LEKID is looking out of the camera into the room, and (blue lines) a mirror is placed under the window so that the LEKID sees the cold interior of the cryostat.*

between 310 and 380 K. The resonant frequency and quality factors were calculated as a function of temperature and are displayed in the first two graphs in Figure 4.6. The responses were linear with temperature as expected for blackbodies emitting in the Rayleigh-Jeans portion of the spectrum. The same LEKID was measured with a fixed readout tone from the same homodyne readout system and the quasiparticle lifetime was measured at three different temperatures. No significant variation was noted with temperature and the lifetime was typically measured at 30 to 40 μs. The results are included the third graph in Figure 4.6.

The resonant frequency response of the KID was sampled repeatedly with a fixed readout tone centred at the resonant frequency for one minute at a slow rate (2 kSPS) and for one second at a high rate (200 kSPS). Measurements were also taken with the readout tone positioned away from the resonance. This was repeated with the cryostat

window open and with the cryostat window blocked by a mirror – in the first case, the KID sees the warm (300 K) room and in the second case, the KID sees the cold interior of the cryostat (100-150 K). The resulting noise power spectra are displayed in Figure 4.7. The on-resonance spectra show the characteristic mid-frequency white noise and high frequency roll-off. This is not present in the off-resonance spectra. The white noise at −110 dBc comes from the cryogenic amplifier and the white noise at −85 dBc is due to fluctuations in quasiparticle density. The high- and low-load results are almost identical except that the roll-off occurs at a very slightly higher frequency in the high optical loading case. There is an argument here that the increased optical power dominates the quasiparticle dynamics and that, therefore, the detector sensitivity is fundamentally limited by the photon noise.

4.3   350 GHz design

The 350 GHz pixels are based on the dual polarisation design from [Roesch et al., 2012]:

- The aluminium is 40 nm thick and sits on a 500,μm high resistivity silicon substrate.

- Absorption is through the back of the substrate with the aluminium facing down towards the base of the array packaging which is covered in a layer of silicon-carbide-blackened epoxy to absorb stray optical radiation. See Figure 4.8.

- The 1450×1450 μm absorber element is a 4th order Hilbert design that is sensitive to both polarisations. The line width of the Hilbert meander is 3 μm and the unit section length is 95 μm. See Figure 4.9.

- The HFSS model indicates around 35% coupling efficiency in a 10% wide band centred at 350 GHz. See Figure 4.10.

- The interdigital capacitor has 6 fingers, each 6 μm wide and separated by 9 μm.

- The linear array is composed of 152 LEKIDs separated by $F\lambda/2$ in the horizontal axis so that the focal plane is Nyquist sampled. See Figure 4.11.

- The resonators are capacitively coupled to a co-planar waveguide transmission line with a central conductor width of 485 μm, free space gaps of 24 μm and a return conductor width of 21 μm in the vicinity of the LEKIDs. The coupling capacitor is 21 μm wide, 1450 μm long, and separated by 6 μm from the return conductor (ground plane).

- The resonant frequencies are deigned to be around 1.5 GHz with 3 MHz spacing between resonances, providing an overall readout bandwidth of 450 MHz.

- The total inductor volume is 2900 μm^3 and the spatial filling factor is 3.45%.
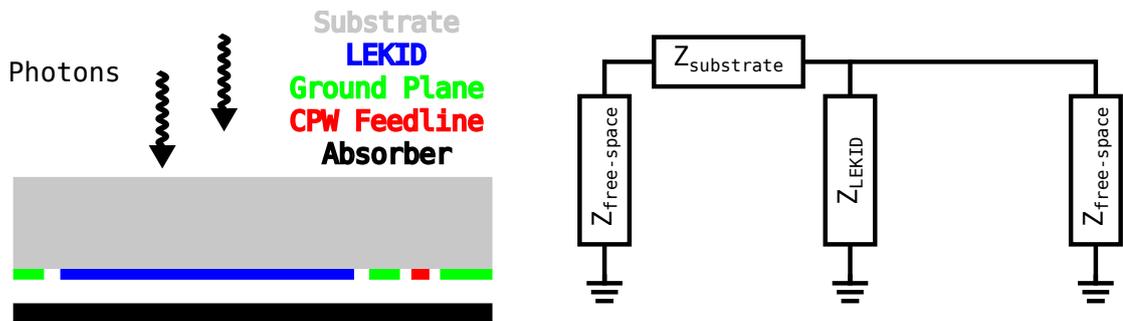
Figure 4.8: *The 350 GHz LEKID optical coupling design: (Left) the aluminium LEKIDs sit between the substrate and the absorber and are back-illuminated through the substrate: (right) the equivalent transmission line model of electromagnetic wave propagation. Figure adapted from [Doyle, 2008] with permission.*
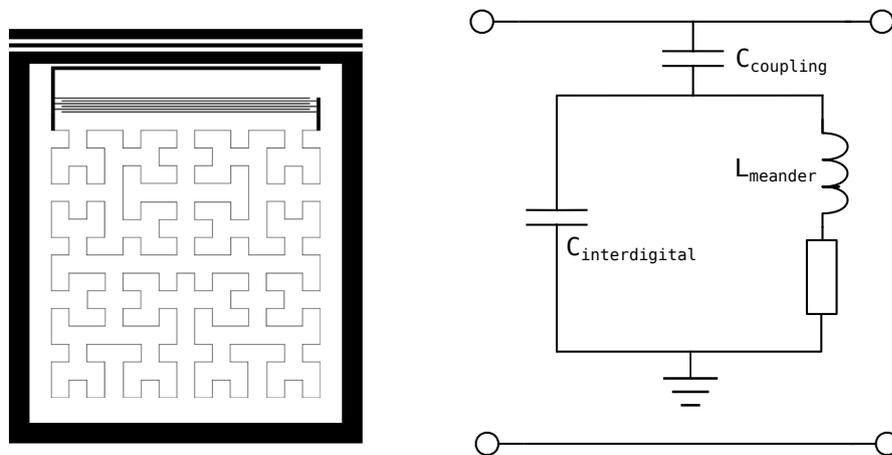
Figure 4.9: *The 350 GHz LEKID devices are capacitively coupled to a CPW transmission line and the absorber is based on a 4th order Hilbert curve: (left) the photomask, and (right) the equivalent microwave readout circuit.*
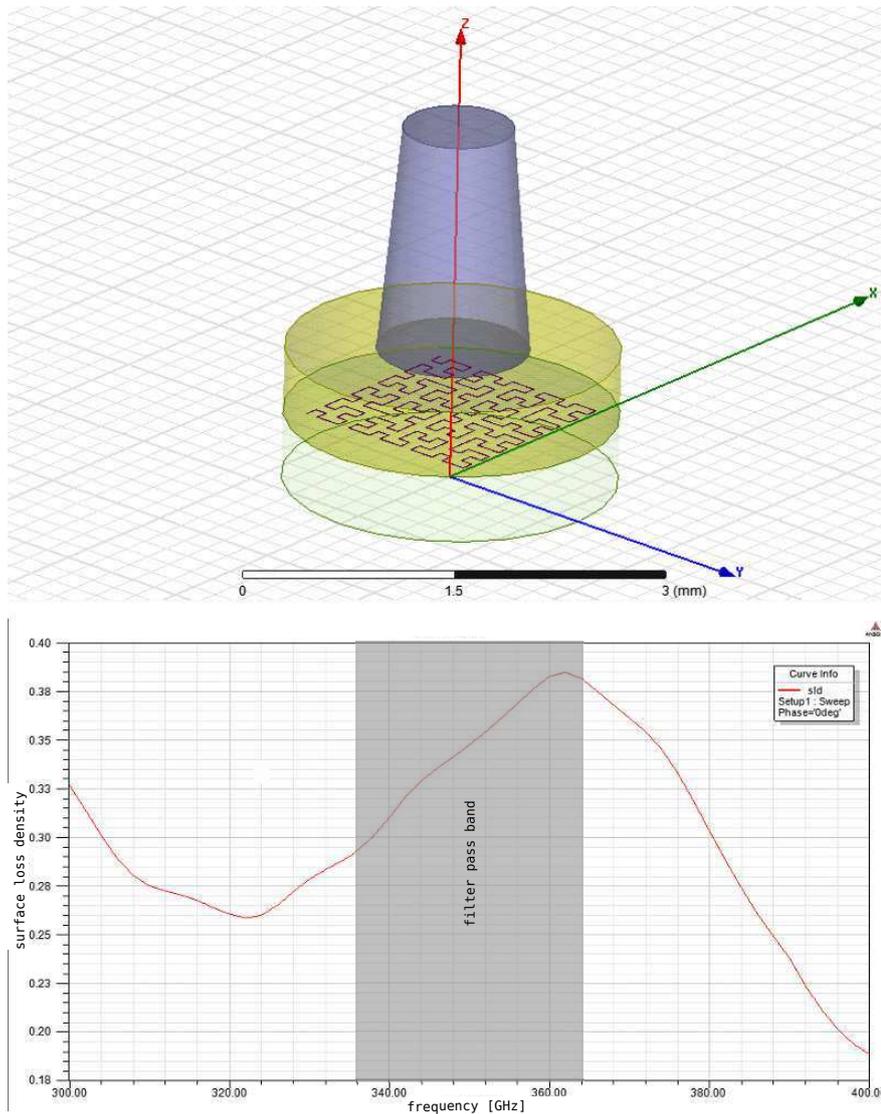
Figure 4.10: *Optical coupling efficiency of the 350 GHz LEKID meander including a section of the feedhorn is modelled in HFSS: (top) the model, and (bottom) the results: ~35% efficiency at 350 GHz.*
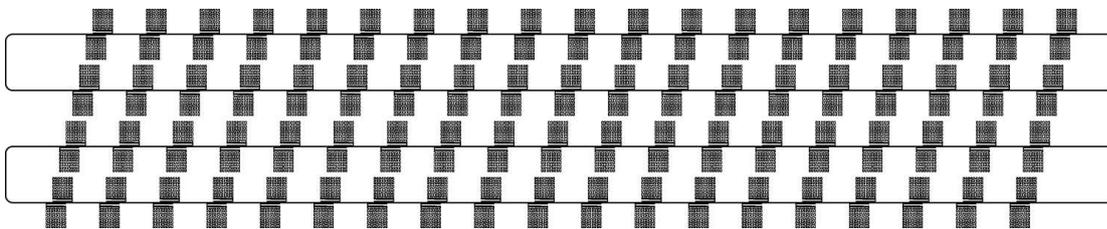
Figure 4.11: *The 350 GHz array layout (ground plane not shown): 8 rows of 19 LEKIDs are spaced evenly along the CPW transmission line. The rows are offset by roughly one eighth of the single-row pixel separation. This arrangement is also sufficient to Nyquist sample the horizontal axis of the image plane given the diffraction limited spot-size expected from the optics system at 350 GHz.*

4.4   350 GHz characterisation

A similar characterisation to that performed on the 150 GHz array was performed on the 350 GHz array.

From the microwave analysis in Figure 4.12, the resonant frequencies span from 1.25 to 1.73 GHz – a slightly higher range than the expected 450 MHz but still within the 500 MHz bandwidth of the NIKEL readout electronics system. Some broad band features are present in the S_{21} plot that interfere with a number of resonators. The total number of KIDs identified with an automatic KID finding algorithm was 138, of which 8 are too closely coupled to be of use. The brings the useful detector yield in this array to 130 out of 150 or 86.6%.

Again, the resonator quality factors are limited by coupling losses to around 20000-30000. The distribution varies across the the array with lower frequency KIDs tending to have higher Q-factors.

Measurements of the overall sensitivity using the multiplexed readout system were taken and these results are included in the final chapter of this thesis. The quasiparticle lifetime was measured for one 30 μs detector.

An additional dark characterisation was performed with a dedicated low-background LEKID test system based in our group here in Cardiff. Twelve random selected KIDs were measured as a function of bath tmerperature in, effectively, a completely dark environment. The results are presented in Figure 4.6. The first graph shows the classic Mattis-Bardeen response of fractional frequency shift (x) to bath temperature. The second graph shows the expected drop in quality factors with bath temperature. The third graph compares the two results in order to fit the proportionality constant(β).

4.5   Detector noise

4.5.1   Generation-recombination noise

Thermal lattice vibrations within a superconductor break Cooper pairs and generate quasiparticles, as will any external power sources that couple in such as optical photons or electronic readout signals. Quasiparticle excitations will relax and recombine back into Cooper pairs in a mean time $\tau_{qp} \propto 1/n_{qp}$ where n_{qp} is the quasiparticle density. These intrinsic thermal fluctuations in n_{qp} are indistinguishable from any externally sourced fluctuations and will always be present in the superconductor, although at exponentially lower levels as the temperature is reduced. This sets a fundamental limit to the sensitivity of a kinetic inductance detector.

For a LEKID at a temperature $T \ll T_c$, with an inductive meander of volume V_L, the number of quasiparticles present N_{qp} due to thermal (phonon) excitations is given by

$$N_{qp}(T) = n_{qp}(T)V_L = 2N_0 V_L \sqrt{2\pi k_B T \Delta} \ \exp\left(\frac{-\Delta}{k_B T}\right) , \qquad (4.1)$$
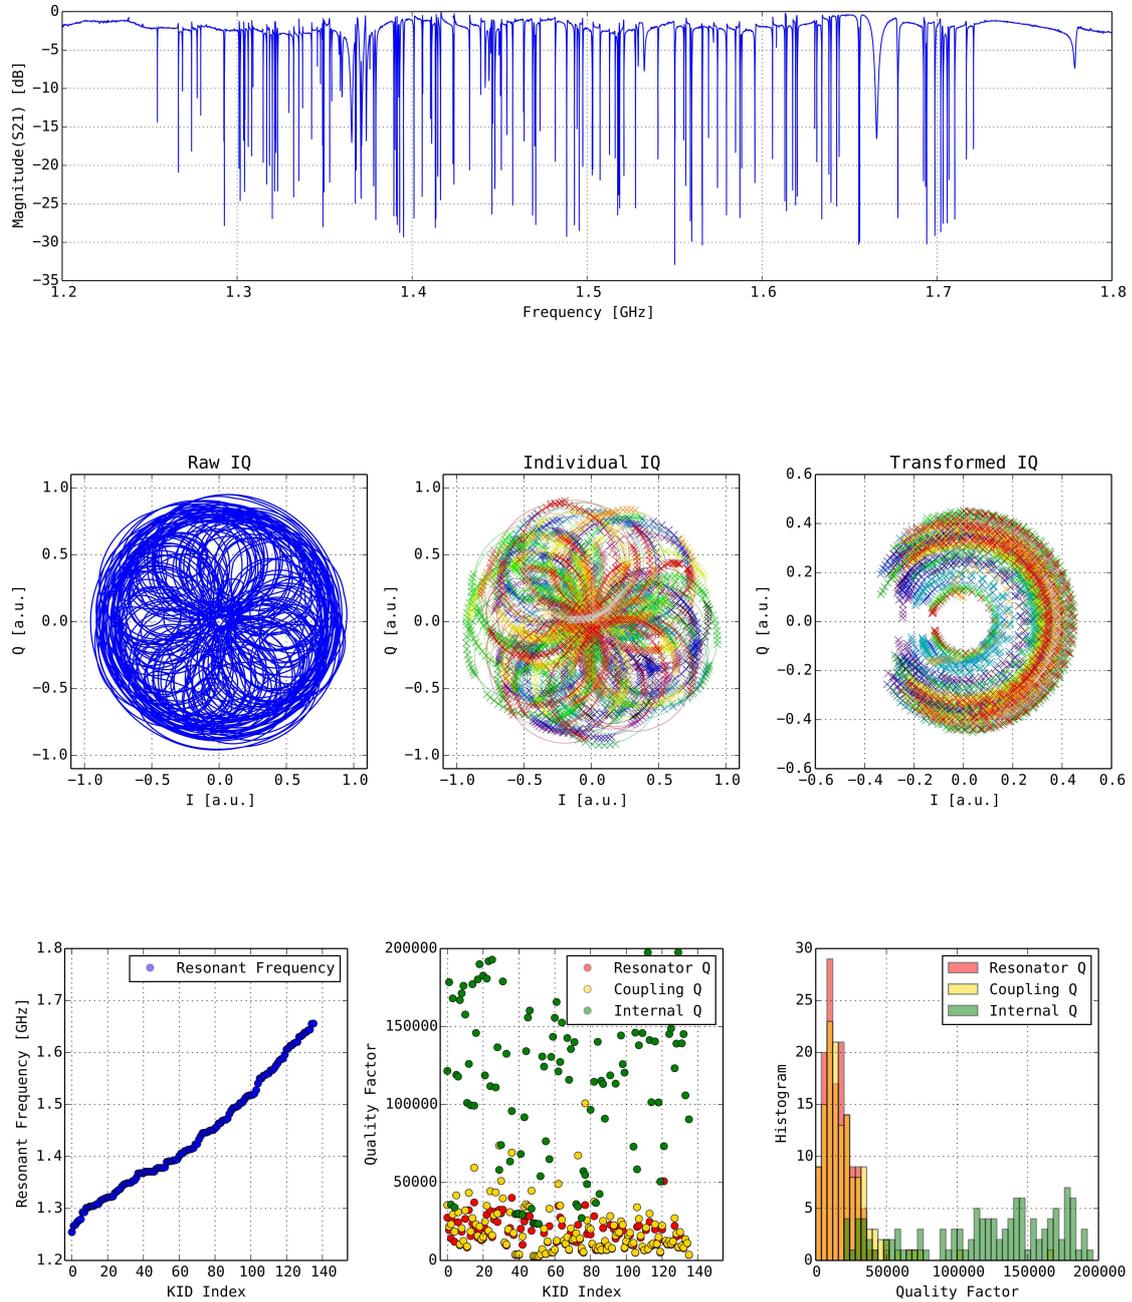
Figure 4.12: *Microwave characterisation of the 350 GHz LEKID array with a VNA sweep driven at −70 dBm: (top) the magnitude of $S_{21}(f)$, (middle) the IQ plane representation, and (bottom) extracted resonant frequencies and quality factors.*

Table 4.1: *A comparison of the detector array characteristics between the two optical bands.*

Frequency	150 GHz	350 GHz
Wavelength	2 mm	850 µm
Filter bandwidth	10%	10%
Number of pixels	60	152
Array layout	2×30	8×19
Optical coupling	Direct	Through substrate
Optical efficiency	30%	35%
Polarisation sensitivity	Single pol.	Dual pol.
Microwave coupling	Inductive	Capacitive
Microwave transmission line	Microstrip	CPW
Microwave frequency separation	2 MHz	3 MHz
Microwave centre frequency	1.5 GHz	1.5 GHz
Microwave bandwidth	120 MHz	450 MHz
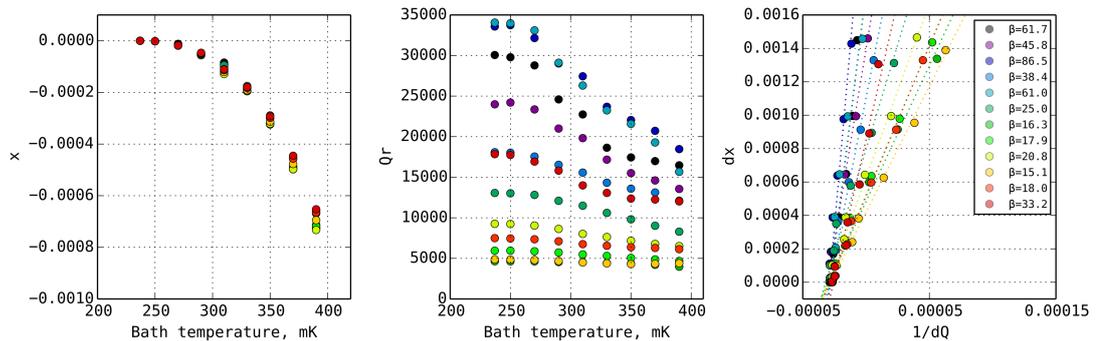
Figure 4.13: *The response to variations in bath temperature power for a random sample of 12 LEKIDs. (left) the fractional frequency response, (centre) the dissipative response, and (right) fitting β.*

and the quasiparticle lifetime τ_{qp} is

$$\tau_{qp}(T) = \frac{\tau_0}{\sqrt{\pi}} \left(\frac{k_B T_c}{2\Delta}\right)^{5/2} \sqrt{\frac{T_c}{T}} \, \exp\left(\frac{\Delta}{k_B T}\right) \quad [\text{s}] \,, \tag{4.2}$$

where the material dependent parameters N_0, τ_0, Δ, and T_c are respectively the normal state single spin density of states at the Fermi surface, the electron-phonon interaction time, the superconducting gap energy, and the superconducting critical temperature. Typical values for the thin film aluminium used in our devices is given in Table 4.2. For a given device, the proportionality between N_{qp} and τ_{qp} is described by the temperature independent constant,

$$N_{qp}\tau_{qp} = N_0 \tau_0 V_L \frac{k_B^3 T_c^3}{2\Delta^2} \quad [\text{s}] \,, \tag{4.3}$$

Table 4.2: *Material dependent parameters for estimating GR noise of the 350 GHz thin aluminium LEKIDs described in this thesis.*

Parameter	Symbol	Value	Unit
Normal state single spin density of states	N_0	1.7×10^{10}	$\text{µm}^{-3}\,\text{eV}^{-1}$
Electron-phonon interaction time	τ_0	80	ns
Superconducting energy gap	Δ	180	µeV
Superconducting critical temperature	T_c	1.4	K
LEKID inductor volume	V_L	3000	µm^3

Generation-recombination (G-R) noise is characterised by the power spectral density of fluctuations in the number of quasiparticles due to thermal interactions as a function of frequency, as given by

$$S_{N_{qp}}(f) = \frac{4 N_{qp}\tau_{qp}}{1 + (2\pi f \tau_{qp})^2} \quad \left[\text{Hz}^{-1}\right] \,. \tag{4.4}$$

The noise spectrum is typically flat but rolls off when the sampling interval ($\Delta\tau = 1/f_s$) is of order τ_{qp} or faster.

The responsivity of a KID is defined as the change in quasiparticle number for a given change in absorbed optical power and can be be shown to be equal to

$$R_{N_{qp}} = \frac{dN_{qp}}{dP_{abs}} = \frac{\eta_e \tau_{qp}}{\Delta} \quad \left[\text{W}^{-1}\right] \,, \tag{4.5}$$

where η_e is the quasiparticle production efficiency. For a standard BCS superconductor, $\eta_e \approx 0.6$ when absorbing photons with $h\nu \gtrsim 4\Delta$ [de Visser et al., 2015].

The noise equivalent power, NEP, referring to the equivalent absorbed optical signal

power required to equal the G-R noise power is then

$$NEP_{GR}(f_s) = \frac{2\Delta}{\eta_e}\sqrt{\frac{N_{qp}}{\tau_{qp}}}\frac{1}{\sqrt{1 + (2\pi f_s \tau_{qp})^2}} \quad \left[\text{W Hz}^{-\frac{1}{2}}\right] . \tag{4.6}$$

However, this model is only valid in the case of a dark detector. Any KID that is under an optical load will have an excess of quasiparticles and, therefore, a reduced quasiparticle lifetime and increased G-R noise. A corrected estimate for the NEP can be made by applying a model where the KID sits at an elevated effective temperature T_{eff} that is dependent on the incident optical power. The power transferred from Cooper pair states to the quasiparticles gas by phonon interactions is given, to first order, by

$$P_{phonon} = \frac{N_{qp}\Delta}{\tau_{qp}} \quad [\text{W}] . \tag{4.7}$$

Thus, we can write the optically unloaded quasiparticle number as a function of the phonon power,

$$N_{qp,\text{ unloaded}} = \sqrt{\frac{N_0 \tau_0 V_L P_{phonon}}{2}}\left(\frac{k_B T_c}{\Delta}\right)^{3/2} . \tag{4.8}$$

The optically loaded quasiparticle number is then simply

$$N_{qp,\text{ loaded}} = \sqrt{\frac{N_0 \tau_0 V_L (P_{phonon} + P_{abs})}{2}}\left(\frac{k_B T_c}{\Delta}\right)^{3/2} . \tag{4.9}$$

Then, T_{eff} can be found by solving Equation 4.1 for T, the loaded quasiparticle lifetime can be computed from Equation 4.2 as $\tau_{qp,\text{ loaded}} = \tau_{qp}(T_{\text{eff}})$, and the corrected G-R NEP can be computed from Equation 4.6.

For the $350\,\text{GHz}$ detectors, the noise equivalent power due to G-R noise is estimated using the method described above at $NEP_{GR} = 6.5 \times 10^{-17}\,\text{W}/\sqrt{\text{Hz}}$.

4.5.2 TLS noise

Amorphous structure in the thin film lattices, especially at the boundaries between the film, the substrate, the oxide layer and free space, cause non-uniformities in the electric potential around a KID. So called Two-Level-System (TLS) interactions occur where electrons tunnel between the various minima of the potential. This leads to localized variations in the quasiparticle density and dielectric constant of the superconductor, and manifests as noise in the resonant frequency of the KID.

The effects of TLS noise can be mitigated by driving KID resonators with more powerful readout tones such that the noisy TLS states become 'saturated' and fluctuations are reduced. The detectors in use in this thesis already require high power levels and TLS noise has never been an issue.

The noise power spectrum of TLS fluctuations has a $1/f$ distribution and typically dominates on time-scales lower than $\sim 10\,\text{Hz}$. This has been a problem for astronomical

systems that are modulated at low frequencies but, again, this is less of an issue in this system as the detectors are modulated at a much higher rate ($\sim 500\,$Hz).

4.6   Summary

In summary, two detector arrays were designed and fabricated within the group here at Cardiff. The microwave performance of the arrays was as expected although some overlap between resonance lead to the reduced detector yields of 85% on each array. Responses to optical power and bath temperature have been presented and no unexpected effects were noted. The detector sensitivity is estimated to be of order 6×10^{-17} W/$\sqrt{\text{Hz}}$. This results is discussed further in the overall sensitivity analysis in Chapter 7. From the offset, the detector arrays performance was suitable for imaging purposes and there was never any need for any further investigations into any possible performance limitations. Development was focussed on the more pressing issues of electronic readout and software development.

Chapter 5

The Electronic Readout System

Kinetic inductance detectors exhibit variations in complex electrical impedance when light is absorbed. Electronic readout systems are responsible for measuring and recording the magnitude of this change. In this chapter I present an overview of the main KID readout techniques, including mechanisms for single pixel and multiplexed readout. I go on to describe the design and characterisation of the two multiplexing systems used in our terahertz camera.

5.1 KID readout techniques

KID readout is achieved by monitoring the complex forward transmission, S_{21}, of a probe signal that couples to KID resonators. If a_1 and b_1 are, respectively, the transmitted and reflected components of the probe signal at the input to a detector array, and a_2 and b_2 are the transmitted and reflected components at the output, then the S_{21} scattering parameter is defined by the following linear relationship:

$$\begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}. \tag{5.1}$$
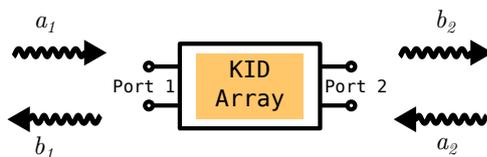
Figure 5.1: *Scattering parameters describe the relationships between transmitted and reflected waves at the ports of an n-port network. The S-parameters for a two port network such as a KID array are defined in Equation 5.1.*

For an inductively coupled KID resonator, S_{21} as a function of angular frequency ω

is given by [Doyle, 2008, F. S. Porter, 2002, Zmuidzinas et al., 2002, Barry, 2014]

$$S_{21}(\omega) = \frac{2}{2 + Z_0/Z_{in}(\omega)} = \frac{Q_r/Q_i + 2jQ_r x_r}{1 + 2Q_r x_r} \,, \tag{5.2}$$

from, where Z_0 is the characteristic impedance of the transmission line (typically $50\,\Omega$) and Z_{in} is the KID's complex impedance. Q_r and Q_i are the total and internal (i.e. loaded and unloaded) quality factors of the KID, and x_r is the coupling-modified fractional frequency shift, which is a function of ω and of the resonant frequency ω_r of the KID.

Absorbed optical power modifies Z_{in} which produces measurable changes in Q_i (dissipative response) and ω_r (frequency response). The frequency response is typically characterised by the standard fractional frequency shift, $x = \delta\omega_r/\omega_r$. The dissipative response is typically characterised by the change in the inverse of the internal quality factor, δQ_i^{-1}. The two responses are related by a frequency dependent factor β, where $\beta(\omega) = 2\,\delta Q_i^{-1}/x$.

Detector characterisation is achieved by measuring $S_{21}(\omega)$ across a range of frequencies (by shifting the frequency of the probe signal) and then fitting to the parameters that define the resonance. This is actually quite a time consuming process and becomes impractical when reading out KIDs at high rates. Instead, a fixed frequency probe is used and the frequency and dissipation responses are then estimated or inferred from the variations in S_{21} at that frequency alone. For a probe signal set at ω_r, variations in the S_{21} amplitude are a proxy for the dissipative response and variations in the S_{21} phase are a proxy for the frequency response.

S_{21} is usually measured in terms of I and Q, where $S_{21} = I + jQ$. In this IQ representation, the amplitude and phase as functions of time are given by

$$A(t) = \sqrt{I(t)^2 + Q(t)^2} \tag{5.3a}$$

$$\phi(t) = \tan^{-1}\left(\frac{Q(t)}{I(t)}\right) . \tag{5.3b}$$

Schematics of optical response to optical in amplitude, A, and phase, ϕ, are presented in Figure 5.2. A different representation of the phase that gives a greater response for a given change in resonant frequency is that of the phase measured from the centre of the resonance circle (I_c, Q_c) in the IQ plane, given by

$$\phi(t) = \tan^{-1}\left(\frac{Q(t) - Q_c}{I(t) - I_c}\right) , \tag{5.4}$$

and indicated in the right hand part of Figure 5.2. In practice, calculating the phase is computationally expensive due to the inverse tangent operations. A more convenient
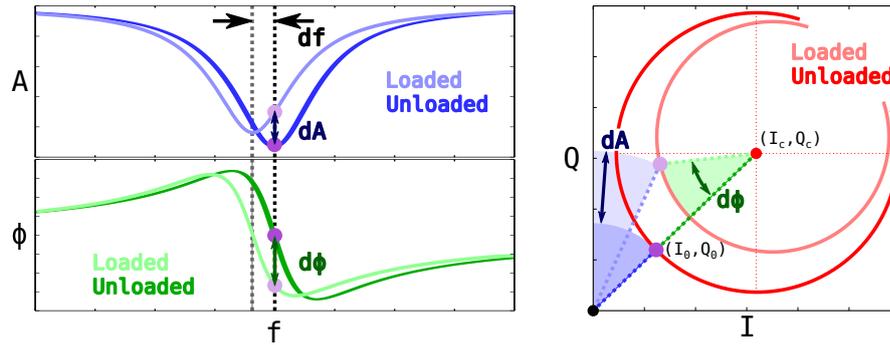
Figure 5.2: *KID response: a fixed frequency readout tone indicated by the dotted black line takes on the amplitude and phase values indicated by the purple dots. When the KID resonance changes due to an optical load, the amplitude and phase measured at the probe frequency will vary as indicated by the blue and green arrows.*

proxy for the detector frequency shift, referred to as δf, is given by

$$\delta f(t) = \frac{(I_0 - I(t)) \left.\frac{dI}{df}\right|_{f_0} + (Q_0 - Q(t)) \left.\frac{dQ}{df}\right|_{f_0}}{\left.\frac{dI}{df}\right|_{f_0}^2 + \left.\frac{dQ}{df}\right|_{f_0}^2} \quad \text{[Hz]} , \tag{5.5}$$

where I_0 and Q_0 are the optically unloaded I and Q values at the readout tone frequency; $\left.\frac{dI}{df}\right|_{f_0}$ and $\left.\frac{dQ}{df}\right|_{f_0}$ are the gradients of the optically unloaded $I(f)$ and $Q(f)$ sweeps at the readout tone frequency; and, $f = \omega/2\pi$. This approximation to the frequency shift of the resonator involves computationally inexpensive, linear, algebraic operations and is valid in the limit of frequency shifts lower than a couple of line-widths. A derivation is provided in Appendix A.1.

Calculation of the $\phi(t)$ and/or $\delta f(t)$ values requires prior knowledge of the location of the resonance circle centre in the IQ plane and/or the gradients of the $I(f)$ and $Q(f)$ curves. As such, frequency sweeps must be performed and the values extracted prior to the start of any fixed frequency probe readout.

Everything described so far for single KID readout also holds for multiple KID readout. The only difference is that the single-tone probe signal is replaced by a super-position of tones where each tone frequencies matches one of the KID resonant frequencies. That is,

$$s(t) = \sum_{i=1}^{N} a_i e^{j(\omega_i t + \varphi_i)}, \tag{5.6}$$

where N is the number of channels, a_i is an individual tone's amplitude, ω_i is an individual tone's frequency, and φ_i is the phase offset for an individual tone.

When this multiplexed signal is fed into the KID array's transmission line, the resonators are free to modulate their corresponding tones. Once the signal has left the array we must demultiplex it to extract the individual modulations. This is generally achieved by digitising the signal and applying either an FFT or an array of lock-ins to
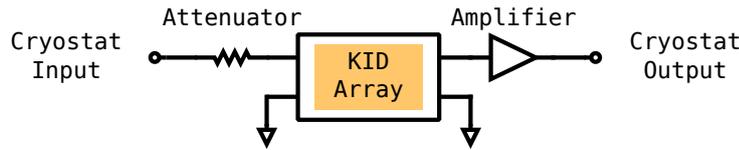
Figure 5.3: *A schematic of the cryogenic electronics requirements for a typical KID system. Theoretically, 1000s of KIDs can be read out through a single cryogenic amplifier before saturation becomes a problem.*

extract the I and Q components of the individual readout channels, and then performing the appropriate signal transformations (see Equations 5.3-5.5) to compute the individual detector responses.

5.1.1 Cryogenic electronics

Kinetic inductance detectors are superconducting devices that require cooling to sub-Kelvin temperatures for operation. The main advantage that KIDs have over other low temperature detectors is that the frequency domain multiplexed readout approach requires only one physical line into the detector array and only one line out. Other contemporary low temperature detector technologies generally require $\mathcal{O}(N)$ wires where N is the number of detectors. The incorporation of modern cryogenic time-division multiplexing systems can reduce this to around $\mathcal{O}(\sqrt{N})$, however, these systems are not cheap, are not always 100% reliable, and still require considerable cryogenic engineering efforts. KID systems, therefore, have to the potential to match the performance of other modern systems at considerably lower costs with, potentially, greater reliability and vastly reduced overall complexities.

The power levels of KID readout tones have to be sufficiently low so as not to drive electrical current densities above the superconducting critical point. The levels required by the KIDs in our system are typically -70 to $-60\,\mathrm{dBm}$ or 10 to $100\,\mathrm{pW}$ – which is relatively high compared to other KID systems. In order to maintain good dynamic range in the readout tones, they are generated at considerably higher levels than this and attenuated down prior to reaching the detector arrays. The tones then require amplification before they are digitised and processed. The attenuation and amplification are best performed within the cold stages of the cooling platforms as the reduced noise temperatures maintain good dynamic range in the probe signal. The Caltech CITLF4 amplifiers used in this camera have noise temperatures of around $5\,\mathrm{K}$ between 0.1 and $6\,\mathrm{GHz}$ when operated at temperatures below $20\,\mathrm{K}$. Broad-band attenuators that operate at typical KID frequencies commonly available but good quality low-noise cryogenic amplifiers cost at least a few thousand pounds and have hold times of many months. A schematic of typical cryogenic requirements for a KID readout system is given in Figure 5.3.
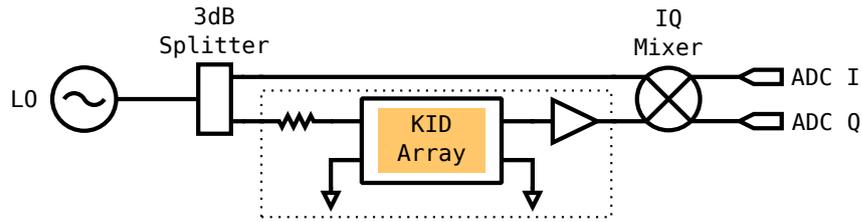
Figure 5.4: *A schematic of a homodyne readout system.*
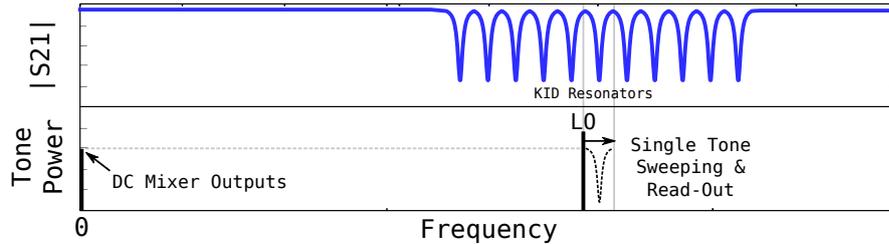
Figure 5.5: *Homodyne sweeps are performed by repeatedly stepping the LO frequency until the desired bandwidth has be probed. The DC mixer outputs trace out the S_{21} curves.*

5.1.2 Homodyne readout

The simplest practical way to read out a KID device is with the homodyne mixing system illustrated in Figure 5.4. A variable frequency signal generator or local oscillator (LO) outputs a readout tone which is split in two by a power divider. One half of the signal is passed through the detector array and the other is not. The two are combined in an IQ mixer and the resultant I and Q components are digitised. Because the LO and RF input ports of the IQ mixer receive the signals of the same frequency, the I and Q outputs will be DC levels with $I + jQ$ proportional to the S_{21} of cryogenic system.

KID resonators can be characterised, one at a time, as a function of frequency by stepping the LO frequency while synchronously sampling the ADCs as indicated in Figure 5.5. Single channel fixed tone readout is trivial and – with fast enough ADCs – enables the measurement of KID quasiparticle lifetimes.

The downside to the homodyne readout system is that multiplexing is impossible and only one KID can be measured at a time. Although, full array characterisation can still be achieved on sensible time-scales with a decent automatic measurement and control system.

5.1.3 Heterodyne readout

A heterodyne readout system makes multiplexed readout possible when the KID resonant frequencies are beyond the reach of the digital electronics – as has been the case with most KID system to date. Initially, a digital multiplexed probe signal is computed and converted to analog with a DAC. The tone frequencies are limited to some intermediate frequency (IF) range determined by the DAC bandwidth. This IF signal is up-mixed
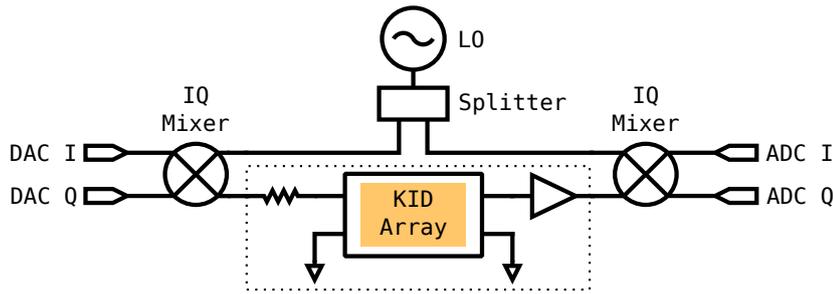
Figure 5.6: *A schematic of a heterodyne readout system. Image rejection relies on a 90°phase shift present between each reaout tone's I and Q signals.*
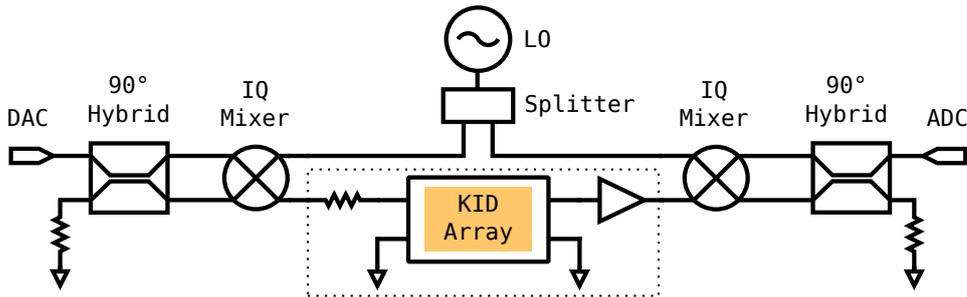
Figure 5.7: *Another schematic of a heterodyne readout system. This time, image rejection is performed by combining IQ mixer IF ports with quadrature hybrid couplers.*

with an LO to the required RF frequencies and passed into the KID array. The output is then down-mixed, with the same LO, back to the IF band where it is digitised and demultiplexed.

Care must be taken to reject unused sidebands when mixing, otherwise unwanted noise will be mixed back down on to the readout tones. If two DACs are available, sideband rejection in the IQ mixer can be established by ensuring, digitally, that the required 90° phase offsets are present between the I and Q signals. The heterodyne readout system in Figure 5.6 relies on this method for image rejection. Sideband rejection is also possible using analog components as illustrated in Figure 5.7. The quadrature hybrid coupler on the left hand side splits the probe signal into I and Q components with the correct phase offsets to ensure image rejection in the IQ mixer.

Sweeping KIDs with a heterodyne readout system is pleasingly straightforward. Changing the LO frequency instantly changes the frequencies of all of the up-mixed tones. By sweeping just the LO frequency, all of the KID resonators in an array can be swept out simultaneously. This is indicated in Figure 5.8.

5.2   The Altera readout system

The Altera readout system is a complete frequency domain multiplexing readout system with both analog and digital electronic components. The system can generate, acquire and process multiplexed readout signals with up to 150 tones in up to 80 MHz of IF
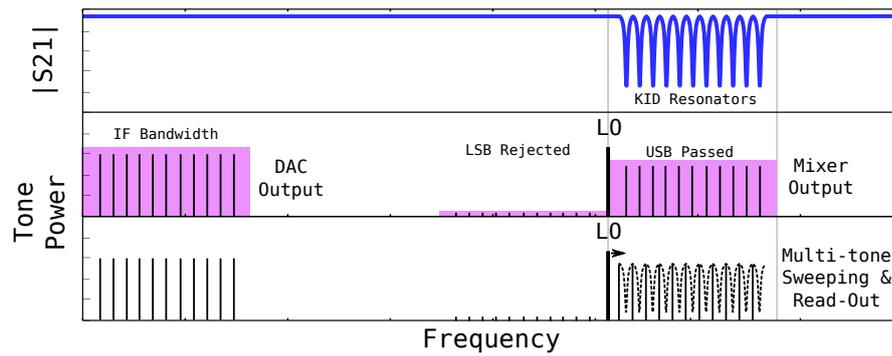
Figure 5.8: *A multiplexed heterodyne readout system enables simultaneous multi channel sweeping simply by shifting the frequency of the local oscillator. Suppression of either the upper sideband (USB) or the lower sideband (LSB) is vital to reduce excess noise after down-mixing.*

bandwidth. The digital aspect is based on an Altera Stratix II DSP Development Kit which was initially developed as part of my Masters research. The Altera Corporation[1] and their University Program is gratefully acknowledged for the donation of the FPGA kit and the associated development software. The analog system includes an external LO input, a pair of IQ mixers and pair of hybrid couplers for RF signal conversion. The single-sideband heterodyne mixing circuit is identical to that described in Figure 5.7. The primary RF bandwidth is limited by the IQ mixers 1.5 to 4.5 GHz, which is suitable for readout of the KIDs used in our camera. An additional set of mixers with 4.5 to 9 GHz of bandwidth are included in the system that are activated by switching an external set of jumpers, should the need for higher frequency readout arise. The readout system also houses an encoder readout circuit for monitoring the orientation of the field scanning mirror. The contents of the readout system are visible in the photograph in Figure 5.9.

The digital signal processing system is implemented on a Stratix II FPGA and was developed with the Quartus II IDE. The FPGA device came mounted to a development board that contains amongst other things, two 160 MSPS DACs, two 140 MSPS ADCs, an 10/100 Ethernet port, a USB port, a VGA port, programmable LEDs, flash memory, RAM, a 100 MHz crystal oscillator, numerous TTL input lines, and a JTAG port for firmware uploading. The FPGA architecture was developed with subsystems for multiplexed readout tone generation, analog and digital conversion, Fourier analysis, channelisation, averaging/decimation, data transfer over Ethernet, absolute encoder readout, and other general purpose inputs and outputs (GPIO):

- Tone generation is implemented with a circular RAM buffer connected to the DAC. The multiplexed readout signal is computed externally and copied into the buffer. The contents of the buffer are then continuously and repeatedly passed - element by element - to the DAC. The waveform upload process is not instantaneous due to the large size of the buffer, but it does allow for the generation of arbitrary

---

[1]Altera University Program:
https://www.altera.com/support/training/university/overview.html
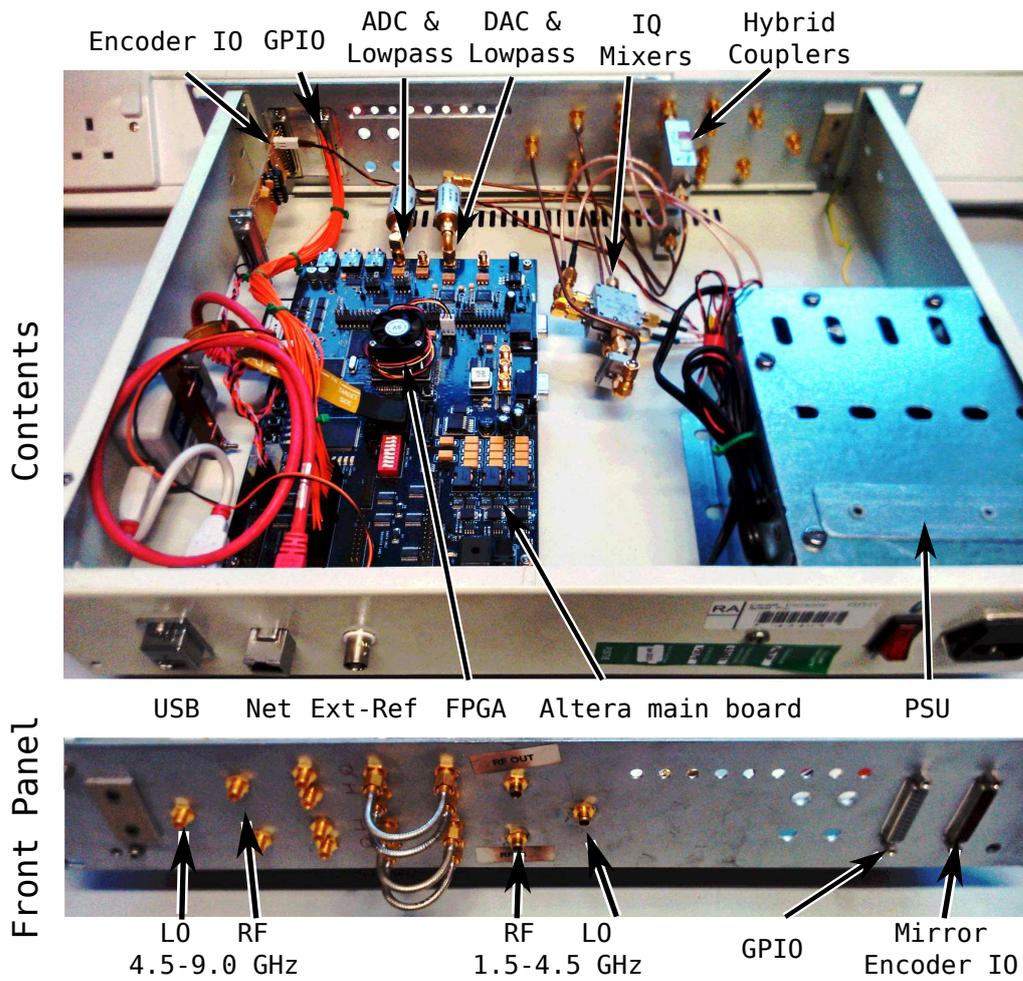
Figure 5.9: *A photo of the contents of the Altera multiplexed readout system.*

waveforms including single or multiple tones of any frequency, amplitude and phase, as well as broad band chirps or even white noise.

- Fourier analysis is implemented with a propriety realtime FFT engine. The FFT can be programmed with a range of different sizes and data types. After a short initial hold-period, it outputs complex spectra at the same rate as data is input.

- Channelisation and decimation are implemented directly into the FPGA logic. FFT frequency bin selection is performed based on a list of tone frequencies that must be uploaded to the system. Complex output channels are stored in memory and coadded together an arbitrarily set number of times before passing to data transfer server.

- An embedded, reduced instruction set microprocessor based on the NIOS architecture is implemented with the µC/OS-II real-time operating system. The OS kernel is implemented with a number of features, including an Ethernet driver with custom *NicheStack* TCP/IP network stack for communication with the control station, and a DMA controller for direct data access. A user space task stack is set up with a command server and a data server. The command server listens for and actions requests from the control station. The data server generates and sends a continual stream of data packets onto the network for acquisition by the control station. Data packets contain demultiplexed IQ data, mirror encoder data and other auxiliary information.

- The base clock for the system comes from the 100 MHz crystal oscillator. A proprietary PLL routine is used to multiply this up to 160 MHz for driving the DACs. A further 80 MHz derived clock is generated to drive the ADC, the FFT engine and the channeliser/decimator, while a separate clock is in place to drive the embedded processor.

The limited IF bandwidth of the Altera system meant that full array readout was not initially possible. The 60 pixel, 150 GHz array was spread out over 120 MHz centred at around 1.5 GHz but the Altera system has only 80 MHz of instantaneous bandwidth. The FPGA firmware system was upgraded to implement an additional time division multiplexing (TDM) system. The tone generation, FFT, channelisation, and averaging subsystems were synchronised with the local oscillator via an external trigger input on the synthesiser. Then, at regular intervals, the LO would switch frequency and intermittently change the RF coverage of the KID array. This successfully enabled full array readout but at the expense of reducing the overall sample rate of the system. The TDM system implementation allowed for an arbitrarily large number of multiplexing divisions. This meant that full readout of the 150 pixel, 350 GHz array became possible – up to 8 separate sub-bands were utilised in the characterisation of that array. A schematic of the TDM approach is illustrated in Figure 5.10
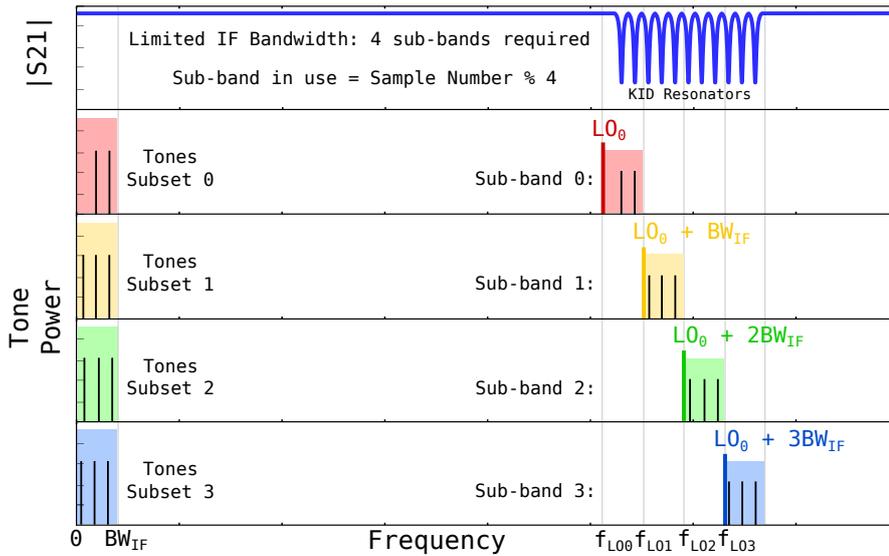
Figure 5.10: *A schematic of the process of time division multiplexing of a multiplexed heterodyne readout system. A time-division factor of 4 is presented in this example.*

The overall demultiplexed sample rate of the Altera system was typically 150 Hz, which meant only 75 samples per second with a 2-channel TDM in place. This was eventually upgraded to 750 Hz which enabled full 350 GHz array readout with 5-chanel TDM at 150 samples per second. While not fast enough for video rate imaging, this rate was sufficient to enable the development and testing of the real-time imaging software.

5.3 The NIKEL readout system

The NIKEL readout system was acquired as a replacement for the Altera system. The system was developed originally for the readout of the NIKA millimeter-wave astronomical KID camera ([Bourrion et al., 2011], [Bourrion et al., 2012], and [Monfardini et al., 2014]) and was adapted to work with our terrestrial terahertz imaging camera. The hardware specifications are a vast improvement on the Altera system with 1 GSPS ADCs and DACs providing 500 MHz of IF bandwidth, the capability for multiplexed generation and subsequent processing of up to 400 readout tones, and a full demultiplexed output sample rate of up to 500 Hz.

The heterodyne mixing system is implemented in a method similar to that described in Figure 5.6 but with a single ADC input. The second IF port of the final IQ mixer is terminated with a matched load meaning that sideband rejection is only implemented in the up-mixing stage. The readout system was adapted to include IF and RF amplifiers and RF variable attenuators for gain control of the incoming and outgoing multiplexed probe signals. Further additions included an encoder readout system for the scanning mirror, and a controller to drive an LED calibration source that is located within the cryostat. A photo of the contents of the readout system is included in Figure 5.11.

The digital system is based on a custom designed main board with a combination of
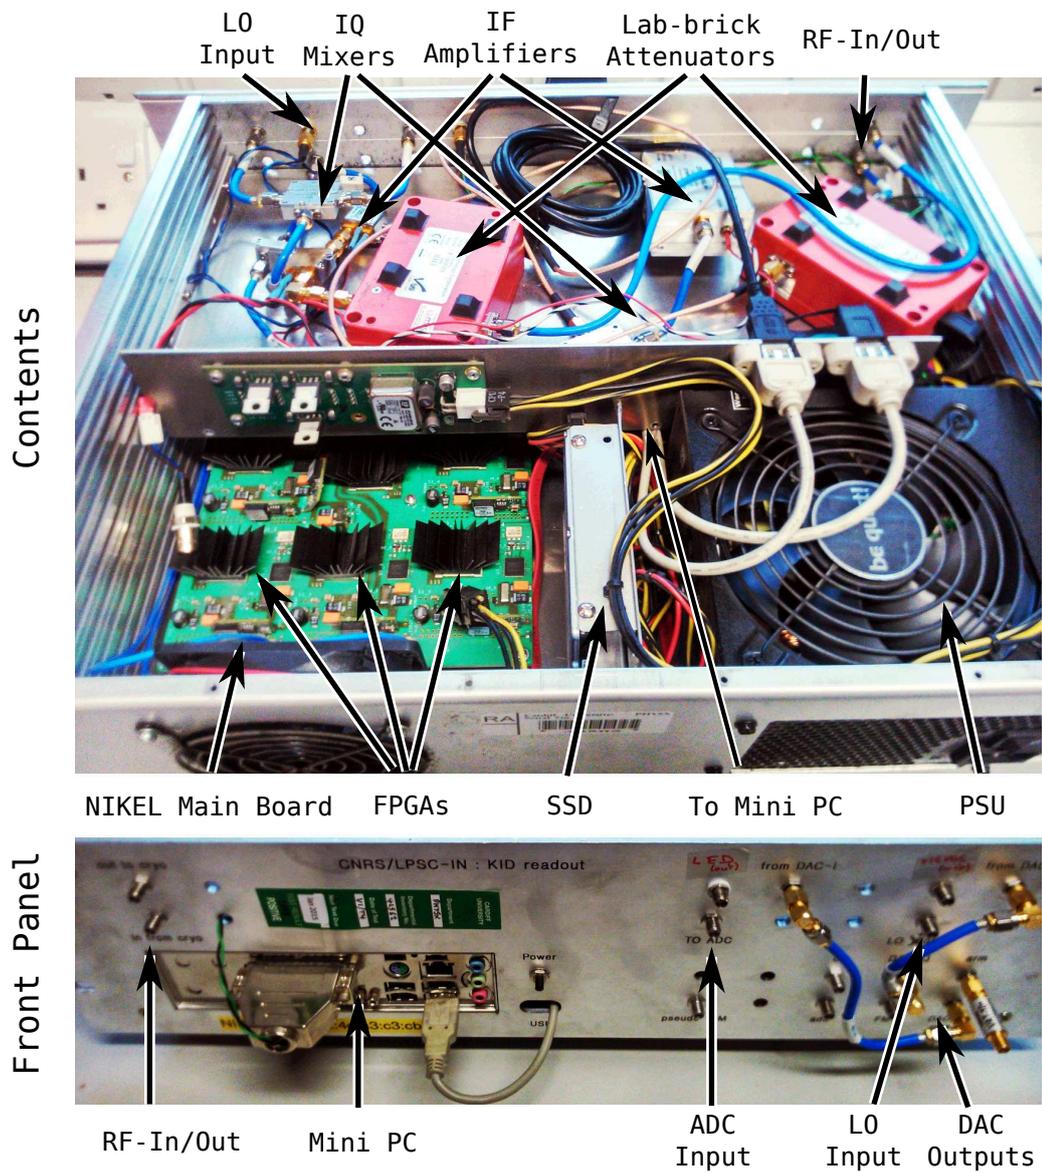
Figure 5.11: *A photo of the contents of the NIKEL multiplexing system.*

six Xilinx Virtex-6 FPGAs with a single ADC, five DACs plus a combiner, a reference clock and PLL, DC/DC power converters and a USB controller for communication with a mini PC that is housed with readout unit. The NIKEL provides an optional means for multiplexed output signal modulation but we have not implemented this in our system. Further details are available in [Bourrion et al., 2012]. Many thanks are due to Dr Bideaud for his efforts in setting up software interfaces on the NIKEL server system.

Some key differences from the Altera system include the tone generation system, the spectral analysis, and the command and data server systems:

- Tone generation is implemented digitally with arrays of CORDIC sin/cos-generators and digital attenuators for gain adjustment. Five FPGAs can generate up to 80 tones each which are passed to five corresponding DACs prior to combination and heterodyne mixing.

- Spectral decomposition is performed with a five-channel polyphase-filter-bank and direct digital downconversion (DDC) mixers. The system is similar to a array of lock-in amplifiers and is applied digitally with a series of finite impulse response filters, frequency shifters and multipliers, to match the input and reference frequencies for each tone.

- The main control system for NIKEL electronics is located in the mini PC that is directly connected to the main readout board. This system was modified to listen and action request from the main camera control system. An API was developed to control the various components of the system, including for example to tone generation, or frequency sweeping. A data server was set to up to gather the I and Q data for each tone and stream them over the network to the control station. Auxiliary data such timestamps, tone frequencies, mirror positions, etc., are included in the data packets.

The NIKEL system was something of a black-box when it arrived and many aspects of its operation were not, initially, well understood. A full characterisation of the output multiplexed readout signal was performed to ensure the correct placement of tones into the RF band. Spectrum analyser measurements of the DAC outputs were carried out to characterise the tone powers and frequencies. Details of the five sub-bands contained in the DAC outputs are described in Table 5.1.

Table 5.1: *Frequency limits of the 5 NIKEL DAC bands*

	Measured		Recommended	
	Start [MHz]	Stop [MHz]	Start [MHz]	Stop [MHz]
DAC-0	$0.185^\dagger$	$120^\dagger$	1	100
DAC-1	87.5	$230^\dagger$	100	200
DAC-2	187.5	$315^\dagger$	200	300
DAC-3	287.5	$410^\dagger$	300	400
DAC-4	387.5	499.9985	400	499

† Not a sharp start/stop. Values given are of the $3\,\mathrm{dB}$ cut-on/off.

Tones are set in the NIKEL by specifying a tone frequency f_{Tone}, a tone gain G_{Tone} and a DAC band. Each DAC band also has a variable gain G_{DAC} that affects all the tones in that band. A model, based on empirical measurements, was developed to estimate the output tone power given the requested tone parameters, the model is given in Equation 5.7 and the measurements are given in Figure 5.12. Derived coefficients describing the linear relationship between G_{DAC}, G_{Tone} and the measured RMS tone voltages are given in Table 5.2, along with power estimates, power measurements and the error for a set of example tones.

$$
P_{\mathrm{Tone}} = 10 \log_{10} \left(\frac{1000}{50} \left[G_{\mathrm{Tone}} G_{\mathrm{DAC}} \frac{dV_{\mathrm{RMS}}}{dG_{\mathrm{DAC}}} \right]^2 \right) \pm 0.5 \quad [\mathrm{dBm}] \tag{5.7}
$$

Table 5.2: *Tone power estimation for the NIKEL electronics. The gains used in this example are $G_{\mathrm{DAC}} = 10000$ and $G_{\mathrm{Tone}} = 5$.*

	$dV_{\mathrm{RMS}}/dG_{\mathrm{DAC}}$	Estimated Power	Measured Power	Error
	[V/unit]	[dBm]	[dBm]	[dB]
DAC-0	8.09×10^{-8}	-34.85	-34.57	$+0.28$
DAC-1	6.95×10^{-8}	-36.17	-34.57	-0.63
DAC-2	8.67×10^{-8}	-34.25	-33.81	$+0.44$
DAC-3	5.64×10^{-8}	-37.98	-38.23	-0.25
DAC-4	4.19×10^{-8}	-40.57	-40.33	$+0.24$
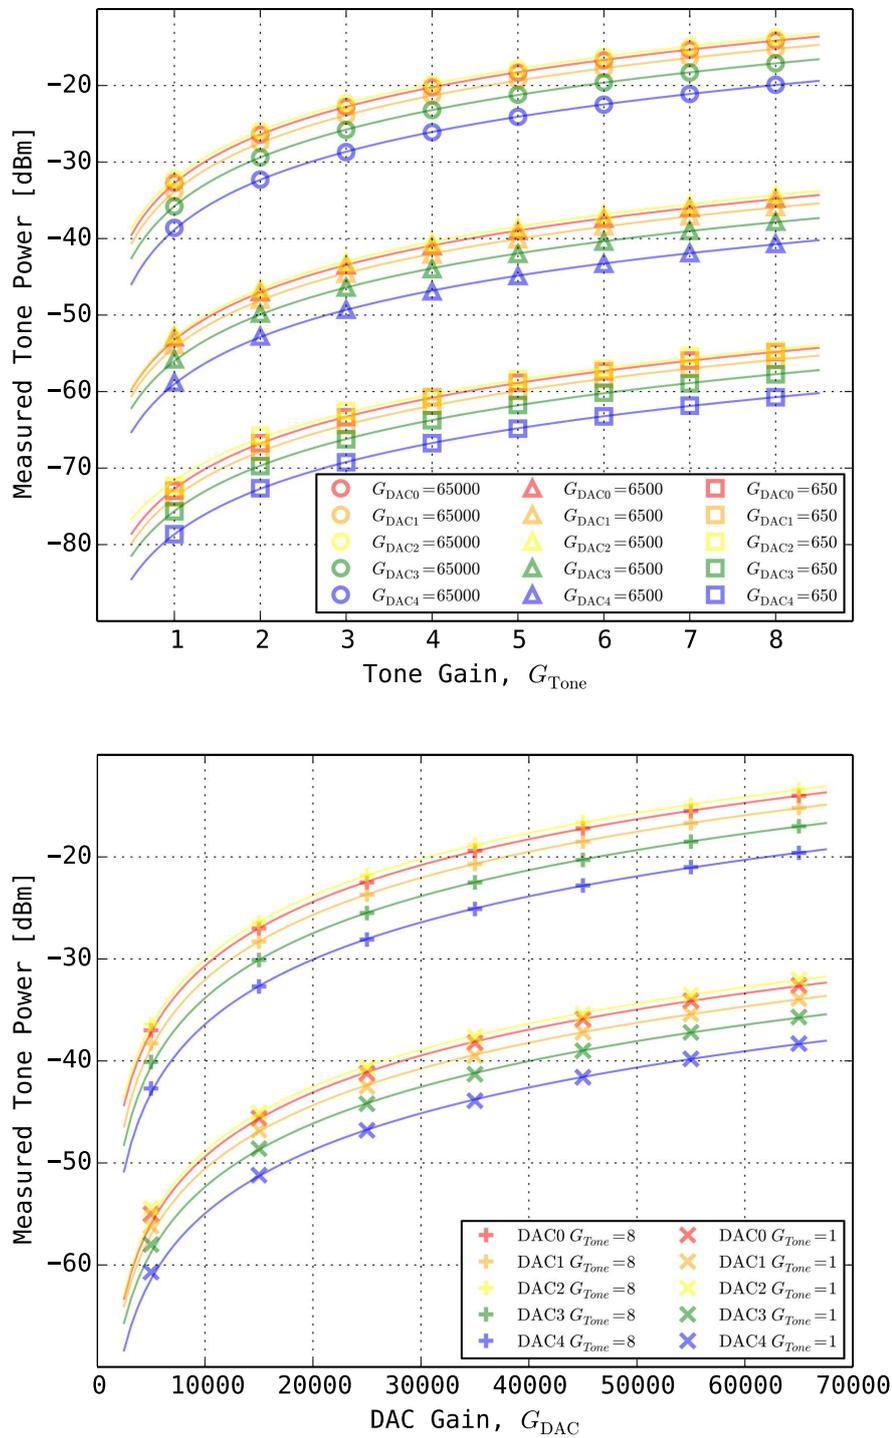
Figure 5.12: *Measurements (points) and results from the power estimation model (curves) of the readout tone output power measured from the DAC-I port on the NIKEL electronics main board.*

5.4   Amplifier noise

The power levels of the tones in the probe signal used to read out the LEKIDs in this system are typically $100\,\mathrm{pW} < P_g < 1\,\mathrm{nW}$ (-70 dBm$< P_g <$-60 dBm) so it is necessary to insert microwave attenuators along the RF chain to reduce the readout powers sufficiently. As mentioned previously, these are fairly high levels when compared to other KID systems, however, they are still sufficiently low to require amplification prior to sampling. Inclusion of a cryogenic, high gain, low noise amplifier (LNA) close to the detector array improves the dynamic range in the outgoing probe signals by taking advantage of the reduced Johnson noise at low temperatures. The power spectral density of the voltage noise imparted by an amplifier with input impedance Z_0 (typically $50\,\Omega$) and noise temperature T_a is given by

$$S_V = 4k_B T_a Z_0 \quad \left[\mathrm{V}^2\,\mathrm{Hz}^{-1}\right] . \tag{5.8}$$

In computing an NEP for the amplifier noise, it is convenient to convert the voltage noise into the corresponding noise in the detector response. In terms of the resonator fractional frequency shift $x = (f - f_0)/f$ it can be shown that

$$S_x = S_V \left(\frac{dx}{dV}\right)^2 = \frac{k_B T_a Q_c^2}{P_g Q_r^4} \quad \left[\mathrm{Hz}^{-1}\right] , \tag{5.9}$$

where Q_c and Q_r are the coupling and resonator quality factors of the detector under test, and P_g is the power of the readout tone at the input to the detector. The responsivity to optical power in this scheme is

$$R_x = -\frac{dx}{dP_{abs}} = \frac{\beta \eta_e \tau_{qp}}{2\Delta N_{qp} Q_i} \quad \left[\mathrm{W}^{-1}\right] , \tag{5.10}$$

where Q_i is the internal quality factor of the KID resonator and β is the ratio of the real and imaginary parts of the complex conductivity of the superconductor. For most materials, $\beta > 1$ and can be measured by comparing the dissipative (real) and frequency (imaginary) responses of a resonator under a range of loads.

Finally, the NEP, or the equivalent variance in absorbed optical power that would be required to equal the noise power due to the thermal fluctuations in the cold amplifier is

$$NEP_{\mathrm{amp}} = \frac{Q_c Q_i}{Q_r^2} \frac{2\Delta N_{qp}}{\beta \eta_e \tau_{qp}} \sqrt{\frac{k_B T_a}{P_g}} \quad \left[\mathrm{W}\,\mathrm{Hz}^{-\frac{1}{2}}\right] , \tag{5.11}$$

Given this approach and taking typical parameters for the detectors in our system, the expected amplifier noise equivalent power is computed to be $NEP_{\mathrm{amp}} = 8.5 \times 10^{-18}\,\mathrm{W}/\sqrt{\mathrm{Hz}}$.

5.5 Digitisation noise

The digital-to-analog and analog-to-digital converters (DACs and ADCs) that are used in KID readout systems have limited dynamic range and so it is important to make sure that any signals do not saturate these components or that quantisation noise is not dominant. A common measure of the signal-to-noise or dynamic range of any signal, especially from a DAC or ADC, is the phase noise. The single sided phase noise spectral density $S_\varphi(f)$ of a time domain carrier signal $x(t)$ is defined as the signal variance relative to the carrier power measured in a $1\,\mathrm{Hz}$ band offset from the carrier frequency by some amount f. It is usually quoted in units of decibels relative to the carrier per Hertz (dBc/Hz) and is closely related to jitter or timing/phase offset variations in the timestream.

In general, in a KID system, the noise introduced by the readout electronics should not be higher than the noise from the cryogenic amplifier. An amplifier with noise temperature T_a and N input tones of power P_g will impart a phase noise equal to

$$S_{\varphi,\,\mathrm{amp}} = 10 \log_{10} \left(\frac{2 k_B T_a}{N P_g} \right) \quad \left[\mathrm{dBc\,Hz^{-1}} \right] . \tag{5.12}$$

This phase noise reduces with increased number of tones or increased power per tone, placing more stringent requirements on the dynamic range of the remainder of the readout system. Generally, state-of-the-art DACs outperform state-of-the-art ADCs and have higher bit counts for a given sample frequency, so the ADC becomes the limiting component. For example the 16-bit, $1\,\mathrm{GSPS}$ DACs in the NIKEL electronics used for this project have quoted noise spectral densities (NSDs) of -174 dBFS/Hz, while the 12-bit, $1\,\mathrm{GSPS}$ ADCs have -147 dBFS/Hz. The dBFS unit is the decibel at full scale, meaning the NSD was measured on a carrier that fills the full dynamic range (uses all the bits) of the converter. Ideally,

$$\mathrm{NSD_{FS,ADC}} \leq S_{\varphi,\,\mathrm{amp}} - 3\,\mathrm{dB} . \tag{5.13}$$

It is becoming more common for DAC and ADC manufacturers to quote noise spectral densities for their components but, more often than not, only signal-to-noise ratios (SNRs) and other related figures of merit are presented. SNR in dBFS can easily be converted to a noise spectral density. First, invert the SNR to compute the quantisation noise floor at the full scale from a $0\,\mathrm{dB}$ reference, then divide by the Nyquist bandwidth. For example, an ADC sampling at f_s samples per second has a bandwidth Δf_s and

$$\mathrm{NSD_{FS}} = 0\,\mathrm{dB} - \mathrm{SNR} - 10 \log_{10} \Delta f_s \quad \left[\mathrm{dBFS\,Hz^{-1}} \right] . \tag{5.14}$$

To calculate an absolute level for the floor of the NSD, a reference to the full scale range is required. If the quoted ADC full scale input voltage is V_{FS} volts then

$$\mathrm{NSD_{dBV}} = 20 \log_{10} V_{\mathrm{FS}} - \mathrm{SNR} - 10 \log_{10} \Delta f_s \quad \left[\mathrm{dBV\,Hz^{-1}} \right] , \tag{5.15}$$

where $1\,\mathrm{dBV}$ is $10^{1/20}\,\mathrm{V}$ and V_{FS} is typically 1.5-2 V. In a KID readout system, the actual noise floor will be higher than this due to the presence of multiple tones in the readout signal. If N tones are present,

$$\mathrm{NSD}_{\mathrm{MUX}} = \mathrm{NSD}_{\mathrm{dBV}} + 10\log_{10} N \quad \left[\mathrm{dBV\,Hz}^{-1}\right] . \tag{5.16}$$

For an ideal converter, fluctuations in the least significant bit (LSB) set the quantisation noise limit. The full scale noise floor or the full scale dynamic range is set by the resolution, ie. the number of bits N_b, as

$$\mathrm{SNR}_{\mathrm{ideal}} = 20\log_{10} 2^{N_b} + 10\log_{10}\tfrac{3}{2} \quad [\mathrm{dBFS}] . \tag{5.17}$$

In practice, a converter has inherent noise and jitter that reduces the SNR. As such, devices are specified with an effective number of bits (ENOB) which relates the actual SNR performance to the ideal performance, where

$$\mathrm{SNR}_{\mathrm{actual}} = 20\log_{10} 2^{\mathrm{ENOB}} + 10\log_{10}\tfrac{3}{2} \quad [\mathrm{dBFS}] . \tag{5.18}$$

For example, consider the 12-bit, $1\,\mathrm{GSPS}$ ADC with 9.5 ENOBs used in the NIKEL readout electronics. There are 2^{12} possible output levels across a $500\,\mathrm{MHz}$ bandwidth. It ideally would provide $\sim$74 dBFS signal-to-noise, so the noise spectral density would be -161 dBFS/Hz. In practice, the actual signal-to-noise is only $\sim$59 dBFS and the noise spectral density is -146 dBFS/Hz. Signal processing that occurs after digitisation can reduce the noise level below the quantisation limit, for example, channelisation and decimation with finite impulse response filters can serve to reduce the noise.

As with the previous noise sources, we desire a readout NEP or a limiting sensitivity due to readout electronics noise. The method for estimation is similar to that used in the amplifier NEP calculation. We already have the voltage noise power spectral density due to quantisation in the ADC, which we convert out of dB units,

$$S_V = 10^{\mathrm{NSD}_{\mathrm{MUX}}/10} \quad \left[\mathrm{V}^2\,\mathrm{Hz}^{-1}\right] . \tag{5.19}$$

Then for a given set of detector parameters this is converted to a variation in the detector response S_x as

$$S_x = S_V \left(\frac{dx}{dV}\right)^2 = 10^{\mathrm{NSD}_{\mathrm{MUX}}/10}\frac{Q_c^2}{4Q_r^4 P_g Z_0} \quad \left[\mathrm{Hz}^{-1}\right] , \tag{5.20}$$

where Q_r and Q_c are the resonator and coupling quality factors, P_g is the power in the readout tone at the input to the detector, and Z_0 is the impedance of the line at the input to the detector. The responsivity to optical power is the same as in Equation 5.10, so the NEP, or equivalent absorbed optical signal power required to equal the readout

noise in a 1 Hz bandwidth is

$$NEP_{\text{readout}} = \frac{Q_c Q_i}{Q_r^2} \frac{\Delta N_{qp}}{\beta \eta_e \tau_{qp}} \sqrt{\frac{10^{\text{NSD}_{\text{MUX}}/10}}{P_g Z_0}} \quad \left[\text{W Hz}^{-\frac{1}{2}} \right] . \qquad (5.21)$$

A number of other factors can affect the noise such as improper filtering, saturation of active devices, mains pick-up, and a variety of other kinds of pick-up, however these are (often) easily remedied. The computed limit for the digitisation noise equivalent power expected in this camera for a typical set of detector parameters is $NEP_{\text{readout}} = 2.8 \times 10^{-18}\,\text{W}/\sqrt{\text{Hz}}$.

5.6 Summary

A review of some the important methods for single- and multi-tone KID readout has been presented, including cryogenic requirements, homodyne mixing systems and heterodyne mixing systems.

The two electronic readout system developed for our terahertz camera have been introduced and some of the key details of their modes of operation have been described. Particular focus was made on the implementation of a time division multiplexing system to extend the effective bandwidth of the Altera system. This technique could be effective in other systems where bandwidth is limited. Additional attention was made to the tone generation system of the NIKEL electronics system. Accurate knowledge of tone powers is vital for the effective operation of a KID readout system.

Noise sources and limiting sensitivity estimates from the cold amplifier and the digital electronics system were estimated and a further discussion of the contribution to the overall camera sensitivity is given in chapter 7.

Chapter 6

The Software System and Image Generation

6.1 Overview

The successful development of the terahertz camera required the engineering of a considerable amount of computer software, ranging from low level hardware controls and data manipulation routines to higher level application design and image processing algorithms. Software systems in general can be viewed from a variety of different perspectives including those of end-users, developers, testers, maintainers, etc. In this chapter I focus on some of the key software components from a mostly technical standpoint in order to communicate the details of component design and implementation.

As far as any end-user/camera operator is concerned, the application software required to switch the camera on and take pictures should be straightforward and easy to use. Assuming that the cryogenic system is already cold, the procedure to begin imaging is simply to power up the electronics and field scanning mechanism, log-in to the control station, start the main camera application, and then wait for initialisation and calibration before the real-time video appears on the display. Interaction with the GUI enables image contrast adjustment, single frame snapshots, and a frame stacking mode where new frames are continuously averaged together in order to improve the SNR. Further options are available by modification of system configuration files prior to running the main application. Adjustment of these advanced settings can affect the image output format, calibration parameters, timeline and/or image filtering, the frame update rate and the integration time per frame.

In actuality, the underlying system is a bit more complex than this. A combination of command line interfaces, graphical interfaces, configuration files, Labview VIs, compiled and interpreted programs, embedded systems firmwares, and proprietary software applications are utilised to control all aspects of the camera. A view of the general architecture is given in Figure 6.1 where the system is broken down into four key concurrently running components:
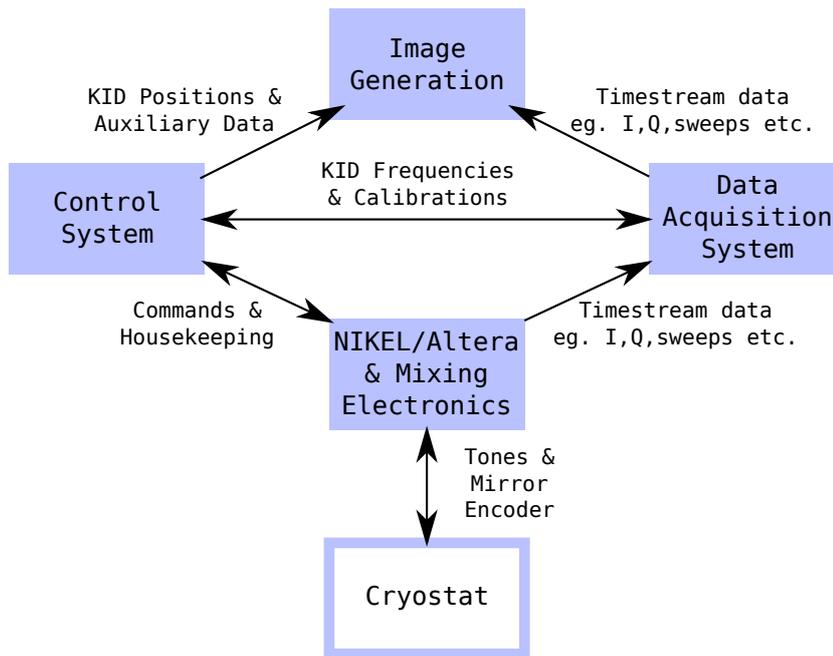
Figure 6.1: *Overview of the software architecture*

- The control system is responsible for setting up and running the camera as a whole, especially with regard to controlling the readout electronics systems, initialising and starting data acquisition, performing frequency sweeps and calibrations, and providing other general information to the image generation system.

- The data acquisition system is responsible for reading and saving raw data streams from the readout electronics. It provides a means for direct access to timelines of detector response and other information channels.

- The readout electronics firmwares are responsible for the generation and acquisition of the multiplexed probe signals that enter the cryostat as well as monitoring the scanning mirror orientation through the absolute encoder.

- The image generation system provides the user with a graphical interface that displays the real-time video. It has access to real-time and auxiliary data and has tools for calibration, timeline filtering, map making, image processing and real-time display.

6.2 Control system

This is, essentially, a collection of purpose built Python modules with functions to control and run all aspects of the camera. In this section I describe some of the most important functionalities, from basic operations to more complex procedures and algorithms.

6.2.1 Basic functionalities

Communications

Communication with the NIKEL and Altera electronics – as well as with other hardware devices such as local oscillators, variable attenuators and power supplies – is typically performed over Ethernet. Functions based around the Berkeley Sockets API are in place to create and manage IP network sockets, to establish reliable communications over TCP connections, and to send and receive multicast UDP datagrams[1].

Command servers in both the Altera (TCP based) and the NIKEL (UDP based) systems were set up to listen for data packets, to parse the contents for certain pre-defined strings, and to respond with the appropriate pre-programmed actions. Examples of the some of the most commonly used commands include requests for data transfers, requests for adjustment of readout tone frequencies or amplitudes, requests to change sample rates and averaging levels, requests to change the local oscillator frequency and requests to change the RF attenuator levels.

UDP data servers in the readout systems were set up to provide continuous, low latency streams of the demultiplexed detector outputs and other information channels. Data packet formats are described in the previous chapter and transfer rates reach up to 5 Mbit/sec from the Altera and 25 Mbit/sec from the NIKEL.

Initialisations

The control software system contains startup scripts to configure various hardware devices upon power-up and to setup and run the data acquisition and imaging software. The readout electronics systems are initialised remotely. The Altera system, including the on-board microprocessor and OS, is started over USB through the Quartus II sotware suite; the NIKEL system is started through an SSH session. Once the command server of the system in use has loaded, requests can be sent to initialise the system, setup the multiplexed readout tones and begin the transmission of the demultiplexed detector outputs. A standard, client-side initialisation script is outlined in Listing 6.1.

Setting tones

As described in the previous chapter, the tone generation systems vary considerably between the two readout systems. The Altera system utilises a single circular waveform buffer, while the NIKEL has multiple CORDIC-based generators spread across five DAC sub-bands. The software system handles requests for the setting of tones appropriately and, in both cases, the input lists of tone frequencies and powers are first validated against the known hardware limitations. With the Altera system, the multiplexed waveform is computed locally and uploaded to the FPGA – along with the list of tone frequencies so that the firmware knows which FFT bins to pass to the data server. With the NIKEL

---

[1]See the `IP(7)` and `SOCKET(7)` man pages.

```
1  from nikel import *
2
3  join()
4    #register as a client with the readout system
5  set_synth_hardware(HW_SYNTH_DEFAULT)
6    #specify which synthesiser is in use for this run
7  set_synth_freq(1.5e9)
8    #set the frequency of the chosen synthesiser to 1.5GHz
9  set_synth_power(16)
10   #set the output of the synthesiser to 16dBm
11 set_synth_state("ON")
12   #switch the sythesiser output on
13 set_attenuator_hardware(HW_ATTS_DEFAULT)
14   #specify which attenuators are in use for this run
15 set_attenuations(60,60)
16   #set the attenuation values to a safe level
17 set_averaging(1)
18   #set the number of data packets to average (reduces data rate)
19 set_tones(TONES_DEFAULT)
20   #set the tone frequencies and amplitudes to the default values
21 start()
22   #start streaming data packets
```

Listing 6.1: *A typical control system initialisation script for the NIKEL readout electronics.*

system, the input tones list is divided into five sub-lists of tones-per-sub-band and these sub-lists are uploaded to the FPGA where the CORDIC generators are updated immediately. Tone validation is extended to include the physical limits of each sub-band of the NIKEL system.

Raw data access

Control software modules have access to the data acquisition system's main storage structure as this is necessary for the user to view sweep data and perform such tasks as KID finding, IQ and tone power calibrations, and imaging. Collecting the raw data from the readout electronics is generally the job of the data acquisition system, however, the control system can also request access to this data. This is useful, for example, when performing frequency sweeps and for general system characterisation where individual data samples and even complete FFT outputs are required.

Frequency sweeping

Frequency sweeping of the readout tones across KID resonators is one of the most important jobs that the readout systems perform. Access to the sweep products is vital for two primary readout activities; first, finding the frequencies of KID resonators and second, characterising the KID responses in order to calculate the final calibrated timelines. When implementing a heterodyne mixing system (as is the case with the present

Altera and NIKEL systems) sweeping is straightforward in that only the local oscillator frequency need be shifted – all of the readout tone frequencies and demultiplexed output channels remain fixed during the course of the sweep. The control system commands and monitors the sweeping of the LO and has methods for collecting and processing the sweep data. Sweeps are defined with pairs of start/stop or centre/span frequencies and either a step-size or a number of points.

The sweep process generally involves discretely stepping the LO frequency while synchronously sampling the demultiplexed output. However, RF synthesisers have finite settling times after change of frequency and so, when collecting data, it is necessary to either introduce a time delay prior to manually sampling the readout or to ignore a fixed number of samples if manual sampling is not available. With the Altera system, data samples are requested between steps of the LO independently from the main data acquisition system and are easily collected and saved during the sweep process. With the NIKEL system, sweep samples are automatically flagged in the data acquisition system's main storage structure and are easily extracted by searching through the storage structure.

If the KID resonant frequencies are initially unknown, they can be identified by sweeping-out and inspecting a wide range of frequencies – this is similar to how a spectrum analyser operates. Sweeping across the full bandwidth of the KID arrays used in this system ($\sim$500 MHz) can be a very slow process. A typical sweep of 1 kHz step-size requiring at least 1000 seconds to complete if sampling at 500 SPS. This can be greatly sped up by sweeping with a number of tones spaced evenly across the IF bandwidth. The LO only needs to sweep a short range equal to the tone separation and the full bandwidth is sampled by concatenating the outputs from each of the readout tones. In the example given, the sweep time can be reduced to five seconds if sweeping with 200 tones.

When the KID frequencies are known, small sweeps about the KID centres are used to infer their properties and to calculate optical responses. Firstly, the LO frequency is set close to the first KID frequency, then multiple readout tones are added to line up with the KIDs' centre frequencies. The LO is then swept about its initial frequency and the readout electronics then returns the I,Q and frequency data for each KID. For every sweep taken, the data is saved, processed to extract the key parameters, and optionally displayed in either IQ, amplitude or phase format.

KID finding

Given a set of sweep data, including I, Q and frequency information, identifying the locations of the KID resonances is another fundamentally important process. A manual inspection of amplitude or phase sweeps is usually the most reliable way to identify the resonant frequencies but this takes a long time and can be a challenge if very low-Q and/or strongly coupled resonators are present. Automatic KID finding algorithms will

also struggle if the quality of the data is poor but they take practically no time at all.

A standard automatic algorithm involves applying initial digital filters to the raw I and Q data to remove high frequency noise and low frequency ripples. Then, an element-wise search through either the amplitude data or the gradient of the phase data highlights any regions beyond some given threshold parameter. The resonant frequency is found by searching each region for the position of its extremal value. This method by itself is fairly robust but is subject to false positives if any glitches are present in the sweep, so a deglitching pass is performed on the raw I and Q timelines prior to any digital filtering. This method is also prone to missing one or more resonances if they are closely coupled together – the stronger/deeper of the two resonances dominates and the weaker/shallower is skipped over.

6.2.2 Tone locking

Every time the camera is run, a new calibration sweep of the detectors is required. Previous calibrations are invalidated by a combination of low frequency, thermal variations in the electronics system gains, and small variations in KID resonant frequency caused by stochastic variations in magnetic field trapping in the superconductors during cool-downs. Two variations on a single method are implemented in the control software system to set readout tones onto KID resonances and perform the calibration sweep. This process is referred to as *locking-on* to the KIDs and the general algorithm is as follows:

1. Find the approximate KID frequencies.

   Variation 1: The KIDs are found with a full sweep, see Figure 6.2a.

   Variation 2: KID frequencies are taken from a list.

2. Set tones at the approximate frequencies and find precise KID frequencies with a short sweep that spans a typical KID bandwidth. See Figure 6.2b.

   Variation 1: All the data from the sweep is utilised.

   Variation 2: Selected regions of individual sweeps are excluded to isolate closely coupled KIDs from their nearby neighbours.

3. Set tones at the new frequencies, perform a second short sweep across the resonances and reset the LO to the centre of the sweep. See Figure 6.2c

   Variation 1: All the data from the sweep is utilised.

   Variation 2: Selected regions of individual sweeps are excluded to isolate closely coupled KIDs from their nearby neighbours.

4. Save sweeps and calibration parameters and apply them to incoming raw I and Q timelines to get the calibrated phases and/or frequency shifts.

The first variation was upgraded to the second variation in order to help isolate closely coupled resonators. The first variation is more useful when the KID environment
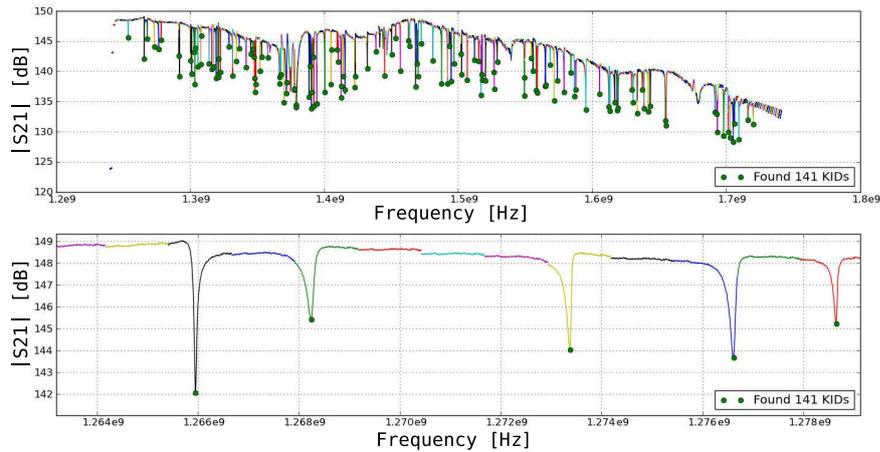
(a) *Finding KID resonant frequencies: a low resolution frequency sweep over 500 MHz is taken by concatenation of IQ data from 150 evenly spaced tones. This is searched through to identify the resonant frequencies. Full view (top) and zoom (bottom).*
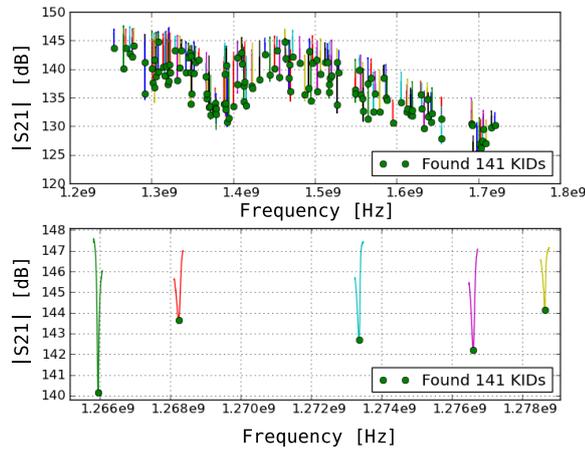
(b) *Refined sweep: a high resolution sweep with tones centered at the frequencies found in the initial full sweep to precisely measure the resonant frequencies. Full view (top) and zoom (bottom).*
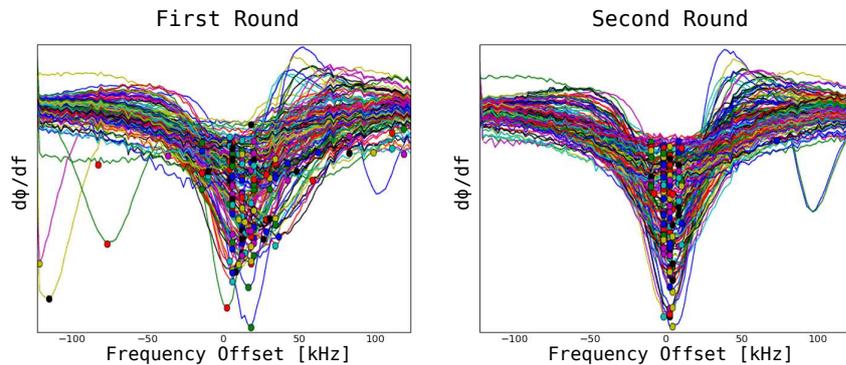
(c) *Calibration sweep: Tones are set to the precise resonant frequencies and a second round of high resolution sweeping provides the IQ data necessary to perform the response calibration. The plots are of the gradient of the phase slope (dφ/df) for each of the 141 found KIDs found in this run.*

Figure 6.2: *Three stages of an automatic tone locking algorithm.*

is regularly changing – as it is during a typical development process. The second variation relies on the KIDs not moving around too much between runs and is more useful in a fixed, stable system.

6.2.3   Tone power optimisation

A KID's performance is strongly dependent on the power of the readout tone used to probe it. Not enough power causes an increase in TLS noise and too much power causes the response becomes strongly non-linear. It is important to make sure the power levels are just right. Furthermore, the KID arrays used in this camera show strong variations in the ideal power levels between individual KIDs. Two methods have been implemented to optimise the power levels across the full array. The first method is a manual method where a large data set is visually inspected and the ideal powers are selected by hand. The second method is based on an iterative process to maximise the signal-to-noise ratio measured from a fixed intensity calibration source situated within the cryostat. Both methods show overall improvements in the detector performance, however, neither method is actually ideal. Further investigations and developments are required to fully optimise the readout tones.

Method 1

Working on the assumption that KIDs perform best when the tone power maximises the detector phase response while ensuring that no strongly non-linear effects are in play, we initially characterise the phase response as a function of readout power. A large dataset is collected and we manually identify the ideal power levels. These values are saved and used for future runs. While it is shown that the phase response is consistently improved, the detector noise levels have not been fully characterised as of yet. As such, it is not confirmed if the signal-to-noise ratios are fully maximised. Figure 6.3 offers a view of the frequency versus power versus phase response dataset. Six KIDs are visible and the chosen power levels are indicated by the white arrows.

Method 2

A $1.5\,\mu m$ LED was installed at the $4\,K$ stage of the cryostat and coupled into the $350\,GHz$ focal plane unit via a multimoded fibre-optic cable. Pulsing the LED produces a large response in some of the detectors and is used to directly measure the detector signal-to-noise ratios (SNRs). A second LED was installed in the same way in an attempt to provide more uniform illumination across the array however some channels remained relatively insensitive to the pulses while others showed very large responses. The outline of the iterative optimisation method is as follows:

1. Lock-on to the KIDs with powers at the lowest level and check SNR with calibration source.
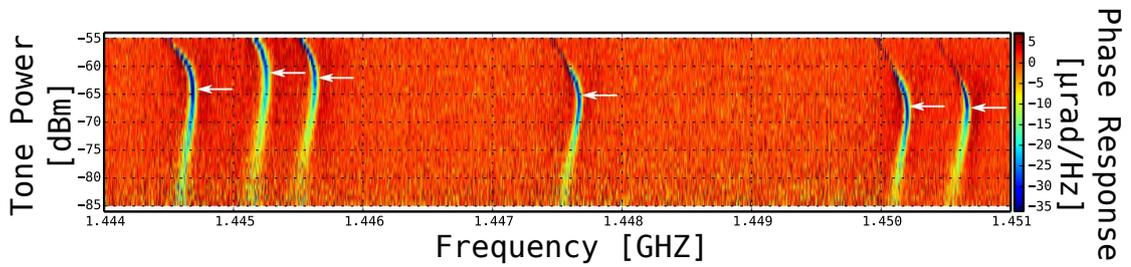
Figure 6.3: *Manual identification of optimal readout power: Full sweeps were taken at a range of readout powers and the optimal frequencies and powers of individual tones were set by clicking on the desired region of the resulting interactive graph as indicated by the white arrows in the figure.*
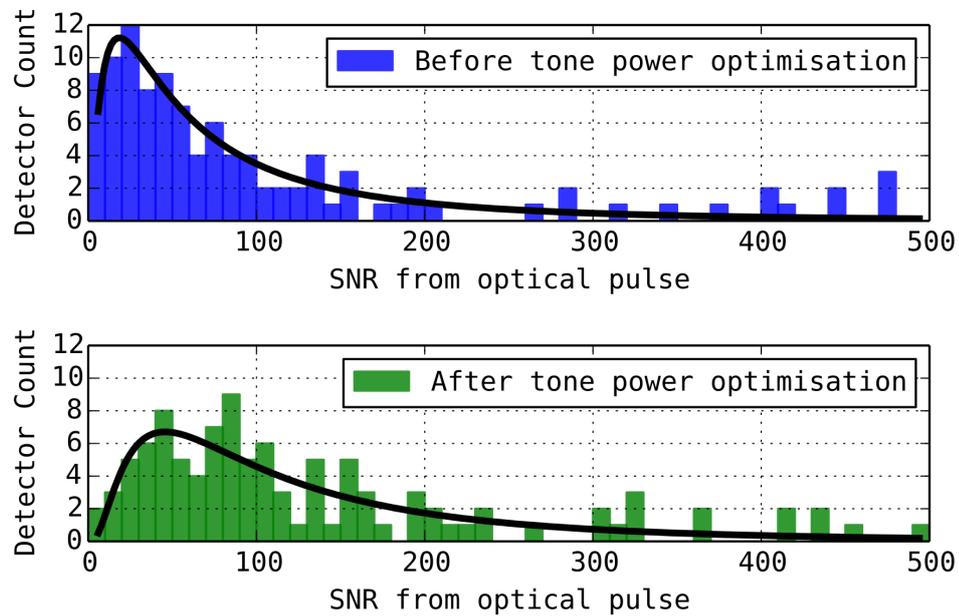
Figure 6.4: *Histograms indicating the effect of the tone power optimisation routine on detector SNR.*

2. Raise each of the tone powers by one unit, lock-on again and measure SNR again.

3. For each KID: if the SNR decreases, drop the power on that tone by one unit.

4. Repeat items 2 and 3 until all tone powers are optimal.

This method is fully automatic and works well to improve the SNR, especially in the most poorly performing detectors. Histograms of the measured SNRs with and without optimised tone powers are given in Figure 6.4. A more uniform array illumination by the calibration LED would, no doubt, provide further improvements.

6.3 Data acquisition system

The data acquisition system is responsible for the continual collection, parsing and storage of data from the readout electronics. The main application is a multi-threaded C program that is started by the control system and continues indefinitely unless it receives a signal to terminate – in which case it gracefully shuts down, freeing any memory and closing any open files and sockets. In the main thread, a UDP socket listens for data packets from the readout electronics and copies the packet contents into a shared FIFO buffer. In a second thread, packets are popped from the FIFO and parsed for information content that is then written to disk in the main storage structure. Here they are free to be accessed by either the control system or the image generation system.

For the purpose of efficiency, the data acquisition system makes use of slow and fast channels within individual data packets. Fast channels update every sample and are always written to the Dirfile. The rate is defined by the amount of averaging applied on the readout electronics and is anywhere up to 477 SPS. Slow channels do not update every sample so consecutive data packets will contain repeated data. They are written to the Dirfile at the slow rate which is defined to always be 1 SPS.

6.3.1 Storage

The main storage structure consists of Dirfiles. These are directories of binary data files (referred to as *raw fields*) with an ASCII format file to specify the binary data formats[2]. For each new run of the camera, the data acquisition systems creates a new Dirfile named by the time of creation in the system data directory, i.e. `/data/<unix-time>.d/` where `<unix-time>` is the number of seconds since midnight, 01/01/1970. A symbolic link to the most recently created Dirfile is written at `/data/etc/dirfile.lnk` for ease of access. Data samples that have been parsed from the packets are written to their own fields (files) within the Dirfile structure. The list of raw fields and their formats is given in Table 6.1.

---

[2]Dirfile Standards: `http://getdata.sourceforge.net/dirfile.html`

Table 6.1:   *The data acquisition system's raw fields, as present in the transmitted and received data packets and as specified in the Dirfile format file.*

Field Name	Type	Format	Rate
i000...399	RAW	INT32	Fast
q000...399	RAW	INT32	Fast
freq000...399	RAW	UINT32	Slow
gain000...399	RAW	UINT32	Slow
dac_gain_0...4	RAW	UINT32	Slow
averaging	RAW	UINT32	Slow
_adc_lev	RAW	UINT32	Fast
_dac_lev	RAW	UINT32	Fast
_flags	RAW	UINT32	Fast
_lo_freq	RAW	UINT64	Fast
_mirror_encoder	RAW	UINT16	Fast
_sample	RAW	UINT32	Fast
_stage_x	RAW	FLOAT32	Fast
_stage_y	RAW	FLOAT32	Fast
_stage_z	RAW	FLOAT32	Fast
_time	RAW	UNIT64	Fast
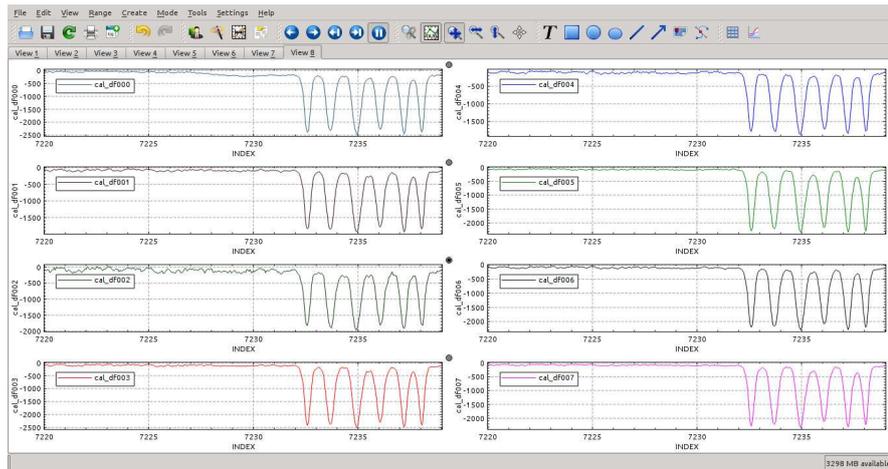
Figure 6.5: *KST plot: Real-time display of calibrated timelines of resonant frequency shift information for 8 channels sampled at 477 Hz with 20 seconds worth of samples per channel. Modulating the field of view by waving a hand in front the camera causes shifts of around 2000 Hz at resonant frequencies of roughly 1.25 GHz.*

6.3.2 Access and display

Dirfile format files are easy enough to read directly but using calls to the GetData library[3] allows us to take advantage of a special class of *derived fields*. These fields can be specified in the format file and enable direct access to mathematically manipulated forms of the raw fields. Magnitudes and phases of complex numbers, multiplications, divisions, linear combinations, bitmasking, timeshifting, multiplexing and a range of other functions are available. These provide us with a means to access an arbitrary number of samples of fully calibrated detector response timelines with incredible easily and minimal CPU and RAM overheads.

Timestream analysis is enhanced by using the real-time plotting software KST[4]. It can simultaneously access and display millions of samples from thousands of files of various formats (including Dirfiles) and has an extensive toolkit for processing, filtering and manipulating timestreams and their spectra. A screenshot show showing the real-time display of detector response from 8 channels is presented in Figure 6.5. With a sufficiently high resolution computer screen it is possible to monitor all of the 150 detector timelines, or all of the 150 power spectra, at the same time.

6.4 Image generation system

The real-time imaging system provides a graphical front-end to the camera operator with a main window to display video frames. The frames are generated and displayed in real-time as the raw data arrives into the system, although the option is available to playback previously recorded videos. The application was designed and written by Dr.

---

[3]The GetData Project: `http://getdata.sourceforge.net/`
[4]KST: `https://kst-plot.kde.org/`

Papageorgiou in Python using a combination of NumPy, SciPy, Matplolib and PIL for image generation, and TkInter for the GUI. In this section I briefly outline the workings of the system.

Before starting the system, some prerequisite information for the production of individual image frames must be available, including:

- A list of KID spatial position in the focal plane, in order to correctly map detector outputs to image pixels.

- The location of detector response and mirror position timelines in the file system, in order to gather data for processing.

- The dimensions of the required image frame and pixel size, measured in image plane coordinates.

- The frame update rate, as set by the mirror oscillation rate.

- The integration time per frame.

- An optional starting frame number and playback speed if reviewing previously recorded frames.

Once the application is started, the GUI loads up and an automatic calibration process begins. Part of this process is the measurement of flat-field correction coefficients for each detector. This is done over the course of a few scanning-mirror cycles by comparing detector responses to a warm calibration bar that is positioned horizontally across the top of the FOV, see Figure 6.6 Any noisy or broken channels are identified and excluded from the remainder of the image generation process.

At this point the program's main loop is started. A diagram representing the key software classes and their methods and attributes is presented in Figure 6.8. An outline of the main loop is given here:

1. For each detector:

    (a) Read the detector and mirror data
    (b) Calculate responses
    (c) Perform timeline filtering
    (d) Perform flat-field calibration
    (e) Bin detector responses into the image array

2. Apply any additional image filters

3. Apply colour and contrast settings

4. Interpolate over missing pixels - see Figure 6.7

5. Update the display with the image

6. Listen for and handle any GUI events

7. Restart main loop
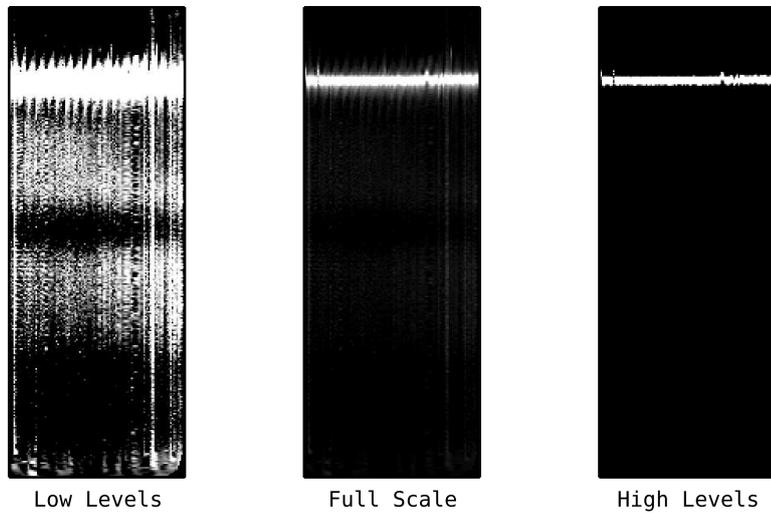
Figure 6.6: *Flat field calibration: A blackened metal bar is heated by an evenly spaced arrangement of electrical resistors to a few degrees above ambient room temperature. Detector outputs are calibrated between this and a fixed region of the background to provide a uniform response across the field of view.*
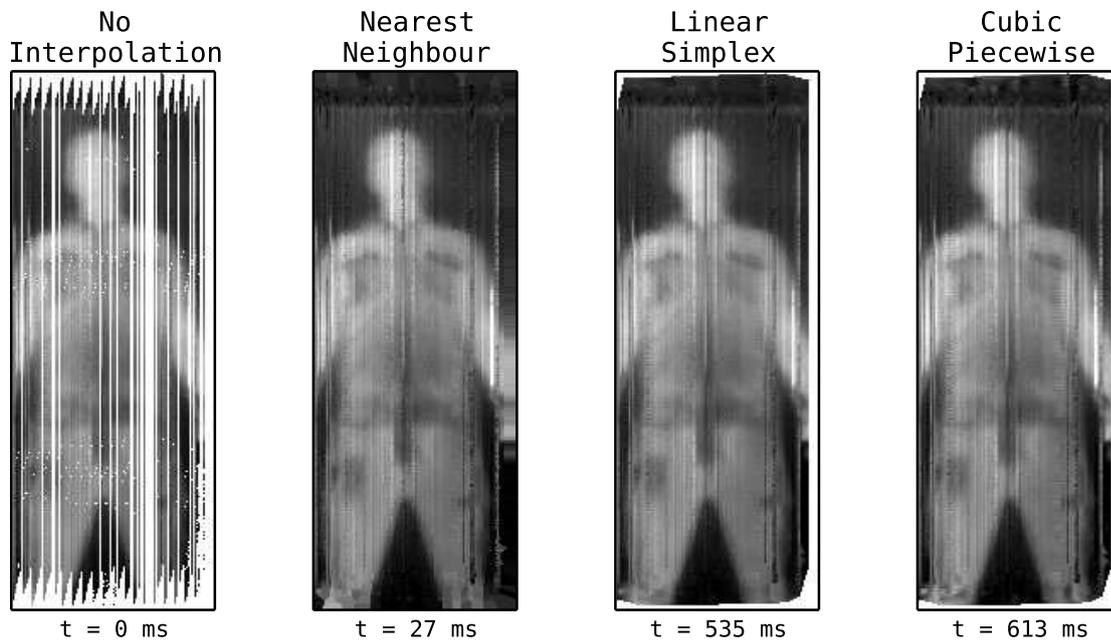
Figure 6.7: *Image interpolation: Three methods of image interpolation are presented against the case with no interpolation. The processing times for each method are included as computed in single core on an Intel i7 2.8 GHz CPU. The outputs of each method are very similar, so the nearest neighbour method is implemented as this is the fastest.*
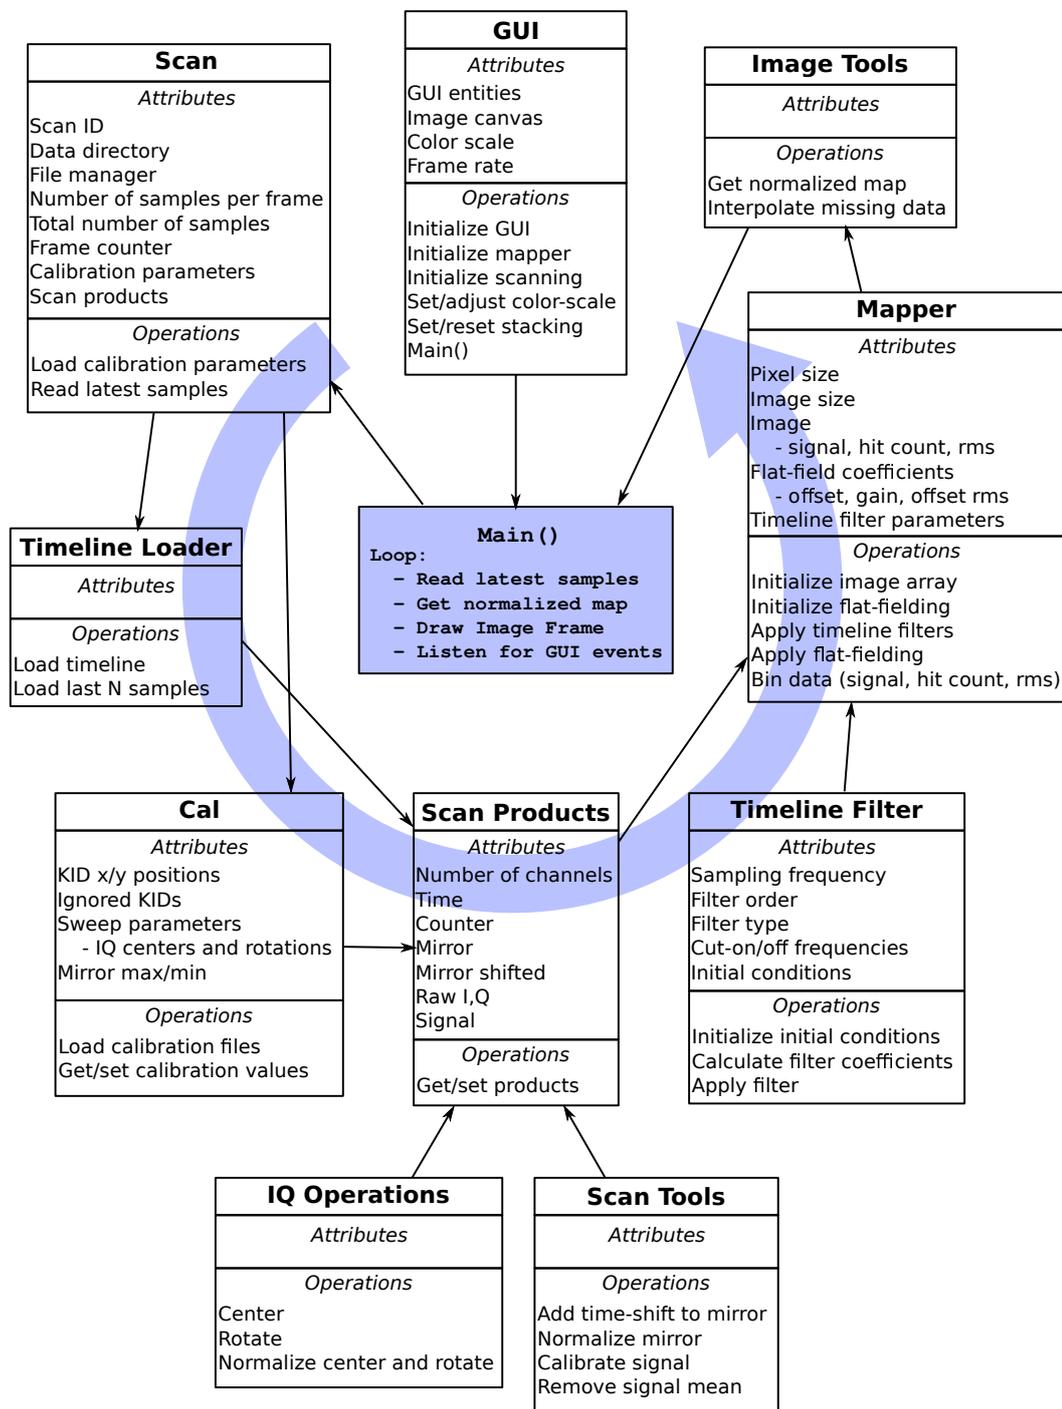
Figure 6.8: *The imaging software broken down by class. For each class the attributes (information that is stored) and operations (functions that can be performed) are described. The direction of data flow through over each iteration of the main loop is indicated by the large arrow arrow. The small arrows show data flow between the classes.*

Image quality

Figure 6.9 presents three video frames taken at 350 GHz over the course of development of the image generation process.

- The left hand frame has image-plane pixel sizes of 10×10 mm, contains data from noisy detector channels and uses an un-optimal flat-field calibration.

- The centre frame has a smaller pixel size of 5×5 mm, excludes most of the noisy pixels and has improved flat-field coefficients.

- In the right hand frame, the field scanning mechanism was adjusted to scan out the full 2 m vertical FOV. While this does not considerable affect the data quality, the image quality is improved greatly as the subject is not longer cut of at the knees. The varying contrasts introduced by the subject's clothing layers add improved aesthetics to the frame.
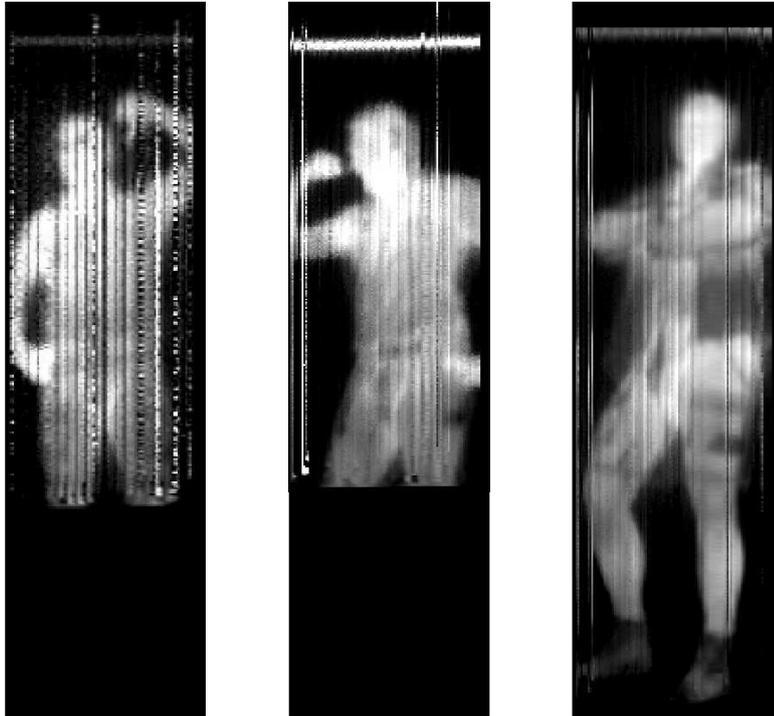
Figure 6.9: *Variations in image quality.*

6.5 Concealed object detection

The following figures demonstrate the detection of concealed items by terahertz imaging. Figure 6.10 contains the first images ever taken by the camera (in the 150 GHz band), and Figure 6.11 has a more recent 350 GHz image. The camera images are presented next to visible and FLIR thermal near-infrared (NIR) images to highlight the unique optical transmission properties of common materials at terahertz frequencies.
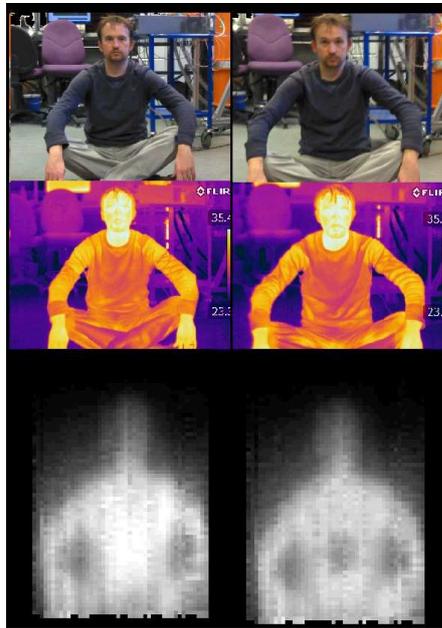
Figure 6.10: *Dr. Doyle posing for the camera in the 150 GHz band with (right) and without (left) a metallic item concealed beneath his clothes. The presence of the object is notable in the terahertz (bottom) but not in the thermal NIR (middle) or visible (top).*
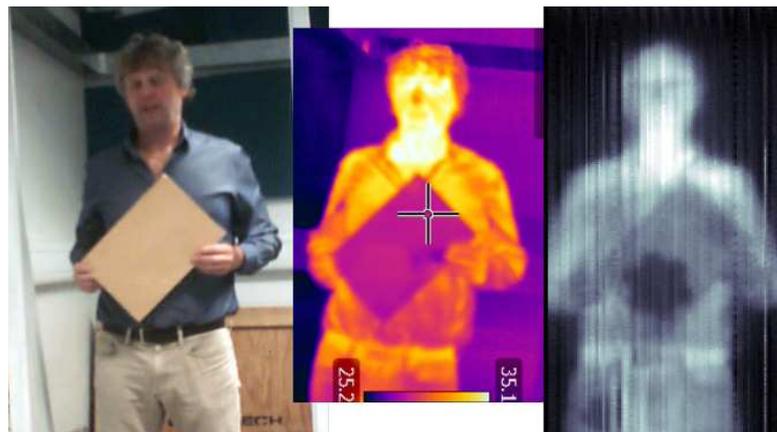
Figure 6.11: *Ken Wood holding a paper envelope in the visible (left), thermal NIR (centre), and terahertz (right). The star shaped band-pass filter in the envelope is clearly visible in the terahertz.*

6.6 Video

Frames from two videos are presented in the following two figures:

- In the first case (Figure 6.12) visible, NIR, and 350GHz videos were recorded simultaneously and stitched together into a single film. Individual frames from one minute of a 2 FPS video are arranged in mosaics separated by optical band for clarity. The frames show a person moving around and removing their jacket. In the first half of this film, a cardboard box in the foreground obscures the field of view. This does not prevent us from seeing the person clearly in the terahertz. Some of the the items concealed beneath the persons clothing are visible, even through the cardboard box.

- The second case (Figure 6.13) is more of a time-lapse than a true video – roughly one frame per minute was saved and recorded into the video file. The purpose of the film was to demonstrate that the heat signatures imparted by living creatures confined in small spaces are detectable by terahertz cameras. Of course a standard thermal imaging system could easily repeat this test, however the terahertz system has the advantage that it could be used to screen containers through the side of soft sided transporters without requiring a direct, visible line of sight. In the film, at t=0, somebody climbs into the box and closes the lid. Within only a couple of minutes the box is seen to warm up. After 30 minutes, the person climbs out and box begins to cool down. The temperature inside the box was measured with a thermometer at regular intervals and is displayed on the individual video frames.

6.7 Summary

The overall architecture of the software system has been presented, including description of the control, data acquisition and image generation systems. Details of various algorithms have been outlined, including the methods used to lock readout tones onto KID resonators and optimise the tone powers. The image processing application has been described and video frames have been presented with various factors affecting the image quality having been described. Analysis of image quality can be somewhat subjective and in the next chapter I describe the camera performance quantitatively in terms of the detector and system sensitivities.
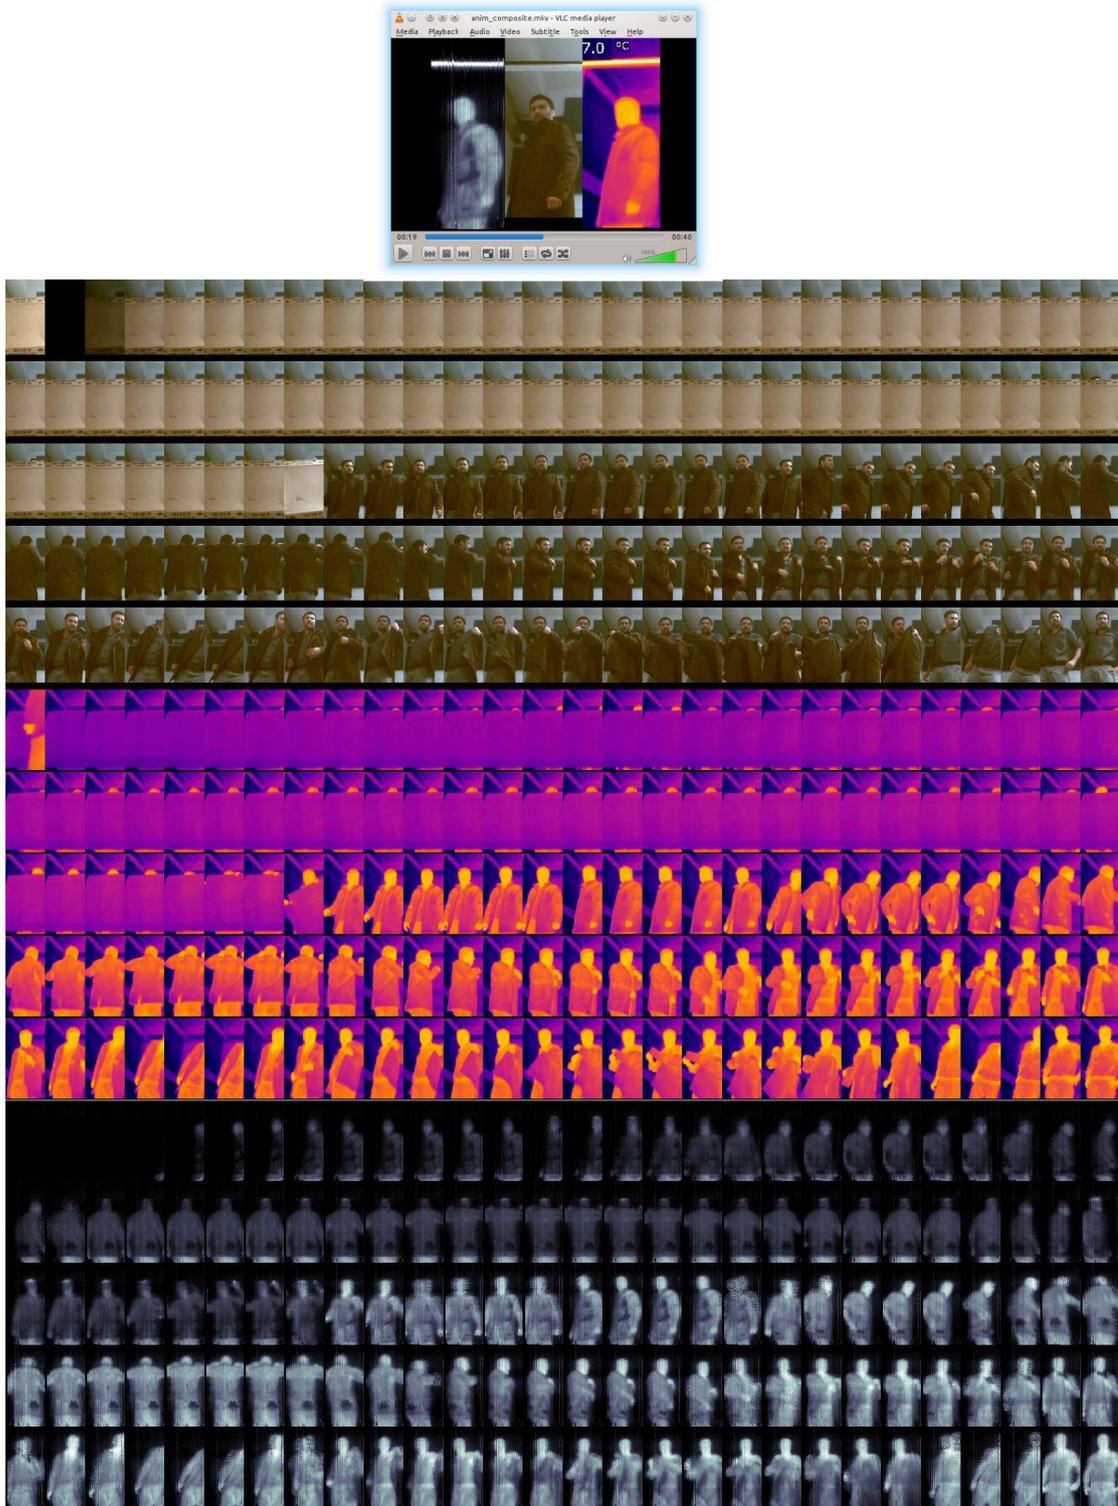
Figure 6.12:   *Individual frames from a 2 FPS video where visible, NIR and terahertz images were played side-by-side. In the first half of the video, the FOV is blocked by a cardboard box. This does not prevent us from clearly seeing the person moving in the background in the terahertz, or indeed, from seeing the items concealed beneath his clothing layers..*
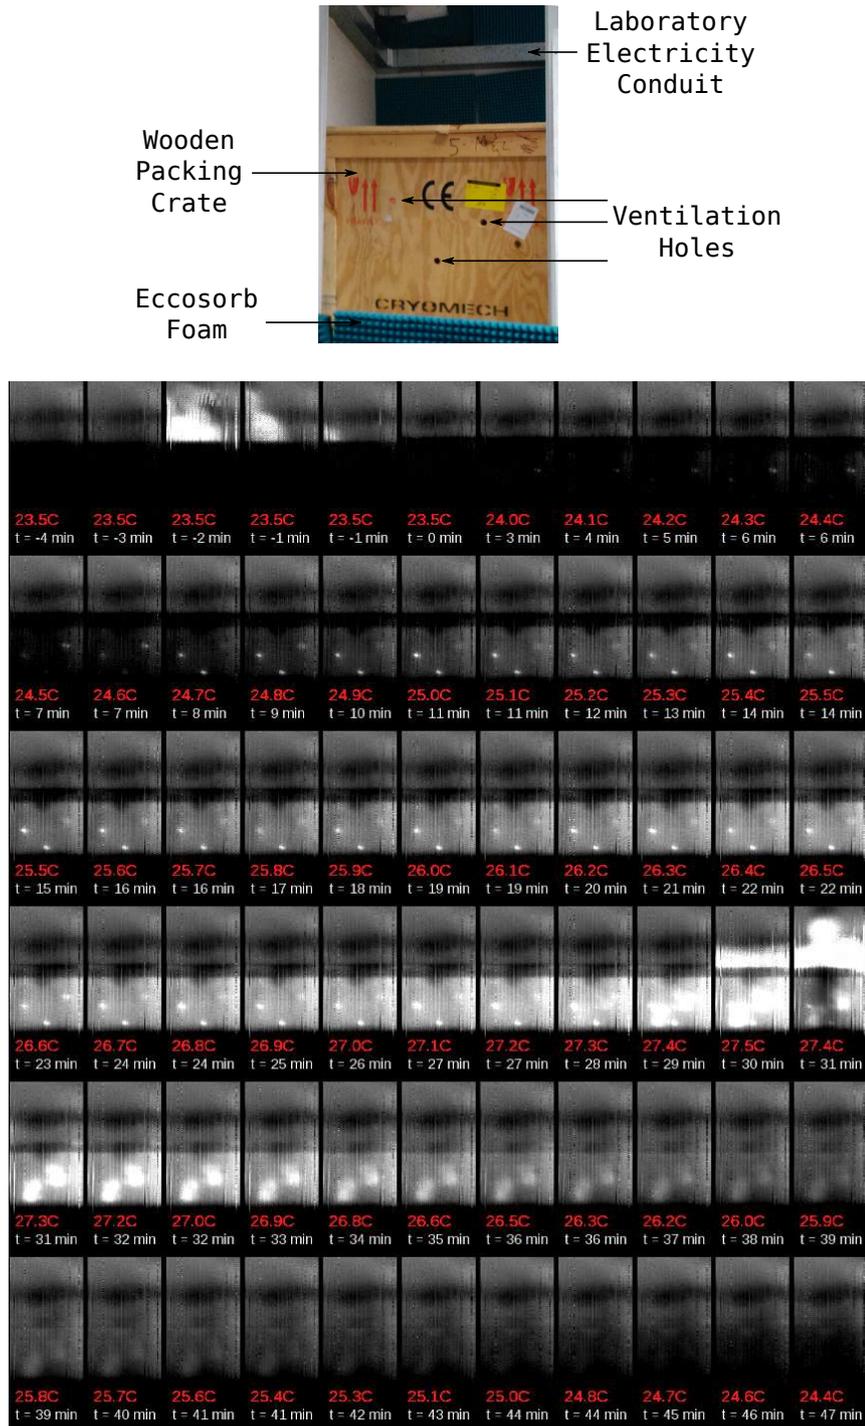
Figure 6.13: *Frames from a time lapse video taken one minute apart for nearly an hour. A wooden packing crate is seen to warm up immediately after somebody climbs in and closes the lid. The signal to noise level is sufficient to enable the rapid detection of concealed animals or even people. Though this is indeed possible with a regular infrared camera, the advantage of a THz system is that the heat source will remain visible even through multiple layers of infrared-opaque material such as further wooden panelling or even the soft-sided exteriors of heavy goods vehicles.*

Chapter 7

Camera Performance and Future Considerations

The goal of this project was to develop a video rate passive imaging system to demonstrate LEKID focal plane array technology as a viable option for applications requiring ultra-high sensitivity at terahertz frequencies. The camera was designed to operate in the same vein as airport body scanners or other stand-off scanning systems utilised in the security industry. Accordingly, a successful system would provide a means for the real-time identification of concealed items on a moving human target.

The following chapter concludes this thesis with an analysis of the overall performance of the terahertz camera in its most recent configuration. A number of factors can affect the ultimate sensitivity, as have been described in previous chapters of this thesis. Expectations for the various limiting sensitivities are recapped and compared with the measured noise properties. The limitations that currently prevent the system from achieving full video rate imaging are discussed, a comparison is given against other commercially available terahertz imaging systems, and finally, important considerations for a next generation system are presented.

7.1 Sensitivity

7.1.1 The ideal system sensitivity

An ideal passive terahertz camera would not introduce any significant noise into the images it generates. The limiting noise source ought to be that of the random fluctuations in the photon distribution due to the variations in light source itself and not the camera. A discussion of photon noise and the associated sensitivity limits is give in appendix B.2. A target noise equivalent differential temperature (NEDT) was put forward of $\sim 0.1\,\mathrm{K}$ NEDT per frame as this is close to the expected background noise limit due to random temperature fluctuations of a human body if it were sampled at a $25\,\mathrm{Hz}$ frame rate with a $1\,\mathrm{cm}$ spatial resolution at a focal distance of 3-5 m, a $1{\times}2\,\mathrm{m}$ field of view, a $25\,\mathrm{Hz}$

refresh rate, and with operation in a 10% wide optical band centred on the 350 GHz atmospheric window. The photon noise level is estimated here given these parameters and the goal would be to ensure that this remains the limiting factor, with detector noise and electronic noise not becoming dominant.

A 0.5 m primary lens with a focal length of 4.5 m provides a diffraction limited resolution of 1.04 cm and a single-moded throughput $A\Omega = 7.6 \times 10^{-7}\,\mathrm{m^2\,sr}$. Assuming, initially, no losses in lens transmission, an ideal top-hat optical filter profile with $\Delta\nu = 35\,\mathrm{GHz}$, and a detector sensitive to both polarisations, the expected power received from a 300 K blackbody source is $\sim$290 pW. The total photon noise NEP is estimated from equation B.16 to be $1.65 \times 10^{-15}\,\mathrm{W}/\sqrt{\mathrm{Hz}}$, which is dominated by wave noise with $\frac{NEP_{\mathrm{shot}}}{NEP_{\mathrm{wave}}} = 0.23$. This is equivalent to an noise equivalent temperature, NET, of $1.26\,\mathrm{mK}/\sqrt{\mathrm{Hz}}$ or $1.13\,\mathrm{mK}\sqrt{\mathrm{s}}$ assuming the Rayleigh-Jeans approximation for the responsivity of optical power to temperature (see appendix B.1.

A fully Nyquist sampled image requires $4 \times \frac{1 \times 2}{0.0104^2} = 74000$ pixels, and the requirement for a 25 Hz frame rate means the integration time per image pixel is 0.5 µs if a single detector is used. The estimated NEDT per frame using Equation B.7 is then 1.55 K for a photon noise limited detector. The detector count for this instrument was set at 150 as this would provide an NEDT of 0.127 K per frame at 25 Hz, and the expected integration time per detector per pixel would then be 107 µs. This NEDT is close to the desired 0.1 K and the modest detector count allows for a broad $\sim$3 MHz spacing between resonators which minimises the likelihood of resonator clash within the 500 MHz of readout bandwidth.

7.1.2 The overall system sensitivity estimate

In practice the system is not perfect and factors such as the camera optical efficiency, the detector absorption efficiency and stray light need to be accounted for. The optical efficiency was estimated in chapter 3 to be $\eta_{opt} = 22\%$. An estimate of the detector efficiency can be made by combining estimates for the detector absorption efficiency and the horn efficiency. The HFSS simulation of the 350 GHz detector architecture in chapter 4 indicates a 35% peak in-band efficiency to incoming plane waves, and a mid-range estimate of the feed-horn efficiency is 70%, which combine to give a total detector efficiency $\eta_{det} = 25\%$. The total power measured by the detector is then expected to be 31 pW (17 pW from the source and 14 pW stray) so the photon noise limiting sensitivity is $NEP_\gamma = 2.88 \times 10^{-16}\,\mathrm{W}/\sqrt{\mathrm{Hz}}$. This expected power is low enough that the shot noise component now marginally dominates over the wave noise, with $\frac{NEP_{\mathrm{shot}}}{NEP_{\mathrm{wave}}} = 1.5$.

Estimates for the remaining NEPs are presented in Table 7.1. The values are calculated according to the formulae presented in this thesis and, where applicable, are based on typical system properties (detector Q-factors, β-factors etc.) that were measured during operation of the camera. Optical illumination is not identical across the whole focal plane and individual detectors do not perform identically, so the parameters

Table 7.1: *The total expected system sensitivity in NEP and NET. The results are computed from the average of a number of typical detectors. Photon noise is clearly the limiting factor.*

Sensitivity	γ	G-R	Amplifier	Readout	**Total**
NEP [W/$\sqrt{}$Hz]	2.9×10^{-16}	6.5×10^{-17}	8.5×10^{-18}	2.8×10^{-18}	$\mathbf{3.0\times10^{-16}}$
NET [mK$\sqrt{}$s]	3.66	0.82	0.11	0.035	**3.75**

vary by about a factor of two-to-three from KID to KID. Nevertheless, it is clear that it is photon noise that remains the most significant noise contribution.

Using a Rayleigh-Jeans responsivity of $\left(\frac{dP}{dT}\right)_{\text{RJ}} = 5.64 \times 10^{-14}\,\text{W/K}$, the estimated overall NET per detector is $3.75\,\text{mK}\sqrt{}\text{s}$. Given that the measured extent of the field of view in the object plane is 0.6×2.15 (thus the number of pixels per image frame is 41100) the expected NEDT per frame is $0.09\,\text{K}$ for a frame rate of $2\,\text{Hz}$ (the integration time per pixel is $1.1\,\text{ms}$). So the target NEDT has been met but unfortunately this is at the expense of the target frame rate.

7.1.3 The measured sensitivity

Sensitivities as function of frequency were found for each detector by dividing the power spectral densities of the detector frequency response timelines by the detector responsivities. The responsivities were found by recording the detector response between the hot calibration bar and a section of the background and then comparing against the independently measured temperature difference. The power spectral densities were computed from FFTs of 60 seconds worth of samples that were taken while staring at a fixed source of constant temperature. The 152 NET spectra are presented in the upper part Figure 7.1.

The majority of the channels show similar performance with white noise down to a few Hertz. The sensitivity at the white noise level was sampled for each detector at approximately $100\,\text{Hz}$ and the values are displayed in the histogram in the lower part of the figure. The distribution is approximated by a log-normal fit and the modal average of the sensitivities was found to be $6.1\,\text{mK}\sqrt{}\text{s}$.

This value is roughly a factor of two higher than the sensitivity we expect for the camera in its present configuration and the likelihood is that there is leak of high frequency (infrared) radiation through one of the band defining filters. Only a few infrared photons from the $\sim$150 K cold lens are necessary to add a significant optical load to the detectors, thus increasing the quasiparticle density (as well as the photon shot noise).

With the detectors twice less sensitive than estimated, we expect the images to be twice less sensitive as well. The new estimate is simply $2\times0.09 = 0.18\,\text{K}$ per frame at 2 FPS. In fact, an NEDT of $0.13\,\text{K}$ was measured in the uniform region of the temperature map given in Figure 7.2. The sensitivity map to the right of the temperature map shows
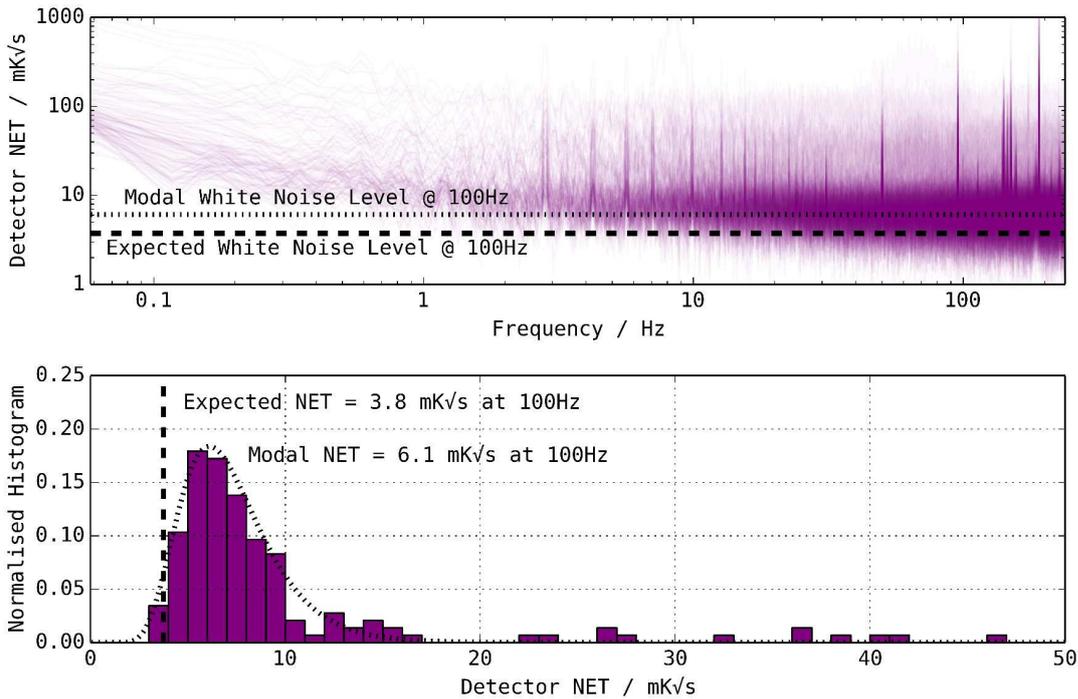
Figure 7.1: *Measured detector sensitivity. (Top) Plots of NET versus frequency for the 152 detectors over 60 seconds of sampling. (Bottom) The histogram of the detector NETs measured at the white noise levels. A log-normal fit indicates a modal value for the NET of 6.1 mK$\sqrt{s}$, which is worse – by around a factor of two – than the expected sensitivity.*

the NEDT per pixel by averaging the per pixel variance over a number of frames and, although it is not directly representative of the single frame NEDT, it does indicate the relative sensitivity as a function of spatial position in the image plane.

7.2   Factors limiting the performance

The current scanning mechanism ultimately limits the frame rate of the system and vertical stripes in the images are one of the worst factors affecting the overall image quality. These are an artefact of the uni-axial scan which become significant when individual detectors malfunction. The 85% array yield is good, however a system that does not cross link between channels will suffer. A redesign of the scanning mechanism to incorporate a fast 2-axis scanner would remedy this.

The image generation algorithm is presently non-optimal as data samples are mapped into single points in the images. In reality, point sources are should appear spread out over a number image pixels. An algorithm utilising full convolution of detector point spread functions should improve image quality.

The camera sensitivity is around 0.13 K per frame at a rate of 2 FPS, however the goal was for full video rate (∼25 FPS) imaging at 0.1 K per frame. The integration time
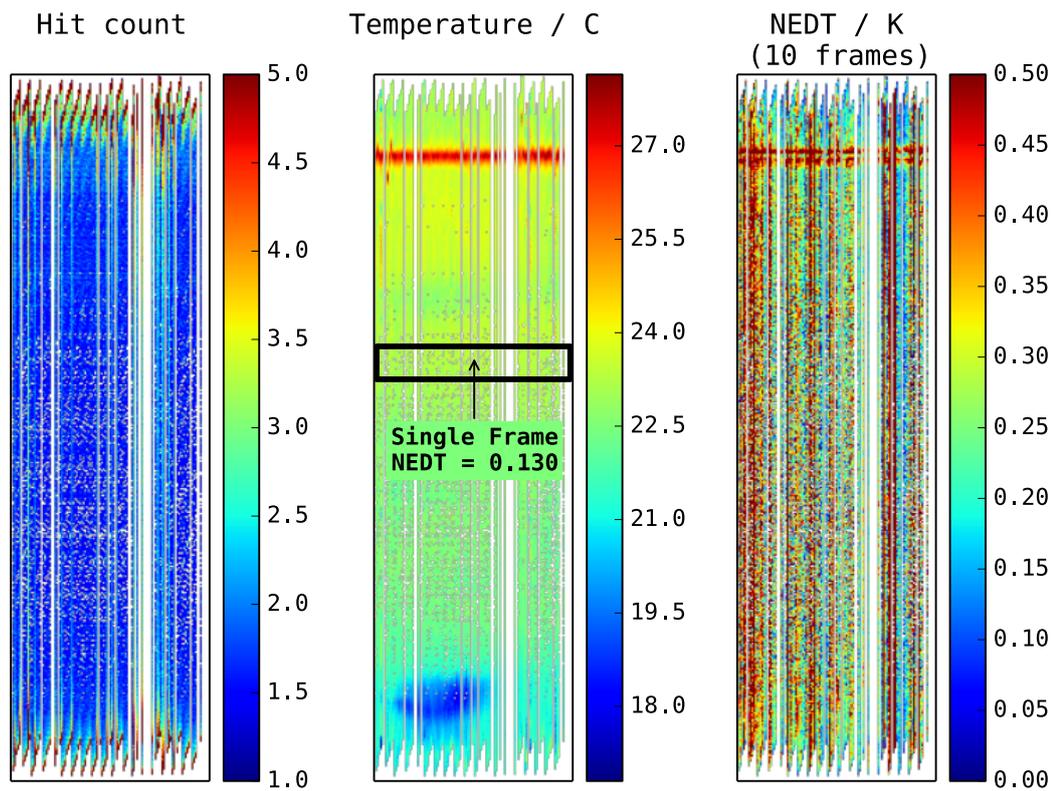
Figure 7.2: *Image sensitivity: Sample count per image pixel (left), a temperature map (centre) and a noise map (right). The NEDT per single frame is calculated from the variance in the temperature map. The highlighted region has a fairly flat background so the NEDT is accurately measured. The per pixel NEDT displayed on the right required sampling over 10 frames in order to gather the required statistics – as such, low frequency noise (see Figure 7.1) is strong and the sensitivity is poorer than for a single frame.*

per image pixel at full video rate will be 12.5 times lower than at present, so a factor of $\sqrt{12.5} = 3.5$ increase in sensitivity is required. There are a number of ways this could be done but each requires major adjustments to the camera.

Stray light is single most significant factor limiting the overall sensitivity. Isolating and fixing the stray infrared leak would provide an immediate factor of two improvement before the estimated photon noise limit is reached. Reduction of the remaining stray light would help to improve the system responsivity, however, this would require a total redesign of the optical system in order to move away from lossy refractive components.

With the current optical configuration, simply operating with 12.5 times more sensors in the focal plane (1850 KIDs) would enable a 25 Hz refresh rate. This would require some effort in array design and fabrication, a new focal plane unit with horn block, new readout electronics (or just four more NIKEL units), and some software updates, but this is totally feasible, in principle. Multi-kilopixel array development is already under way in astronomical cameras such as NIKA-2, and the relatively high optical powers involved with terrestrial imaging place relatively less stringent sensitivity requirements on the individual detectors.

7.3 Performance comparison with other systems

The following table (Table 7.2) presents the key specifications and measured performance of this KID camera next to the specifications of similar terahertz imaging systems. Commercially available systems tend not to use cooled detectors so they can be very compact, however, they lack the high sensitivity of cooled systems. Bear in mind that the system presented in this thesis is only a prototype technology demonstrator and is not currently suited for market deployment.

7.4 Thoughts for future developments

Optics

- LEKIDs are direct detectors and therefore are sensitive to radiation from all directions. As such, slow optics, effective baffles, and cold aperture stops should be in place. Feed horns are shown to be effective at improving beam shape and improving optical coupling.

- Reflective optics can reduce stray light contributions relative to refractive optics. The photon noise level should not be increased by unwanted radiation from emissive components in the optical chain.

- Cross linking between detectors over the course of an imaging scan with a two-axis strategy such as a circular, spiral, or Lissajous pattern is vital to remove dead space from broken/poorly functioning detectors.

Producer	Cardiff AIG	TeraSense	INO	NEC	alphanov	Digital Barriers	IPHT Jena	NIST/VTT
Product	KID cam	TeraFAST-256-HS	IRXCAM-384THZ	IR/V-T0831	TERACAM	Thruvision TS5	–	–
Commercial	no	yes	yes	yes	yes	yes	no	no
Frequency	350 ± 10 GHz	50-700 GHz	94-4000 GHz	1-7 THz	0.1-30 THz	250 ± 20 GHz	350 GHz	350 GHz
Sensors	152	256	384×288	320×240	40×40	?	?	251
Type	LEKID	GaAs	μbolo FPA	μbolo FPA	?	?	TES	TES
Operating T.	0.25 K	300 K	300 K	300 K	300 K	?	0.5 K	<0.5 K
NEP	6×10^{-16} W/$\sqrt{\text{Hz}}$	1×10^{-9} W/$\sqrt{\text{Hz}}$	?	1×10^{-10} W	1×10^{-6} W	?	BLIP	?
NEDT per frame	0.1	?	?	?	?	?	BLIP	?
Frame rate	2 Hz	5000 Hz	48 Hz	30 Hz	?	6 Hz	< 25 Hz	?
Field of view	1×2 m	n/a	F=44mm,f/0.9	$15.0° \times 11.2°$	?	1×1 m at 10 m	< 2 m	?
Range	3.5 m	n/a	?	?	?	6-15 m	3-10 m	16-28 m
Resolution	1 cm @ 3.5 m	n/a	?	?	?	?	1-2 cm	1:4 cm
Dimensions [cm]	$200 \times 100 \times 100$	$30 \times 10 \times 10$	$6.5 \times 5.9 \times 10.5$	$6 \times 6 \times 20$	$6 \times 6 \times 30$	$95 \times 68 \times 31$	$<100 \times 100 \times 200$	$<100 \times 100 \times 200$

Table 7.2: A performance comparison of the KID camera with other similar passive terahertz systems.

- A modular optics approach is sensible. The ability to swap out various optics systems with a variety of focal lengths, depths of fields, resolutions, etc. without requiring any modification of the cryogenic system is desirable. This could be acheived by coupling the detector array to an image plane located externally to the cryostat, perhaps with a small lens in place of a large window.

Detectors

- Improved optical coupling efficiency is not necessarily beneficial when operating in the wave noise limit, however the reduction of the stray light reflections from poorly absorbing detectors is desirable. Good matching of wafer thickness to optical wavelength and the inclusion of anti-reflection coatings are advisable.

- Detectors with higher quality factors are more responsive however the power handling and dynamic range need to be well matched to the expected incident optical power levels.

- Reduction in the overlap of resonators will improve array yields. Moving to lower resonant frequencies, and ensuring more uniform film thickness across the array should inhibit the deviations in resonant frequency and quality factor from the designed values.

- Operation at higher temperatures is desirable. Aluminium KIDs work best at around a few hundred mK and so require sub-Kelvin cooling platforms such as sorption fridges, ADRs or dilutors. Other metals have higher T_c and so it may be possible to operate off of liquid helium or PTCs only.

- KIDs are broadband detectors and, as such, the potential exists for multi-colour imaging. Low resolution imaging spectroscopy at terahertz frequencies is desirable and could be achieved with the selective positioning of bandpass filters over a detector array. Mid resolution ($\mathcal{R} < 1000$) spectroscopy has been achieved with KIDs and designs for imaging spectrometers are under development. However, the number of detecting elements rises very steeply as the spectral and spatial resolutions increase.

Electronics

- Greater ADC/DAC bandwidth and resolution allows for the simultaneous readout of more channels. The accelerating growth in the microwave electronics industry means cutting edge, high speed devices will continue to be affordable in years to come.

- Full video rate imaging requires high data sample rates, so good quality FPGAs with high F_{max} and plentiful DSP blocks are needed for the rapid demultiplexing and filtering of readout signals.

- Good independent control of tone powers is necessary to drive individual KIDs at their point of maximum response.

- The FPGA firmware for KID readout systems should be multi functional - not only should it allow for the real-time detector readout for imaging mode, but it should also provide a means for detector characterisation and other debugging and diagnostic features. Tools for network analysis (frequency sweeping) and high speed sampling (for quasiparticle lifetime measurements) are definitely worth incorporating into firmware designs.

Software

- Software should not affect the performance of any imaging system, however, at high frame rates or with large format images, real-time imaging code can lag. Optimisation of algorithms is important and the use of GPU systems for rapid processing (for example, timeline filtering, real-time PSF convolutions, image filtering) is advisable.

7.5   Conclusion

Kinetic inductance detectors are highly sensitive, broadband, electromagnetic radiation detectors designed specifically to overcome the challenges of large format array fabrication and electronic readout at terahertz frequencies. Lumped element KIDs are a variant originally proposed and developed within the Astronomy Instrumentation Group at Cardiff University. International efforts on KID/LEKID development have focussed primarily on experimental astrophysics research with excellent results to date.

We have designed, built and characterised a prototype passive terahertz video camera based primarily on a 152 pixel array of aluminium LEKIDs with the aim of demonstrating that KID technology is suitable for use in a range of terrestrial applications where ultra-high sensitivity imaging is desired.

The goal of generating images of a sufficient quality to identify objects concealed beneath layers of clothing on human targets has been achieved and the camera successfully operates in real-time at a few frames per second with close to background limited sensitivity. The frame rate limit is currently set by the field scanning mechanism and the sensitivity is limited by excess noise that is probably due to stray infrared radiation leaking into the focal plane.

Pixel counts of KID based imaging instruments are increasing (see Figure 1.5) faster than any other technology in the terahertz. Order of magnitude increases in mapping speeds at background limited sensitivities in a range of astronomical and terrestrial settings are expected within the next decade.

Appendix A

Derivations

A.1 Estimation of resonant frequency shift with fixed tone readout

KIDs respond to changes in incident optical power with changes in complex electrical surface impedance. The affected measurable quantities, resonant frequency, f_0, and quality factor, Q_0, are found by sampling the forward transmission, S_{21}, of a readout tone as a function of frequency across the KID resonator and then fitting to the resulting curve. In practice, it is much more convenient to operate with a fixed frequency readout tone and to estimate changes in f_0 and Q_0 with some linear combination of the sampled $I(t)$ an $Q(t)$ values. Measuring $S_{21}(f)$ around the resonator just once at the beginning of a run can provide enough information for a reasonable first order estimation of the response as a function of time.

The forward transmission scattering parameter has a complex value

$$S_{21} = I + jQ \,, \tag{A.1}$$

where I and Q are the in-phase and quadrature components of the microwave readout signal. To first order, a small change in resonant frequency from the initial value of f_0 can be estimated from the change in S_{21} and the gradient of the S_{21} frequency sweep at f_0 as

$$\delta f = \frac{\delta S_{21}}{\left.\frac{dS_{21}}{df}\right|_{f_0}} \,. \tag{A.2}$$

Expanding this out in terms of I and Q gives

$$\delta S_{21} = \delta I + j\delta Q \,, \tag{A.3}$$

and

$$\left.\frac{dS_{21}}{df}\right|_{f_0} = \left.\frac{dI}{df}\right|_{f_0} + j\left.\frac{dQ}{df}\right|_{f_0} \,. \tag{A.4}$$

Then the small change in resonant frequency is

$$\delta f = \frac{\delta I + j\delta Q}{\left.\frac{dI}{df}\right|_{f_0} + j\left.\frac{dQ}{df}\right|_{f_0}} \ . \tag{A.5}$$

The complex denominator can be rationalised out, leaving

$$\delta f = \frac{\left(\delta I + j\delta Q\right)\left(\left.\frac{dI}{df}\right|_{f_0} - j\left.\frac{dQ}{df}\right|_{f_0}\right)}{\left.\frac{dI}{df}\right|_{f_0}^2 + \left.\frac{dQ}{df}\right|_{f_0}^2} \ , \tag{A.6}$$

which can be expressed in real and imaginary parts as

$$\delta f = \frac{\delta I \left.\frac{dI}{df}\right|_{f_0} + \delta Q \left.\frac{dQ}{df}\right|_{f_0}}{\left.\frac{dI}{df}\right|_{f_0}^2 + \left.\frac{dQ}{df}\right|_{f_0}^2} + j\left(\frac{\delta Q \left.\frac{dI}{df}\right|_{f_0} - \delta I \left.\frac{dQ}{df}\right|_{f_0}}{\left.\frac{dI}{df}\right|_{f_0}^2 + \left.\frac{dQ}{df}\right|_{f_0}^2}\right) \ . \tag{A.7}$$

We are only interested in the real part of the frequency shift so we can discard the imaginary part. Finally, δI and δQ are simply $I(t) - I_0$ and $Q(t) - Q_0$, where I_0 and Q_0 are the values of I and Q with the readout tone fixed at f_0. The response as a function of time can be then estimated as

$$\delta f(t) = \frac{\left(I(t) - I_0\right) \left.\frac{dI}{df}\right|_{f_0} + \left(Q(t) - Q_0\right) \left.\frac{dQ}{df}\right|_{f_0}}{\left.\frac{dI}{df}\right|_{f_0}^2 + \left.\frac{dQ}{df}\right|_{f_0}^2} \ . \tag{A.8}$$

Appendix B

General sensitivity and photon noise calculations

B.1 Sensitivity

The low signal levels present ($\sim$100 pW per detector), along with the need to recognise tiny variations in temperature and emissivity (less than a few percent), on top of the requirements for a large field of view ($\sim$500 sq. deg.) and high frame rate (25 FPS), make sensitivity the key factor in the overall performance of this terahertz camera.

A standard measure of detector sensitivity is the noise equivalent power (NEP) which is defined as the optical signal strength required to equal the noise in a 1 Hz post detection bandwidth, corresponding to a half second of integration. NEP is usually quoted in units of Watts per root Hertz (W/$\sqrt{\text{Hz}}$) and, for a time domain signal $x(t)$ sampled at a frequency f_s, is computed as a function of frequency f by taking the square root of the single sided power spectral density $S_x(f)$ and dividing by the detector responsivity R_x. Responsivity is the change in detector response, dx, for a given change in absorbed optical signal power, dP and may be written explicitly as dx/dP, so that

$$NEP(f) = \frac{\sqrt{S_x(f)}}{R_x} = \sqrt{S_x(f)} \left(\frac{dx}{dP} \right)^{-1} \quad \left[\text{W Hz}^{-\frac{1}{2}} \right] . \tag{B.1}$$

In certain situations it is convenient to replace the power spectral density with the signal variance σ_x^2 although one must remember to divide by the post-detection bandwidth Δf_s to maintain the requirement for the half second of integration. For a Nyquist sampled signal, $\Delta f_s = f_s/2$.

A number of different factors will contribute to the overall sensitivity of a detector. For LEKID based systems, possible noise sources include the cryogenic amplifier, the readout electronics, and the detectors themselves – including quasiparticle generation-recombination (G-R) noise and two level system (TLS) noise. An ideal detector would be sensitive enough that the dominant source of noise comes from the inherent variance of the incident optical signal. Such a detector is said to be photon (γ) noise limited. An

NEP can be computed for each of the sources of noise and the overall sensitivity is then the quadrature sum of the individual NEPs, so for a LEKID based system such as ours,

$$NEP_{\text{total}}^2 = NEP_\gamma^2 + NEP_{\text{GR}}^2 + NEP_{\text{TLS}}^2 + NEP_{\text{amplifier}}^2 + NEP_{\text{readout}}^2 \, . \tag{B.2}$$

It is common to quote detector sensitivity in terms of noise equivalent temperature (NET). This is the minimum equivalent signal temperature required to equal the noise in a 1 Hz bandwidth and is related to the NEP as

$$NEP = NET \frac{dP}{dT} \, , \tag{B.3}$$

where the derivative relates the radiated power P to the source temperature T. At the energy scales within which we operate, an approximation to dP/dT can be made based on blackbody sources emitting in the Rayleigh-Jeans limit, where the blackbody spectral radiance $B_\nu(T)$ varies linearly with temperature, $dB_\nu/dT = 2\nu^2 k_B/c^2$.

For convenience, NET is often quoted in units of Kelvin root seconds (K$\sqrt{s}$) which is the equivalent signal temperature required to equal the noise in a *one second* integration time. Note that conversion to this form involves a division by a factor of $\sqrt{2}$ to account for the extra half second of integration. That is,

$$NET \left[\text{K} \sqrt{\text{s}} \right] = \frac{1}{\sqrt{2}} \, NET \left[\frac{\text{K}}{\sqrt{\text{Hz}}} \right] \, . \tag{B.4}$$

With imaging systems, it is common to quote the sensitivity per image frame. Noise equivalent differential temperature (NEDT, NEΔT, or NETD) is a measure of the temperature difference required to equal the noise in a single frame and is quoted in units of temperature (K), where the integration time is accounted for in the frame rate. For a camera composed of a detector array of N_{det} elements that operates by simply staring at an object, the per-frame NEDT is the average of the per-pixel NEDTs which, in turn, is the average of the detector NETs in K$\sqrt{s}$ divided by the square root of the integration time τ per frame,

$$NEDT_{\text{staring}} = \frac{1}{\sqrt{\tau}} \frac{1}{N_{det}} \sum_{i=1}^{N_{det}} NET_i \quad [\text{K}] \, . \tag{B.5}$$

For a scanning camera such as that presented in this thesis, each detector contributes some signal to a number of image pixels so estimating an NEDT becomes slightly more complicated. A system with N_{det} detectors that scans a scene to generate an image with N_{pix} pixels will have an NEDT per pixel that depends the arrangement of the detectors within the array and the scanning trajectory. It becomes necessary to calculate which detectors have contributed signal to which pixel and how long they have integrated over

that pixel. If the i^{th} detector integrates on the j^{th} pixel for a time $\tau_{j,i}$, then

$$NEDT_{\text{scanning}} = \frac{1}{N_{pix}} \sum_{j=1}^{N_{pix}} NEDT_j \tag{B.6a}$$

$$= \frac{1}{N_{pix}} \sum_{j=1}^{N_{pix}} \sum_{i=1}^{N_{det}} \frac{1}{\sqrt{\tau_{j,i}}} NET_i \quad [\text{K}] . \tag{B.6b}$$

Estimation of the NEDT for a scanning system can be simplified by applying a simple scaling factor to the staring case given by Equation B.5. In general, the factor is the square root of the ratio of image pixel count to detector count and gives

$$NEDT_{\text{estimated}} = \frac{1}{\sqrt{\tau}} \frac{1}{N_{det}} \sqrt{\frac{N_{pix}}{N_{det}}} \sum_{i=1}^{N_{det}} NET_i \quad [\text{K}] , \tag{B.7}$$

where τ is now simply the frame duration.

B.2 Photon noise

Photon noise is a variation in detected optical power that comes from variations inherent in the photon flux itself. Leclercq in [Leclercq, 2007] brings together a number of historical approaches to photon NEP estimation into a unified framework and Zmuidzinas [Zmuidzinas, 2003] provides a comprehensive review of the definitions for photon noise NEP given a a variety of optical configurations. An estimate for the measured photon noise NEP as a function of source power for a system such as ours is given here.

The number of photons per second per unit spectral bandwidth per spatial mode, propagating through free space from a blackbody radiation source is equivalent to the number of photons per standing wave mode in an isothermal cavity of temperature T. Thus, the mean and variance of the photon occupancy number n_i within the i^{th} free space mode can be estimated from the Bose-Einstein distribution derived from statistical mechanics. The mean is

$$\bar{n}_i = \frac{1}{\exp\left(h\nu/k_B T\right) - 1} \tag{B.8}$$

and the variance is

$$\sigma_{n_i}^2 = \bar{n}_i(1 + \bar{n}_i) . \tag{B.9}$$

In the limit $h\nu/k_B T \gg 1$, the variance reduces to $\bar{n}_i$ and the distribution is Poissonian. This is the well known photon shot noise associated with the arrival of independent discrete particles of light. In the opposing limit, $h\nu/k_B T \ll 1$, the variance reduces to $\bar{n}_i^2$ and is interpreted, in a quantum sense, as noise due to the correlated arrival of photons that are coherently 'bunched' together within individual modes. The latter effect can be explained classically by considering the interference of electromagnetic waves and, as such, this variance is commonly known as wave noise.

Spatial, spectral and polarisation modes must be considered in the beam describing the paths from source to detector. The number of spatial modes is estimated by the ratio of the optical throughput or étendue, $A\Omega$, to the étendue of coherence, $\lambda^2 = c^2/\nu^2$. The number of spectral modes is the product of the detector integration time and the spectral bandwidth, $\tau\Delta\nu$. The number of polarisation modes m is between 1 and 2. So, the total number of modes available is

$$g = \frac{mA\Omega\nu^2\tau\Delta\nu}{c^2} \,, \tag{B.10}$$

the mean number of photons expected across all modes is

$$n = g\bar{n}_i \,, \tag{B.11}$$

and the overall variance is

$$\sigma_n^2 = g\sigma_{n_i}^2 = n\left(1 + \frac{n}{g}\right)\left(\frac{g-1}{g+1}\right) \,. \tag{B.12}$$

The final term in the last equation comes from the fact that individual photons can switch between occupation states. This term is only significant when g is very small.

The average power P_s radiated by the source into the beam in a time τ can be calculated from the energy in the modes as

$$P_s = \frac{nh\nu}{\tau} = \frac{g\bar{n}_i h\nu}{\tau} = mA\Omega\Delta\nu\frac{h\nu^3}{c^2}\frac{1}{\exp(h\nu/k_B T) - 1} \tag{B.13a}$$

$$= \frac{1}{2}mA\Omega\Delta\nu B_\nu(\nu, T) \quad [\mathrm{W}] \,, \tag{B.13b}$$

where $B_\nu(\nu, T)$ is the common Planck form for blackbody spectral radiance. Similarly, the mean squared noise in the power radiated per unit integration time $\sigma_{P_s}^2$ can be calculated from the variance as

$$\sigma_{P_s}^2 = \frac{\sigma_n^2 h^2\nu^2}{\tau^2} = \left(\frac{nh^2\nu^2}{\tau^2} + \frac{n^2 h^2\nu^2}{g\tau^2}\right)\left(\frac{g-1}{g+1}\right) \tag{B.14a}$$

$$= \left(\frac{h\nu P_s}{\tau} + \frac{c^2 P_s^2}{mA\Omega\nu^2\tau\Delta\nu}\right)\left(\frac{g-1}{g+1}\right) \quad [\mathrm{W}^2] \,. \tag{B.14b}$$

'Electrical' NEP is defined as the signal power required to equal the noise power in an optically lossless system measured in a $1\,\mathrm{Hz}$ detection bandwidth, $\Delta f_s = 1\,\mathrm{Hz}$. The Nyquist-Shannon sampling theorem requires that $\Delta f_s = 1/2\tau$, so

$$NEP_{\gamma, electrical} = \sqrt{\frac{\sigma_{P_s}^2}{\Delta f_s}} \tag{B.15a}$$

$$= \sqrt{\left(2h\nu P_s + \frac{2c^2 P_s^2}{mA\Omega\nu^2\Delta\nu}\right)\left(\frac{g-1}{g+1}\right)} \quad \left[\mathrm{W\,Hz}^{-\frac{1}{2}}\right] \,. \tag{B.15b}$$

In practice, losses in the system must be accounted for. For example, optical power will be attenuated by lenses and filters in the system so that the power input to the detector $P_{in} = \eta_{opt}P_s$, where the optical efficiency $\eta_{opt} < 1$. Additionally, the detector will not convert all of the power incident upon it into useful signal, so the detected power $P_{det} = \eta_{det}P_{in} = \eta_{det}\eta_{opt}P_s$, where the detector efficiency $\eta_{det} < 1$ also. The photon limited 'optical' NEP measured by a detector for a given source power is then

$$NEP_{\gamma,optical} = \frac{1}{\eta_{det}}\sqrt{\frac{\sigma^2_{P_{det}}}{\Delta f_s}} \tag{B.16a}$$

$$= \sqrt{\left(\frac{2h\nu\eta_{opt}P_s}{\eta_{det}} + \frac{2c^2\eta^2_{opt}P^2_s}{mA\Omega\nu^2\Delta\nu}\right)\left(\frac{g-1}{g+1}\right)} \quad \left[\mathrm{W\,Hz^{-\frac{1}{2}}}\right]. \tag{B.16b}$$

Stray light

Stray light is any form of detected but unwanted optical signal. It will contribute to the power absorbed by a detector and thus it will contribute to the photon noise. However, it will not affect the responsivity of the detector, and so it can lead to a reduction in overall sensitivity.

The ghosts described in Chapter 3 are examples of stray light reflections and more often than not these kinds of reflections can be well baffled with proper shielding or with anti-reflection coatings. Stray emissions are more of a challenge to hold back. They can come from any warm body in the optical path, such as lenses, mirrors, filters or the atmosphere and they will always contribute some load onto a detector. Lowering the temperatures of the emissive optical components is the best way to reduce these emissions. SPICA[1] is a next generation space-based infrared observatory and will utilise a 3.5 m primary mirror that is actively cooled to 4.5 K to reduce emissions and enable ultra-sensitive detection at terahertz frequencies.

---

[1]SPICA: `http://sci.esa.int/cosmic-vision/53635-spica/`

Bibliography

[Ade et al., 2006] Ade, P. A. R., Pisano, G., Tucker, C., and Weaver, S. (2006). A review of metal mesh filters. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6275 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*.

[Appleby et al., 1994] Appleby, R., Gleed, D. G., Anderton, R. N., and Lettington, A. H. (1994). Advances in passive millimeter wave imaging. In *Proceedings of the IEEE*, volume 2211, pages 312–317.

[Barry, 2014] Barry, P. (2014). *On the development of SuperSpec. A Fully Integrated On-chip Spectrometer for Far-infrared Astronomy* . PhD thesis, Cardiff University.

[Becker et al., 2010] Becker, D., Beall, J., Cho, H.-M., Duncan, W., Hilton, G., Horansky, R., Irwin, K., Lowell, P., Niemack, M., Paulter, N., Reintsema, C., Schima, F., Schwall, R., Yoon, K. W., Ade, P., Tucker, C., Dicker, S., and Halpern, M. (2010). A 350-GHz high-resolution high-sensitivity passive video imaging system. In Wikner, D. A. and Luukanen, A. R., editors, *Passive Millimeter-Wave Imaging Technology XIII*. SPIE.

[Bersanelli et al., 2010] Bersanelli, M., Mandolesi, N., Butler, R. C., Mennella, A., Villa, F., et al. (2010). Planck pre-launch status: Design and description of the Low Frequency Instrument. A&A, 520:A4.

[Bourrion et al., 2011] Bourrion, O., Bideaud, A., Benoit, A., Cruciani, A., Macias-Perez, J. F., Monfardini, A., Roesch, M., Swenson, L., and Vescovi, C. (2011). Electronics and data acquisition demonstrator for a kinetic inductance camera. *Journal of Instrumentation*, 6(06):P06012.

[Bourrion et al., 2012] Bourrion, O., Vescovi, C., Bouly, J. L., Benoit, A., Calvo, M., Gallin-Martel, L., Macias-Perez, J. F., and Monfardini, A. (2012). NIKEL: Electronics and data acquisition for kilopixels kinetic inductance camera. *Journal of Instrumentation*, 7:7014.

[Chase Research Cryogenics Ltd.,] Chase Research Cryogenics Ltd.

[Day et al., 2003] Day, P. K., LeDuc, H. G., Mazin, B. A., Vayonakis, A., and Zmuidzinas, J. (2003). A broadband superconducting detector suitable for use in large arrays. *Nature*, 425:817–821.

[de Graauw et al., 2010] de Graauw, T., Helmich, F. P., Phillips, T. G., Stutzki, J., Caux, E., Whyborn, N. D., Dieleman, P., Roelfsema, P. R., et al. (2010). The Herschel-Heterodyne Instrument for the Far-Infrared (HIFI). A&A, 518:L6.

[de Visser et al., 2015] de Visser, P. J., Yates, S. J. C., Guruswamy, T., Goldie, D. J., Withington, S., Neto, A., Llombart, N., Baryshev, A. M., Klapwijk, T. M., and Baselmans, J. J. A. (2015). The non-equilibrium response of a superconductor to pair-breaking radiation measured over a broad frequency band. *Applied Physics Letters*, 106(25):252602.

[Dietlein et al., 2006] Dietlein, C., Luukanen, A., Meyer, F., Popovic, Z., and Grossman, E. N. (2006). Phenomenology of Passive Broadband Terahertz Images. In *ESA Workshop on Millimetre-wave Technology and Applications*, , pages 405–409.

[Doyle, 2008] Doyle, S. (2008). *Lumped Element Kinetic Inductance Detectors*. PhD thesis, Cardiff University.

[Doyle et al., 2008] Doyle, S., Mauskopf, P., Naylon, J., Porch, A., and Dunscombe, C. (2008). Lumped Element Kinetic Inductance Detectors. *Journal of Low Temperature Physics*, 151:530–536.

[European Commission, 2010] European Commission (2010). Communication from the Commission to the European Parliament and the Council on the Use of Security Scanners at EU airports, June 2010.

[F. S. Porter, 2002] F. S. Porter, D., editor (2002). *Multiplexable kinetic inductance detectors*. American Institute of Physics.

[Fitzgerald et al., 2005] Fitzgerald, A. J., Cole, B. E., and Taday, P. F. (2005). Nondestructive analysis of tablet coating thicknesses using terahertz pulsed imaging. *Journal of Pharmaceutical Sciences*, 94(1):177–183.

[Fukunari et al., 2014] Fukunari, M., Wongsuryrat, N., Yamaguchi, T., Komurasaki, K., Oda, Y., Ikeda, R., Kajiwara, K., Takahashi, K., and Sakamoto, K. (2014). Demonstration of sub-megawatt power transmission to microwave rocket. In *Infrared, Millimeter, and Terahertz waves (IRMMW-THz), 2014 39th International Conference on*, pages 1–2. IEEE.

[Galitzki et al., 2014] Galitzki, N. et al. (2014). The Next Generation BLAST Experiment. *Journal of Astronomical Instrumentation*, 3:40001.

[Griffin et al., 2010] Griffin, M. J. et al. (2010). The Herschel-SPIRE instrument and its in-flight performance. A&A, 518:L3.

[Grossman et al., 2010] Grossman, E., Dietlein, C., Ala-Laurinaho, J., Leivo, M., Gronberg, L., Gronholm, M., Lappalainen, P., Rautiainen, A., Tamminen, A., and Luukanen, A. (2010). Passive terahertz camera for standoff security screening. *Appl. Opt.*, 49(19):E106–E120.

[Heinz et al., 2012] Heinz, E., May, T., Born, D., Zieger, G., Brömel, A., Anders, S., Zakosarenko, V., Krause, T., Krüger, A., Schulz, M., Bauer, F., and Meyer, H.-G. (2012). Development of passive submillimeter-wave video imaging systems for security applications. In Salmon, N. A. and Jacobs, E. L., editors, *Millimetre Wave and Terahertz Sensors and Technology V*. SPIE.

[Hollister, 2009] Hollister, M. (2009). *SCUBA-2 instrument: an application of large-format superconducting bolometer arrays for submillimetre astronomy*. PhD thesis, The University of Edinburgh.

[Janssen et al., 2014] Janssen, R. M. J., Baselmans, J. J. A., Endo, A., Ferrari, L., Yates, S. J. C., Baryshev, A. M., and Klapwijk, T. M. (2014). Performance of hybrid NbTiN-Al microwave kinetic inductance detectors as direct detectors for sub-millimeter astronomy. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 9153, pages 91530T–91530T–7.

[Klemencic et al., 2015] Klemencic, G. et al. (Submitted 2015). A continuous dry 300 mK cooler for THz sensing applications. *Review of Scientific Instruments*.

[Knab et al., 2007] Knab, J. R., Chen, J.-Y., He, Y., and Markelz, A. G. (2007). Terahertz Measurements of Protein Relaxational Dynamics. *Proceedings of the IEEE*, 95(8):1605–1610.

[Lamarre, 1986] Lamarre, J. M. (1986). Photon noise in photometric instruments at far-infrared and submillimeter wavelengths. Appl. Opt., 25:870–876.

[Lamarre et al., 2010] Lamarre, J.-M., Puget, J.-L., et al. (2010). Planck pre-launch status: The HFI instrument, from specification to actual performance. A&A, 520:A9.

[Leclercq, 2007] Leclercq, S. (2007). Discussion about Noise Equivalent Power and its use for photon noise calculation. . "http://www.iram.fr/ leclercq/Reports/About_NEP_photon_noise.pdf".

[Mauskopf et al., 2014] Mauskopf, P. D., Doyle, S., Barry, P., Rowe, S., Bidead, A., Ade, P. A. R., Tucker, C., Castillo, E., Monfardini, A., Goupy, J., and Calvo, M. (2014). Photon-Noise Limited Performance in Aluminum LEKIDs. *Journal of Low Temperature Physics*, 176:545–552.

[Mazin et al., 2006] Mazin, B. A., Bumble, B., Day, P. K., Eckart, M. E., Golwala, S., Zmuidzinas, J., and Harrison, F. A. (2006). Position sensitive x-ray spectrophotometer using microwave kinetic inductance detectors. *Applied Physics Letters*, 89(22):222507.

[Mazin et al., 2013] Mazin, B. A., Meeker, S. R., Strader, M. J., Szypryt, P., Marsden, D., van Eyken, J. C., Duggan, G. E., Walter, A. B., Ulbricht, G., Johnson, M., Bumble, B., O"Brien, K., and Stoughton, C. (2013). ARCONS: A 2024 Pixel Optical through Near-IR Cryogenic Imaging Spectrophotometer. PASP, 125:1348–1361.

[Monfardini et al., 2014] Monfardini, A. et al. (2014). Latest NIKA Results and the NIKA-2 Project. *Journal of Low Temperature Physics*, 176:787–795.

[Monfardini et al., 2010] Monfardini, A., Swenson, L. J., Bideaud, A., Désert, F. X., Yates, S. J. C., Benoit, A., Baryshev, A. M., Baselmans, J. J. A., Doyle, S., Klein, B., Roesch, M., Tucker, C., Ade, P., Calvo, M., Camus, P., Giordano, C., Guesten, R., Hoffmann, C., Leclercq, S., Mauskopf, P., and Schuster, K. F. (2010). NIKA: A millimeter-wave kinetic inductance camera. A&A, 521:A29.

[Motte et al., 2010] Motte, F., Zavagno, A., Bontemps, S., Schneider, N., Hennemann, M., di Francesco, J., André, P., Saraceno, P., Griffin, M., Marston, A., Ward-Thompson, D., White, G., Minier, V., Men'shchikov, A., Hill, T., Abergel, A., Anderson, L. D., Aussel, H., Balog, Z., Baluteau, J.-P., Bernard, J.-P., Cox, P., Csengeri, T., Deharveng, L., Didelon, P., di Giorgio, A.-M., Hargrave, P., Huang, M., Kirk, J., Leeks, S., Li, J. Z., Martin, P., Molinari, S., Nguyen-Luong, Q., Olofsson, G., Persi, P., Peretto, N., Pezzuto, S., Roussel, H., Russeil, D., Sadavoy, S., Sauvage, M., Sibthorpe, B.,

Spinoglio, L., Testi, L., Teyssier, D., Vavrek, R., Wilson, C. D., and Woodcraft, A. (2010). Initial highlights of the HOBYS key program, the Herschel imaging survey of OB young stellar objects. A&A, 518:L77.

[National Research Council, 1996] National Research Council (1996). *Airline Passenger Security Screening: New Technologies and Implementation Issues*. National Academies Press.

[Oda et al., 2014] Oda, N., Kurashina, S., Miyoshi, M., Doi, K., Ishi, T., Sudou, T., Morimoto, T., Goto, H., and Sasaki, T. (2014). Microbolometer terahertz focal plane array and camera with improved sensitivity at 0.5-0.6 THz. In *2014 39th International Conference on Infrared, Millimeter, and Terahertz waves (IRMMW-THz)*. IEEE.

[Osiander et al., 2003] Osiander, R., Miragliotta, J. A., Jiang, Z., Xu, J., and Zhang, X.-C. (2003). Mine field detection and identification using terahertz spectroscopic imaging. *Proc. SPIE*, 5070:1–6.

[Pilbratt et al., 2010] 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. (2010). Herschel Space Observatory. An ESA facility for far-infrared and submillimetre astronomy. A&A, 518:L1.

[Planck Collaboration et al., 2011] Planck Collaboration et al. (2011). Planck early results. I. The Planck mission. A&A, 536:A1.

[Poglitsch et al., 2010] Poglitsch, A., Waelkens, C., Geis, N., Feuchtgruber, H., Vandenbussche, B., Rodriguez, L., Krause, O., Renotte, E., van Hoof, C., Saraceno, P., Cepa, J., Kerschbaum, F., et al. (2010). The Photodetector Array Camera and Spectrometer (PACS) on the Herschel Space Observatory. A&A, 518:L2.

[Portieri, 2012] Portieri, A. (2012). Vision-Systems Article.

[Roesch et al., 2012] Roesch, M., Benoit, A., Bideaud, A., Boudou, N., Calvo, M., Cruciani, A., Doyle, S., Leduc, H. G., Monfardini, A., Swenson, L., Leclercq, S., Mauskopf, P., Schuster, K. F., and for the NIKA collaboration (2012). Development of Lumped Element Kinetic Inductance Detectors for NIKA. *ArXiv e-prints*.

[Rosenberg et al., 2014] Rosenberg, M., Bladon, G., Russo, P., and Christensen, L. L. (2014). Astronomy in Everyday Life. *Communicating Astronomy with the Public Journal*, 14:30.

[Seco-Martorell et al., 2013] Seco-Martorell, C., López-Domínguez, V., Arauz-Garofalo, G., Redo-Sanchez, A., Palacios, J., and Tejada, J. (2013). Goya's artwork imaging with Terahertz waves. *Opt. Express*, 21(15):17800–17805.

[Shirokoff et al., 2014] Shirokoff, E., Barry, P. S., Bradford, C. M., Chattopadhyay, G., Day, P., Doyle, S., Hailey-Dunsheath, S., Hollister, M. I., Kovács, A., Leduc, H. G., McKenney, C. M., Mauskopf, P., Nguyen, H. T., O'Brient, R., Padin, S., Reck, T. J., Swenson, L. J., Tucker, C. E., and Zmuidzinas, J. (2014). Design and Performance of SuperSpec: An On-Chip, KID-Based, mm-Wavelength Spectrometer. *Journal of Low Temperature Physics*, 176(5):657–662.

[Shumaker et al., 2013] Shumaker, E., Corcos, D., Kaminsky, N., Elad, D., Morf, T., and Klein, B. (2013). Uncooled and passive 0.5-1.5 THz FPA imager: Design and characterization of focal plane array imager based on antenna-coupled MOSFET bolometers. In *2013 IEEE International Conference on Microwaves, Communications, Antennas and Electronic Systems (COMCAS 2013)*. IEEE.

[Siegel, 2002] Siegel, P. H. (2002). Terahertz technology. *IEEE Transactions on Microwave Theory Techniques*, 50:910–928.

[Titova et al., 2013] Titova, L. V., Ayesheshim, A. K., Golubov, A., Fogen, D., Rodriguez-Juarez, R., Hegmann, F. A., and Kovalchuk, O. (2013). Intense THz pulses cause H2AX phosphorylation and activate DNA damage response in human skin tissue. *Biomed. Opt. Express*, 4(4):559–568.

[Tucker and Ade, 2006] Tucker, C. E. and Ade, P. A. R. (2006). Thermal filtering for large aperture cryogenic detector arrays. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6275 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 0.

[Woodward et al., 2003] Woodward, R., Wallace, V., Arnone, D., Linfield, E., and Pepper, M. (2003). Terahertz pulsed imaging of skin cancer in the time and frequency domain. *Journal of Biological Physics*, 29(2-3):257–261.

[Yates et al., 2011] Yates, S. J. C., Baselmans, J. J. A., Endo, A., Janssen, R. M. J., Ferrari, L., Diener, P., and Baryshev, A. M. (2011). Photon noise limited radiation detection with lens-antenna coupled microwave kinetic inductance detectors. *Applied Physics Letters*, 99(7):073505.

[Zmuidzinas, 2003] Zmuidzinas, J. (2003). Thermal noise and correlations in photon detection. Appl. Opt., 42:4989–5008.

[Zmuidzinas et al., 2002] Zmuidzinas, J., Mazin, B., Vayonakis, A., Day, P., and Leduc, H. G. (2002). Multiplexable Kinetic Inductance Detectors Photoconductor Analogy.

[Zmuidzinas and Richards, 2004] Zmuidzinas, J. and Richards, P. L. (2004). Superconducting detectors and mixers for millimeter and submillimeter astrophysics. *Proceedings of the IEEE*, 92(10):1597–1616.